



## A Flow-based Truncated Denoising Diffusion Model for Super-resolution Magnetic Resonance Spectroscopic Imaging

Siyuan Dong<sup>a,\*</sup>, Zhuotong Cai<sup>b,i,\*</sup>, Gilbert Hangel<sup>c,d</sup>, Wolfgang Bogner<sup>c</sup>, Georg Widhalm<sup>d</sup>, Yaqing Huang<sup>e</sup>, Qinghao Liang<sup>b</sup>, Chenyu You<sup>a</sup>, Chathura Kumaragama<sup>f</sup>, Robert K. Fulbright<sup>f</sup>, Amit Mahajan<sup>f</sup>, Amin Karbasi<sup>a,g,h</sup>, John A. Onofrey<sup>f</sup>, Robin A. de Graaf<sup>b,f</sup>, James S. Duncan<sup>a,b,f,\*</sup>

<sup>a</sup>Department of Electrical Engineering, Yale University, New Haven, CT, USA

<sup>b</sup>Department of Biomedical Engineering, Yale University, New Haven, CT, USA

<sup>c</sup>Department of Biomedical Imaging and Image-Guided Therapy, Highfield MR Center, Medical University of Vienna, Vienna, Austria

<sup>d</sup>Department of Neurosurgery, Medical University of Vienna, Vienna, Austria

<sup>e</sup>Department of Pathology, Yale University, New Haven, CT, USA

<sup>f</sup>Department of Radiology and Biomedical Imaging, Yale University, New Haven, CT, USA

<sup>g</sup>Department of Computer Science, Yale University, New Haven, CT, USA

<sup>h</sup>Department of Statistics and Data Science, Yale University, New Haven, CT, USA

<sup>i</sup>Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an, China

### ARTICLE INFO

#### Article history:

Received xxx

Received in final form xxx

Accepted xxx

Available online xxx

Communicated by xxx

2000 MSC: 41A05, 41A10, 65D05, 65D17

Keywords: MR Spectroscopic Imaging, Super-resolution, Diffusion Models, Normalizing Flow

### ABSTRACT

Magnetic Resonance Spectroscopic Imaging (MRSI) is a non-invasive imaging technique for studying metabolism and has become a crucial tool for understanding neurological diseases, cancers and diabetes. High spatial resolution MRSI is needed to characterize lesions, but in practice MRSI is acquired at low resolution due to time and sensitivity restrictions caused by the low metabolite concentrations. Therefore, there is an imperative need for a post-processing approach to generate high-resolution MRSI from low-resolution data that can be acquired fast and with high sensitivity. Deep learning-based super-resolution methods provided promising results for improving the spatial resolution of MRSI, but they still have limited capability to generate accurate and high-quality images. Recently, diffusion models have demonstrated superior learning capability than other generative models in various tasks, but sampling from diffusion models requires iterating through a large number of diffusion steps, which is time-consuming. This work introduces a Flow-based Truncated Denoising Diffusion Model (FTDDM) for super-resolution MRSI, which shortens the diffusion process by truncating the diffusion chain, and the truncated steps are estimated using a normalizing flow-based network. The network is conditioned on upscaling factors to enable multi-scale super-resolution. To train and evaluate the deep learning models, we developed a <sup>1</sup>H-MRSI dataset acquired from 25 high-grade glioma patients. We demonstrate that FTDDM outperforms existing generative models while speeding up the sampling process by over 9-fold compared to the baseline diffusion model. Neuroradiologists' evaluations confirmed the clinical advantages of our method, which also supports uncertainty estimation and sharpness adjustment, extending its potential clinical applications.

© 2024 Elsevier B. V. All rights reserved.

## 1. Introduction

Magnetic Resonance Spectroscopic Imaging (MRSI) is a non-invasive imaging technique for studying metabolism within the body. Since the study of metabolism plays a key role in understanding a range of diseases including neurodegenerative diseases, cancers and diabetes, MRSI has become a valuable clinical and preclinical tool. MRSI is feasible on any nucleus possessing nuclear spin, and the most commonly used nucleus for in vivo MRSI is proton ( $^1\text{H}$ ) due to its high nuclear magnetic resonance (NMR) sensitivity and abundance in metabolites (De Graaf, 2019).

High spatial resolution MRSI is needed to characterize small lesions and intra-lesion heterogeneity. However, in practice, MRSI is acquired at low resolution due to time and sensitivity restrictions caused by the low metabolite concentrations. To achieve a satisfactory signal-to-noise ratio (SNR) with clinical practice, which often employs 3T scanners and mandates acceptable scan times, the spatial resolution of  $^1\text{H}$ -MRSI is often restricted to approximately  $1\times 1\times 1\text{ cm}^3$ . While advancements in hardware and acceleration techniques have been noteworthy (Bogner *et al.*, 2021), achieving high-resolution MRSI through stronger magnetic fields (e.g., 7T) or extended scan times remains largely impractical for routine clinical use, especially on established 3T MRI systems. Therefore, developing a post-processing approach to generate high-resolution MRSI from low-resolution scans will greatly benefit its clinical applications.

Traditional post-processing approaches for super-resolution MRSI primarily employ model-based regularization, which incorporates high-resolution anatomical information from corresponding MRI scans and enforces spatial smoothness using total variation distance (Lam and Liang, 2014; Jain *et al.*, 2017; Kasten *et al.*, 2016; Hangel *et al.*, 2019). However, there are a few problems with these regularization-based methods: (1) While a correlation exists between MRI intensity and MRSI metabolite concentrations, MRIs lack metabolic information. For this reason, over-reliance on MRI-based regularization might bias the resulting high-resolution MRSI towards the MRI (Dong *et al.*, 2023). (2) Handcrafted regularizations, such as total variation distance, do not truthfully reflect the intrinsic characteristics of high-resolution MRSI, potentially leading to over-smoothed results. (3) Each new acquisition needs to undergo an optimization process, which is time-consuming.

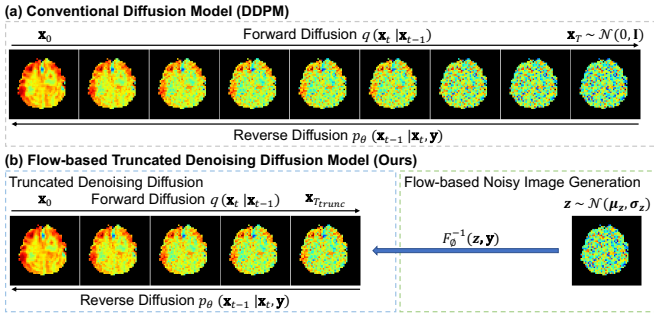
Deep learning has shown great potential in super-resolution tasks for medical imaging (Li *et al.*, 2021) and can effectively circumvent the limitations of conventional techniques. The first deep learning method for super-resolution MRSI (Iqbal *et al.*, 2019) utilized a Dense UNet to map low-resolution MRSI metabolic maps to higher resolution ones by incorporating anatomical information from T1 MRI. However, due to the absence of a public MRSI training dataset, these metabolic maps were synthesized from MRI, which has a fundamental issue that

the simulated MRSI may not authentically capture all metabolite characteristics. To address this issue, Dong *et al.* (2021) introduced an in vivo  $^1\text{H}$ -MRSI dataset acquired from high-grade glioma patients to train a super-resolution MRSI network. This work aimed to restore high-resolution details and enhance visual quality by applying adversarial loss, which was also explored in subsequent research (Dong *et al.*, 2022b; Li *et al.*, 2022b). Imposing adversarial loss (Wang *et al.*, 2020; Chen *et al.*, 2024a) via Generative Adversarial Networks (GAN) (Goodfellow *et al.*, 2020) is often subject to training instability and mode-collapse. To more accurately learn the distribution of high-resolution MRSI images, Dong *et al.* (2022c) proposed a normalizing flow-based network. This framework learns the target distribution via likelihood maximization, which makes the training more stable and interpretable, outperforming GAN-based approaches. Nonetheless, the flow-based methods must be implemented by invertible neural networks (Ardizzone *et al.*, 2018; Lugmayr *et al.*, 2020; Liang *et al.*, 2021; Dong *et al.*, 2022a; Chen *et al.*, 2024b), which has constrained architecture and limit the learning capability.

Recently, diffusion models (Sohl-Dickstein *et al.*, 2015; Ho *et al.*, 2020; Song and Ermon, 2019; Song *et al.*, 2020) have demonstrated superior learning capability than other generative models in a range of tasks (Dhariwal and Nichol, 2021; Yang *et al.*, 2022). Diffusion models can be used to learn conditional image distributions for image-to-image translation tasks, such as image super-resolution (Saharia *et al.*, 2022; Li *et al.*, 2022a) and MRI reconstruction (Chung and Ye, 2022; Peng *et al.*, 2022; Kazerouni *et al.*, 2022). A popular subclass of diffusion models is called Denoising Diffusion Probabilistic Model (DDPM) (Ho *et al.*, 2020; Nichol and Dhariwal, 2021), which involves a forward diffusion process and a reverse diffusion process. The forward diffusion process gradually perturbs the target images with Gaussian noise through a large number of diffusion steps until a pure Gaussian noise is obtained. During the reverse diffusion process, a neural network retraces this process, training to revert the Gaussian noise back to the original image by gradually eliminating the noise. However, an inherent problem with diffusion models is that the sampling process, or the reverse diffusion process, requires iterating through a large number of diffusion steps, typically in the thousands, which is time-consuming. This lengthy process is due to the need to inject only a small amount of noise in each step to ensure that the forward and reverse processes have approximately the same functional form (Ho *et al.*, 2020). One simple way to expedite the sampling process is to use strided steps (Nichol and Dhariwal, 2021), but this may compromise the quality of the samples. Recently, Zheng *et al.* (2022) proposed a Truncated Diffusion Probabilistic Model (TDPM) to shorten the diffusion process by truncating the diffusion chain, which achieved a similar level of performance as DDPM while achieving a significant speedup.

In this paper, inspired by TDPM, we propose a Flow-based Truncated Denoising Diffusion Model (FTDDM) for super-resolution MRSI. This model shortens the diffusion process of DDPM by truncating the diffusion chain, and the truncated steps are estimated using a normalizing flow-based network, as depicted in Fig. 1. In contrast to the TDPM, which uses

\*Corresponding author at: 300 Cedar Street, New Haven, CT, 06519  
*e-mail*: s.dong@yale.edu (Siyuan Dong), zhuotong.cai@yale.edu (Zhuotong Cai), james.duncan@yale.edu (James S. Duncan)



**Fig. 1.** A comparison between (a) the conventional diffusion model DDPM and (b) our method FTDDM.  $\mathbf{x}_0$  is the noiseless high-resolution MRSI metabolic map. The forward diffusion process gradually adds Gaussian noise to  $\mathbf{x}_0$ . The noise is only added to the region of interest, as defined by a quality-filtering mask, to avoid the necessity of suppressing noise from the background during the reverse diffusion process. The reverse diffusion process uses a denoising network with parameters  $\theta$  to retrace the forward diffusion process, provided with any condition images  $\mathbf{y}$ .  $F_\phi^{-1}$  is the inverse of a normalizing flow-based network used to bridge the gap between the pure Gaussian noise  $\mathbf{z}$  and the noisy image at the truncation point  $\mathbf{x}_{T_{trunc}}$ .

an adversarial network, our approach leverages a normalizing flow-based network to bridge the gap between pure Gaussian noise and the noisy image at the truncation point. This is because flow-based networks were trained with likelihood maximization, which is more stable than adversarial training, and demonstrated superior performance over GAN in the recent MRSI super-resolution work (Dong *et al.*, 2022c). In addition, we condition our network on the upscaling factor using Conditional Instance Normalization. This promotes multi-scale super-resolution, eliminating the need to train a separate network for each upscaling factor. Moreover, we incorporate a temperature parameter to flexibly adjust the sharpness of the output images and investigate the trade-off between image visual quality and fidelity. To train and evaluate the proposed method, we develop a  $^1\text{H}$ -MRSI dataset acquired from 25 high-grade glioma patients. Experimental results show that FTDDM outperforms existing generative models, i.e. GAN and flow, while accelerating the sampling process by over 9-fold relative to DDPM. Neuroradiologists were invited to evaluate our method from a range of clinical perspectives, which confirms the clinical advantages of our super-resolved MRSI.

Our most significant contributions are: (1) development of a novel Flow-based Truncated Denoising Diffusion Model (FTDDM) for performant and sampling-efficient super-resolution MRSI; (2) adoption of network conditioning to facilitate multi-scale super-resolution; (3) introduction of a temperature parameter for sharpness adjustment; (4) development of an in vivo  $^1\text{H}$ -MRSI dataset acquired from glioma patients for training deep learning networks; and (5) Integration of neuroradiologists' ratings to assess the quality of super-resolution metabolic maps relative to the ground truth from clinical perspectives.

## 2. Material and methods

We first describe the acquisition and pre-processing of our dataset in Section 2.1. We then present our Truncated Denoising Diffusion model and the corresponding loss function in Sec-

tion 2.2, followed by how multi-scale super-resolution is incorporated into the network in Section 2.3. Section 2.4 describes the design and training of the flow-based model for generating noisy images at the truncation point from Gaussian noise. Finally, the overall algorithm to sample from FTDDM is summarized in Section 2.5.

### 2.1. Dataset

3D high-resolution  $^1\text{H}$ -MRSI, T1-weighted (MP2RAGE) and Fluid Attenuated Inversion Recovery (FLAIR) scans from 25 high-grade glioma patients were acquired on a 7T whole body MRI (Magnetom, Siemens Healthcare) equipped with a 32-channel receive coil array (Nova Medical). IRB approval and informed consent from all participants were obtained. MRSI sequence used spatial-spectral encoding to acquire 3D metabolic maps with a  $64 \times 64 \times 39$  measurement matrix and  $3.4 \times 3.4 \times 3.4$  mm<sup>3</sup> nominal resolution in 15 minutes (Hangel *et al.*, 2020; Hingerl *et al.*, 2020; Dong *et al.*, 2021). This resolution is notably high for MRSI due to the low SNR of metabolite signals. Additional parameters included an acquisition delay of 1.3 ms, a repetition time (TR) of 450 ms and a spectral bandwidth of 2778 Hz.

T1 and FLAIR MRIs were skull-stripped and co-registered to MRSI using FSL v5.0 (Smith *et al.*, 2004). The processed MRSI spectra, as described in Hangel *et al.* (2020); Hingerl *et al.* (2020), were quantified using LCMoDel v6.3-1 (Provencher, 2014) to obtain 3D metabolic maps. For each of these maps, voxels with insufficient quality (SNR < 2.5, or FWHM > 0.15 ppm due to  $B_0$  field inhomogeneity, or with severe extracranial lipid contamination proximal to the skull) were filtered out using a quality-filtering mask. In this work, we focus on 7 metabolites considered as major markers of oncometabolism and proliferation (Hangel *et al.*, 2020): total creatine (tCr), total choline (tCho), N-acetylaspartate (NAA), glutamate (Glu), inositol (Ins), glutamine (Gln), and glycine (Gly).

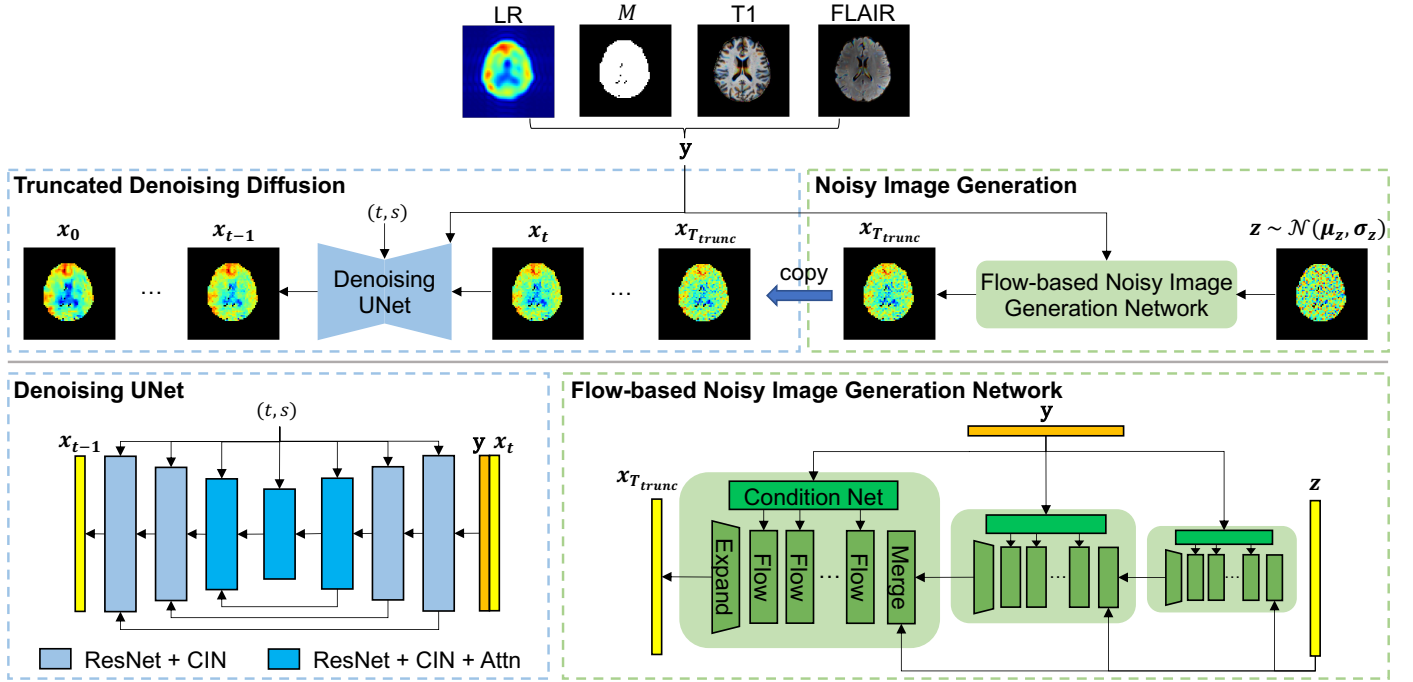
From each 3D MRSI scan of the 25 patients, we selected 9~18 axial slices with clear spatial characteristics of brain. Each slice provided 7 metabolites, resulting in a total of 2275 2D metabolic maps. These  $64 \times 64$  metabolic maps are regarded as high-resolution ground truth, from which we derived low-resolution images using k-space truncation. The size of the truncation window depends on the upscaling factor, which will be discussed later.

### 2.2. Truncated Denoising Diffusion

The conventional diffusion model, DDPM (Ho *et al.*, 2020), adopts a forward and a reverse diffusion process to establish a mapping between a target distribution and a simple Gaussian distribution, as illustrated in Fig.1(a). Starting with a high-resolution metabolic map  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ , the forward diffusion process gradually adds Gaussian noise to  $\mathbf{x}_0$  over  $T$  diffusion steps until a pure Gaussian noise  $\mathbf{x}_T$  is achieved

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}), \quad (1)$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t | \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$$



**Fig. 2. Overview of the proposed method.** The truncated denoising diffusion employs a Denoising UNet to iteratively estimate and remove noise from  $\mathbf{x}_{T_{trunc}}$ , resulting in a noiseless high-resolution MRSI metabolic map  $\mathbf{x}_0$ . The Denoising UNet also takes the condition  $\mathbf{y}$ , which is a concatenation of the low-resolution (LR) metabolic map, a quality-filtering mask  $M$ , T1 MRI and FLAIR MRI, i.e.  $\mathbf{y} = \{\text{LR}, M, \text{T1}, \text{FLAIR}\}$ . The Denoising UNet consists of Residual Network (ResNet) blocks and Conditional Instance Normalization (CIN). The CIN embeds timestep  $t$  and upscaling factor  $s$  into the network. The blocks in the middle have multi-head attention (Attn) modules, following Nichol and Dhariwal (2021).  $\mathbf{x}_{T_{trunc}}$  is generated from the Gaussian noise  $\mathbf{z}$  via the flow-based noisy image generation network, which comprises a series of flow layers across multiple dimensions, in line with Dong et al. (2022c). Each flow layer contains conditional affine coupling, affine injector, invertible  $1 \times 1$  convolution and activation normalization (Lugmayr et al., 2020). The condition images  $\mathbf{y}$  are infused into the flow layers through Condition Networks (Condition Net), which consist of convolution layers and LeakyReLU.

where  $\beta_t \in (0, 1)$  is the predefined noise variance at each timestep  $t$ . Based on the property of Gaussian distribution, we can obtain an analytical form of the forward process for an arbitrary timestep  $t$

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t | \sqrt{\gamma_t} \mathbf{x}_0, (1 - \gamma_t) \mathbf{I}) \quad (2)$$

where  $\gamma_t = \prod_{i=1}^t (1 - \beta_i)$ . The reverse diffusion process counteracts the forward diffusion process by gradually eliminating noise from  $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$  to ultimately approximate  $\mathbf{x}_0$ , given any condition image  $\mathbf{y}$

$$p_\theta(\mathbf{x}_{0:T-1} | \mathbf{x}_T, \mathbf{y}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}) \quad (3)$$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}) = \mathcal{N}(\mathbf{x}_{t-1} | \mu_\theta(\mathbf{y}, \mathbf{x}_t, t), \sigma_t^2 \mathbf{I}),$$

where the mean  $\mu_\theta(\mathbf{y}, \mathbf{x}_t, t)$  is learnt with a Denoising UNet with model parameters  $\theta$  (Ho et al., 2020), and the variance is fixed to  $\sigma_t^2 = \frac{1-\gamma_{t-1}}{1-\gamma_t} \beta_t$  as suggested by Peng et al. (2022). This reverse process can be regarded as a Markov chain with learned Gaussian transitions with an initial state of  $p(\mathbf{x}_T) = \mathcal{N}(0, \mathbf{I})$ . However, to approximate the reverse process with this network parameterized Gaussian distribution, the value of  $\beta_t$  must be sufficiently constrained to ensure that the forward and reverse processes have approximately the same functional form (Sohl-Dickstein et al., 2015; Ho et al., 2020). As a result,  $T$  should be selected large enough (typically thousands of diffusion steps)

to make  $p(\mathbf{x}_T)$  virtually identical to standard Gaussian noise  $\mathcal{N}(0, \mathbf{I})$ . This requirement makes the reverse diffusion process, or the sampling process, extremely slow due to the thousands of network evaluations. To overcome this issue, taking inspiration from Zheng et al. (2022), we shorten the sampling process by truncating the diffusion chain at  $T_{trunc}$  and initiating the reverse denoising from a noisy image  $\mathbf{x}_{T_{trunc}}$  rather than pure Gaussian noise, as depicted in Fig.1(b). The truncated denoising diffusion process is now characterized as

$$p_\theta(\mathbf{x}_{0:T_{trunc}-1} | \mathbf{x}_{T_{trunc}}, \mathbf{y}) = p(\mathbf{x}_{T_{trunc}}) \prod_{t=1}^{T_{trunc}} p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}) \quad (4)$$

The distribution of noisy images at the truncation point  $p(\mathbf{x}_{T_{trunc}})$  is estimated from Gaussian noise using a normalizing flow-based network, which will be detailed in Section 2.4. To fulfill Equation 4, a Denoising UNet  $\epsilon_\theta$  with parameters  $\theta$  is trained to estimate the noise  $\epsilon$  in  $\mathbf{x}_t$  at a given timestep  $t$

$$\mathbb{E}_{(\mathbf{x}_0, \mathbf{y})} \mathbb{E}_{\epsilon, t, s} \|\epsilon - \epsilon_\theta(\mathbf{y}, \mathbf{x}_t, t, s)\|_2^2 \quad (5)$$

where  $\mathbf{x}_t = \sqrt{\gamma_t} \mathbf{x}_0 + (1 - \gamma_t) \epsilon$  and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  based on Equation 2. Timestep  $t$  is uniformly sampled from  $\mathcal{U}(1, T_{trunc})$ .  $s$  represents the upscaling factor, which will be detailed in Section 2.3.

The detailed architecture of the Denoising UNet is shown in Fig.2, adhering to the design presented by Nichol and Dhariwal (2021). Specifically,  $\mathbf{x}_0$  is a single noiseless high-resolution

MRSI metabolic map. The condition  $\mathbf{y}$  is a concatenation of four components: low-resolution (LR) metabolic map, quality-filtering mask  $M$ , T1 MRI and FLAIR MRI, i.e.  $\mathbf{y} = \{\text{LR}, M, \text{T1}, \text{FLAIR}\}$ . The inclusion of the quality-filtering mask enables the network to differentiate between artificial features (voxels with zero intensity) generated by quality filtering and genuine metabolic characteristics. T1 and FLAIR MRIs may offer high-resolution anatomical information beneficial for the super-resolution process (Dong *et al.*, 2021).

### 2.3. Multi-scale Super-resolution

Upscaling factor is defined as the factor that resolution is improved on each dimension. For example, improving an image's resolution from  $16 \times 16$  to  $64 \times 64$  gives an upscaling factor of 4. Depending on the acquisition protocol and the specific MRSI application, experimentally acquired resolutions can vary. Therefore, the upscaling factor often remains unknown prior to training the super-resolution network (Dong *et al.*, 2022b). Most of the previous MRSI super-resolution networks were designed single-scale (Iqbal *et al.*, 2019; Dong *et al.*, 2021, 2022c), which requires training an independent network for each upscaling factor. This is time-consuming and memory-inefficient. To enable multi-scale super-resolution, we condition the Denoising UNet with the upscaling factor by injecting  $s$  to the network. Specifically, the feature maps after the convolution layers in the residual blocks are modulated with Conditional Instance Normalization (Huang and Belongie, 2017)

$$f' = \hat{\sigma}(t, s) \left( \frac{f - \mu(f)}{\sigma(f)} \right) + \hat{\mu}(t, s) \quad (6)$$

where  $\mu(f) \in \mathbb{R}^C$  and  $\sigma(f) \in \mathbb{R}^C$  are the channel-wise mean and standard deviation of the feature map  $f$  with channel number  $C$ . The modulated mean  $\hat{\mu}$  and standard deviation  $\hat{\sigma}$  are learned from  $t$  and  $s$  using a multilayer perceptron (4 fully-connected layers with Sigmoid Linear Units (Elfwing *et al.*, 2018), or SiLUs, activation functions in between). This approach ensures the network uniquely addresses each upscaling factor.

### 2.4. Flow-based Noisy Image Generation

Normalizing flow constructs an unknown target distribution from a simple base distribution using a flow of invertible transformations (Rezende and Mohamed, 2015). Given that  $\mathbf{x}_{T_{trunc}}$  represents a sample from the distribution of noisy images at the truncation point and  $\mathbf{y}$  is the corresponding set of condition images, we use an invertible neural network  $F_\phi$  to transform  $\mathbf{x}_{T_{trunc}}$  into a latent Gaussian variable  $\mathbf{z}$ , i.e.  $\mathbf{z} = F_\phi(\mathbf{x}_{T_{trunc}}, \mathbf{y})$ . Once this transformation is learned, the network can generate samples from the target distribution  $p(\mathbf{x}_{T_{trunc}})$  by sampling  $\mathbf{z}$  from its distribution  $p_{\mathbf{z}}(\mathbf{z})$  and inversely passing through the network  $\mathbf{x}_{T_{trunc}} = F_\phi^{-1}(\mathbf{z}, \mathbf{y})$ . Based on the change of variable formula, the target distribution can be expressed as

$$p_{\mathbf{x}_{T_{trunc}}|\mathbf{y}}(\mathbf{x}_{T_{trunc}}|\mathbf{y}) = p_{\mathbf{z}}(\mathbf{z}) \left| \det \left( \frac{\partial F_\phi(\mathbf{x}_{T_{trunc}}, \mathbf{y})}{\partial \mathbf{x}_{T_{trunc}}} \right) \right| \quad (7)$$

with  $\frac{\partial F_\phi}{\partial \mathbf{x}}$  being the Jacobian matrix (Winkler *et al.*, 2019). This expression can be reformulated into a negative log-likelihood (NLL) loss for network training

$$\mathcal{L}_{NLL} = -\log p_{\mathbf{z}}(\mathbf{z}) - \log \left| \det \left( \frac{\partial (F_\phi(\mathbf{x}_{T_{trunc}}, \mathbf{y}))}{\partial \mathbf{x}_{T_{trunc}}} \right) \right| \quad (8)$$

where the first term equals  $-\frac{1}{2} \left( \left\| \frac{\mathbf{z} - \boldsymbol{\mu}_{\mathbf{z}}}{\boldsymbol{\sigma}_{\mathbf{z}}} \right\|_2^2 + \left\| \log(2\pi\boldsymbol{\sigma}_{\mathbf{z}}^2) \right\|_1 \right)$  because  $\mathbf{z}$  follows a Gaussian distribution with learned mean  $\boldsymbol{\mu}_{\mathbf{z}}$  and standard deviation  $\boldsymbol{\sigma}_{\mathbf{z}}$ , i.e.  $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}}, \boldsymbol{\sigma}_{\mathbf{z}})$ .

Observing that the noisy image  $\mathbf{x}_{T_{trunc}}$  should center around  $\mathbf{x}_0$ , we add a guide loss in the sampling direction. This encourages the network to generate an estimated noisy image  $\hat{\mathbf{x}}_{T_{trunc}}$  that centers around  $\mathbf{x}_0$  (Liang *et al.*, 2021; Dong *et al.*, 2022c). The guide loss is composed of two fidelity-oriented loss functions, pixel loss and structural loss:

$$\mathcal{L}_{guide} = (1 - \alpha) \mathcal{L}_{pixel}(\mathbf{x}_0, \hat{\mathbf{x}}_{T_{trunc}}^{\tau_f=0}) + \alpha \mathcal{L}_{structural}(\mathbf{x}_0, \hat{\mathbf{x}}_{T_{trunc}}^{\tau_f=0}). \quad (9)$$

$\mathcal{L}_{pixel}$  is the pixel-wise L1-distance between two images, and  $\mathcal{L}_{structural}$  measures the Multiscale Structural Similarity (MSS-SIM) between two images (Dong *et al.*, 2021; Wang *et al.*, 2003). The scalar  $\alpha \in (0, 1)$  defines the weight between the two losses.  $\tau_f \in (0, 1)$  acts as a temperature parameter that controls the variance of the sample:  $\hat{\mathbf{x}}_{T_{trunc}}^{\tau_f} = F_\phi^{-1}(\mathbf{z}, \mathbf{y})$  with  $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}}, \tau_f \boldsymbol{\sigma}_{\mathbf{z}})$ . In the guide loss, we set  $\tau_f = 0$  to encourage that the center of the distribution of  $\hat{\mathbf{x}}_{T_{trunc}}$  locates around  $\mathbf{x}_0$ .

The overall loss for the flow-based noisy image generation network is a weighted combination of the NLL loss (Equation 8) and the guide loss (Equation 9)

$$\mathcal{L}_{flow} = \mathcal{L}_{NLL} + \lambda \mathcal{L}_{guide} \quad (10)$$

### 2.5. Sampling from FTDDM

To generate noisy images at the truncation point, we use the trained flow-based noisy image generation network  $F_\phi$ :

$$\hat{\mathbf{x}}_{T_{trunc}} = F_\phi^{-1}(\mathbf{z}, \mathbf{y}), \quad \mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}}, \tau_f \boldsymbol{\sigma}_{\mathbf{z}}). \quad (11)$$

To generate clean high-resolution MRSI images  $\mathbf{x}_0$  from the estimated  $\mathbf{x}_{T_{trunc}}$  through the reverse diffusion process, we use the trained Denoising UNet  $\epsilon_\theta$  and apply the sampling formula derived in the original DDPM paper (Ho *et al.*, 2020) iteratively:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \gamma_t}} \epsilon_\theta(\mathbf{y}, \mathbf{x}_t, t, s) \right) + \sigma_t \mathbf{z}_t \quad (12)$$

where  $t = T_{trunc}, \dots, 1$ , and  $\mathbf{z}_t$  is a latent Gaussian variable for the diffusion model at timestep  $t$ . Observing that the visual sharpness of the generated images is correlated to the variance of the latent Gaussian variable in the flow-based networks (Dong *et al.*, 2022c; Lugmayr *et al.*, 2020), we analogously introduce a temperature parameter  $\tau_d \in (0, 1)$  into our diffusion model, such that  $\mathbf{z}_t \sim \mathcal{N}(0, \tau_d \mathbf{I})$ . In this way, the sharpness of the generated images can be flexibly adjusted by  $\tau_d$ . The overall sampling process of FTDDM is described in Algorithm 1.

**Algorithm 1** Sampling from FTDDM

---

```

1:  $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_z, \tau_f \boldsymbol{\sigma}_z)$ 
2:  $\hat{\mathbf{x}}_{T_{trunc}} = F_{\phi}^{-1}(\mathbf{z}, \mathbf{y})$ 
3:  $\mathbf{x}_{T_{trunc}} = \hat{\mathbf{x}}_{T_{trunc}}$ 
4: for  $t = T_{trunc}, \dots, 1$  do
5:    $\mathbf{z}_t \sim \mathcal{N}(0, \tau_d \mathbf{I})$  if  $t > 1$ , else  $\mathbf{z}_t = 0$ 
6:    $\mathbf{x}_{t-1} = \frac{1}{\sqrt{1-\beta_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} \epsilon_{\theta}(\mathbf{y}, \mathbf{x}_t, t, s) \right) + \sigma_t \mathbf{z}_t$ 
7: end for
8: return  $\mathbf{x}_0$ 

```

---

**3. Experiments and Results***3.1. Implementation Details*

We implemented 5-fold cross-validation on our dataset comprising 25 patients. In each fold, 15 patients were used for training, 5 for validation and another 5 for testing. Across these folds, the average counts of 2D metabolic maps used for training, validation and testing are 1365, 455 and 455, respectively. The training dataset was substantially augmented using random rotation, flipping and translation during training. The validation dataset was used to monitor convergence and save the model with the lowest validation loss. The low resolutions were uniformly sampled during training from 13 values:  $\text{LR} \in \{8 \times 8, 10 \times 10, 12 \times 12, \dots, 30 \times 30, 32 \times 32\}$ , which corresponds to up-scaling factors of  $s \in \{8.0, 6.4, 5.3, \dots, 2.1, 2.0\}$ . T1 and FLAIR images were down-sampled to a consistent resolution of  $64 \times 64$  with the high-resolution metabolic maps and were normalized between  $[0, 1]$  before being input to the network.

The Denoising UNet, shown in Fig.2 follows the same architecture as proposed in Nichol and Dhariwal (2021). It has four stages of downsampling and four stages of upsampling, each with three convolutional residual blocks. The number of channels used at each stage is 64, 128, 192 and 256, respectively. The upscaling factor  $s$  and sinusoidal embedding of timestep  $t$  are passed to the residual blocks via Conditional Instance Normalization. Multi-head attention blocks with four attention heads are used between the residual blocks at  $16 \times 16$  and  $8 \times 8$  feature map resolution levels.

The diffusion models employed a total of  $T=1000$  diffusion steps, as recommended by Ho et al. (2020). Our FTDDM truncated the diffusion chain at  $T_{trunc}=100$ , equivalent to roughly 10-fold acceleration.  $\beta_t$  follows a cosine noise schedule as proposed by Nichol and Dhariwal (2021). The temperature parameters for sampling from the diffusion model and the flow-based network were set to  $\tau_d=0.9$  and  $\tau_f=0.9$ , respectively. The weighting parameters for the loss function are  $\alpha=0.84$  (Zhao et al., 2016) and  $\lambda=10$  (Dong et al., 2022c). The proposed model was trained with Adam optimizer (Kingma and Ba, 2014), 100,000 iterations and batch size of 8. The initial learning rate was set to  $1 \times 10^{-4}$ , followed by a decay per iteration. The methods for comparison were implemented based on their original publications and also trained in a multi-scale manner. All models were implemented with PyTorch and run on NVIDIA GTX 1080 and V100 GPUs.

*3.2. Experimental Studies*

We conducted both quantitative and qualitative comparisons of our proposed method, FTDDM, against other state-of-the-art deep learning models. MUNet-FS-cWGAN (Dong et al., 2022b) is a Multi-encoder UNet (MUNet) architecture trained with conditional Wasserstein GAN (cWGAN). It uses a Filter-Scaling (FS) strategy for multi-scale super-resolution. The Flow Enhancer (Dong et al., 2022c) is a normalizing flow-based invertible network designed to enhance the visual quality of the blurry super-resolution images given by MUNet. DDPM (Ho et al., 2020) is the conventional diffusion model that uses the full diffusion chain  $T=1000$  to iteratively transform Gaussian noise into high-resolution MRSI images (Fig.1(a)), given the condition images  $\mathbf{y}$ . As a baseline method to reduce sampling time, we implemented the simple striding, or respacing, method to reduce DDPM sampling steps from  $T = 1000$  to  $T_{respace}=100$  by evenly skipping the diffusion steps (Nichol and Dhariwal, 2021). DPM-Solver++ is a high-order ODE solver designed to generate high-quality samples from diffusion models within a small number of steps (Lu et al., 2022), which we have implemented for DDPM. Furthermore, we implemented TDPM (Zheng et al., 2022) which uses a truncated diffusion chain and adopts a GAN to estimate the noisy images at the truncation point. For quantitative analyses, we evaluated the super-resolution performance using three metrics, Normalized Root Mean Square Error (NRMSE), Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). Sampling efficiency was measured by the number of function evaluations (NFE) (Zheng et al., 2022).

Our FTDDM consists of two components: a diffusion model (Truncated Denoising Diffusion) and a normalizing flow model (Flow-based Noisy Image Generation). We conducted ablation studies to showcase the significance of each component. Key design elements, including the guide loss  $\mathcal{L}_{guide}$  and the uses of quality-filtering mask  $M$  and MRIs as condition images, were also justified through ablation studies. We studied the selection for  $T_{trunc}$ , testing the performance of FTDDM at values of  $T_{trunc}=20, 50$  and 100. We also attempted to combine DPM-Solver++ with FTDDM to further reduce the number of sampling steps, observing that the remaining sampling steps in the Truncated Denoising Diffusion model can be resolved using DPM-Solver++. Moreover, to confirm that conditioning on the upscaling factor using Equation 6 aids the network in processing each low resolution at the right scale, we studied the network performance across different combinations of actual and conditioned upscaling factors. Furthermore, effect of the temperature parameter  $\tau_d$  on the image sharpness was investigated. The image sharpness, or visual quality, was evaluated quantitatively using Learned Perceptual Image Patch Similarity (LPIPS), which measures the high-level similarity between two images using a pretrained VGG network and aligns closely with human perceptual judgment (Zhang et al., 2018).

To confirm the clinical advantage of our method, we invited 2 experienced neuroradiologists to assess the quality of our super-resolution images relative to ground truth from clinical perspectives. Radiologists' rating was considered as a key evaluation criterion in previous deep learning-based image recon-

**Table 1. Quantitative comparisons of FTDDM against other deep learning methods at three upscaling factors  $s=8.0, 4.0, 2.0$ . NFE: number of function evaluations, Time: Sampling time in seconds (s) for a single image slice, Params: number of model parameters in millions (M), NRMSE: Normalized Root Mean Square Error, PSNR: Peak Signal-to-Noise Ratio, SSIM: Structural Similarity Index.**

| Methods                                   | NFE        | Time (s)    | Params (M)  | $s=8.0$ (LR=8×8)   |                 |                    | $s=4.0$ (LR=16×16) |                 |                    | $s=2.0$ (LR=32×32) |                 |                    |
|---|------------|-------------|-------------|--------------------|-----------------|--------------------|--------------------|-----------------|--------------------|--------------------|-----------------|--------------------|
|   |            |             |             | NRMSE↓             | PSNR↑           | SSIM↑              | NRMSE↓             | PSNR↑           | SSIM↑              | NRMSE↓             | PSNR↑           | SSIM↑              |
| MUNet-FS-cWGAN                            | 1          | 0.05        | 27.5        | 0.414±0.188        | 26.0±2.3        | 0.864±0.043        | 0.344±0.164        | 27.7±2.5        | 0.904±0.036        | 0.238±0.110        | 30.9±2.8        | 0.956±0.019        |
| Flow Enhancer                             | 1          | 0.22        | 37.6        | 0.380±0.193        | 26.9±2.4        | 0.882±0.038        | 0.321±0.175        | 28.5±2.5        | 0.917±0.031        | 0.240±0.140        | 31.1±2.7        | 0.956±0.020        |
| DDPM ( $T=1000$ )                         | 1000       | 12.44       | 30.3        | 0.364±0.174        | 27.2±2.5        | 0.884±0.036        | 0.287±0.136        | 29.3±2.9        | 0.927±0.027        | 0.183±0.086        | 33.2±3.4        | 0.970±0.014        |
| DDPM ( $T_{respace}=100$ )                | 100        | 1.25        | 30.3        | 0.395±0.183        | 26.4±2.3        | 0.868±0.040        | 0.317±0.145        | 28.3±2.6        | 0.911±0.031        | 0.212±0.096        | 31.9±2.9        | 0.958±0.018        |
| DPM-Solver++ ( $T_{solver}=100$ )         | 101        | 1.46        | 30.3        | 0.413±0.183        | 26.0±2.5        | 0.864±0.041        | 0.325±0.147        | 28.1±2.9        | 0.911±0.032        | 0.208±0.094        | 32.1±3.4        | 0.962±0.017        |
| TDPM ( $T_{trunc}=100$ )                  | 101        | 1.27        | 60.6        | 0.382±0.169        | 26.6±2.6        | 0.876±0.038        | 0.290±0.138        | 29.2±3.0        | 0.927±0.028        | 0.175±0.082        | 33.6±3.7        | 0.972±0.014        |
| <b>FTDDM (<math>T_{trunc}=100</math>)</b> | <b>101</b> | <b>1.33</b> | <b>50.5</b> | <b>0.356±0.168</b> | <b>27.3±2.5</b> | <b>0.888±0.036</b> | <b>0.276±0.134</b> | <b>29.6±2.9</b> | <b>0.932±0.027</b> | <b>0.170±0.080</b> | <b>33.9±3.6</b> | <b>0.974±0.013</b> |

struction tasks (Muckley et al., 2021; Dong et al., 2022a; Zhou et al., 2022). To determine the lowest resolution the network can effectively manage, or the “breakpoint” of our method, we asked the neuroradiologists to evaluate the model performance at three extremely low input resolutions of  $6\times 6$ ,  $4\times 4$  and  $2\times 2$ .

### 3.3. Results and Discussion

#### 3.3.1. Quantitative evaluations

Table 1 presents the quantitative comparison of the proposed FTDDM against other deep learning models at three different upscaling factors  $s=8.0, 4.0$  and  $2.0$ . Compared to the GAN-based model (MUNet-FS-cWGAN) and the flow-based model (Flow Enhancer), the diffusion model DDPM excels across all metrics. This highlights the superior learning capability of the diffusion model within this set of generative models. However, DDPM requires 1000 NFE during sampling, translating to a sampling time of approximately 12.44 seconds for a single image slice on our devices. Reducing the NFE to 100 with the simple respacing method ( $T_{respace}=100$ ) downgrades the performance, as skipping the diffusion steps yields images that are less denoised. Using the same number of sampling steps, DPM-Solver++ with  $T_{solver}=100$  fails to perform better than the respacing method. However, we show in Section 3.3.4 that DPM-Solver++ can outperform the respacing method when fewer sampling steps are used. The truncation method TDPM at  $T_{trunc}=100$  outperforms the respacing method and offers comparable performance with the DDPM at  $T=1000$ . TDPM requires 101 NFE, due to an extra GAN function evaluation. Our proposed method, FTDDM, achieves even better performance than TDPM, indicating the flow-based network’s superiority over GAN in generating noisy images at the truncation point. For a single image slice, FTDDM at  $T_{trunc}=100$  has a sampling time of approximately 1.33 seconds, which is over 9-fold acceleration relative to DDPM. Since evaluating the flow-based network for generating noisy images at the truncation point takes extra time (one additional NFE), the speed increase is slightly below 10-fold. Wilcoxon signed-rank tests were conducted between each pair of scores for all methods in Table 1, for which the results indicate that all differences are significant with P values  $< 0.001$ .

#### 3.3.2. Qualitative evaluations

The quantitative improvement given by FTDDM is reflected by the better image quality and tumor visualization in qualitative analysis. The qualitative results of different methods at

two different upscaling factors  $s=4.0$  and  $s=8.0$  are displayed in Fig.3 and Fig.4, respectively. The low-resolution images are interpolated to the same resolution as the high-resolution images ( $64\times 64$ ) by filling the high-frequency components of the k-space with zeros, i.e. zero-filled images. Compared to the ground truth high-resolution images, the zero-filled low-resolution images appear blurry and lack high-frequency image details. While MUNet-FS-cWGAN and Flow Enhancer can recover a large portion of high-frequency features, they either fail to capture the tumor contrast (tCr map from Fig.3) or distort the tumor shape (Gly map from Fig.3 and both maps from Fig.4) when compared to the ground truth. DDPM has greater learning capability and reconstructs the tumors more consistent with the ground truth. Additionally, images given by DDPM more accurately capture the characteristics of white matter and gray matter. TDPM, due to its GAN dependence, can exhibit training instability and occasionally introduce pronounced artifacts, as observed in the tCh map from Fig.4. Among the diffusion models, FTDDM reconstructs tumor characteristics closest to ground truth and yields the lowest overall error. The error maps of FTDDM exhibit less bright spots compared to the other methods. The uncertainty map (framed in red) for FTDDM was calculated from the standard deviation of 50 repeated samples and acts as a tool for error estimation. Typically, uncertainty is more pronounced around the brain’s periphery, suggesting the model detected a higher variance in these areas. This observation aligns with De Graaf (2019) and Ziegs et al. (2023), which mentions the stronger signal distortion due to extracranial lipid contamination and  $B_0$  field inhomogeneity proximal to the skull.

#### 3.3.3. Ablation studies

Table 2 shows the ablation studies on the diffusion model and the normalizing flow model in FTDDM at an upscaling factor of  $s=4.0$ . FTDDM is equivalent to a normalizing flow model when the truncation point is selected as  $T_{trunc}=0$  and is equivalent to a diffusion model when  $T_{trunc}=1000$ . Using the flow model alone underperforms the combined model by a large margin due to the missing diffusion process. Conversely, using the diffusion model alone requires a large number of sampling steps while still underperforming the combined model. Ablation studies on specific design elements at an upscaling factor of  $s=8.0$  are presented in Table 3. Omitting the quality-filtering mask from the condition images makes the model interpret artificial features introduced by quality filtering as genuine image

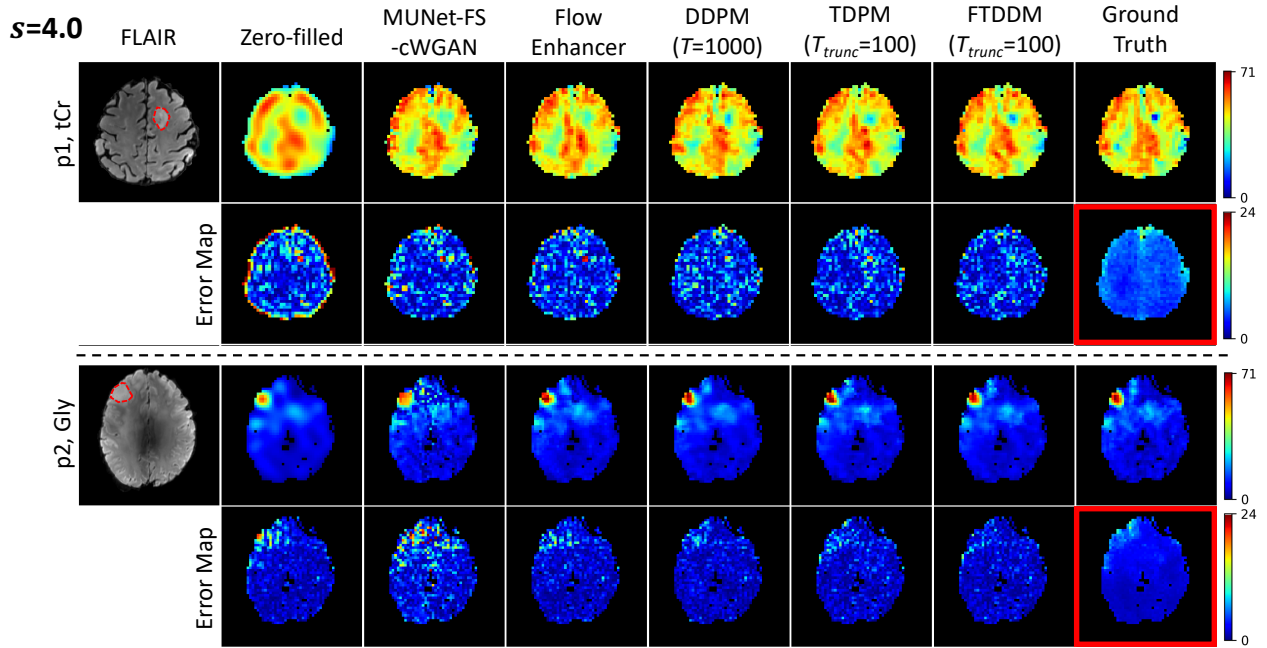


Fig. 3. Qualitative comparisons of FTDDM against other methods at upscaling factor  $s=4.0$ . The two examples are: a tCr image from patient p1 and a Gly image from patient p2. FLAIR MRI provides the corresponding anatomical reference, with the tumor delineated by the red dashed line. Each metabolic map is shown alongside with its error map, except for ground truth. Note that the images below ground truth, framed in red, are the standard deviation maps of 50 FTDDM samples and can be used for uncertainty estimation (they are not error maps of ground truth).

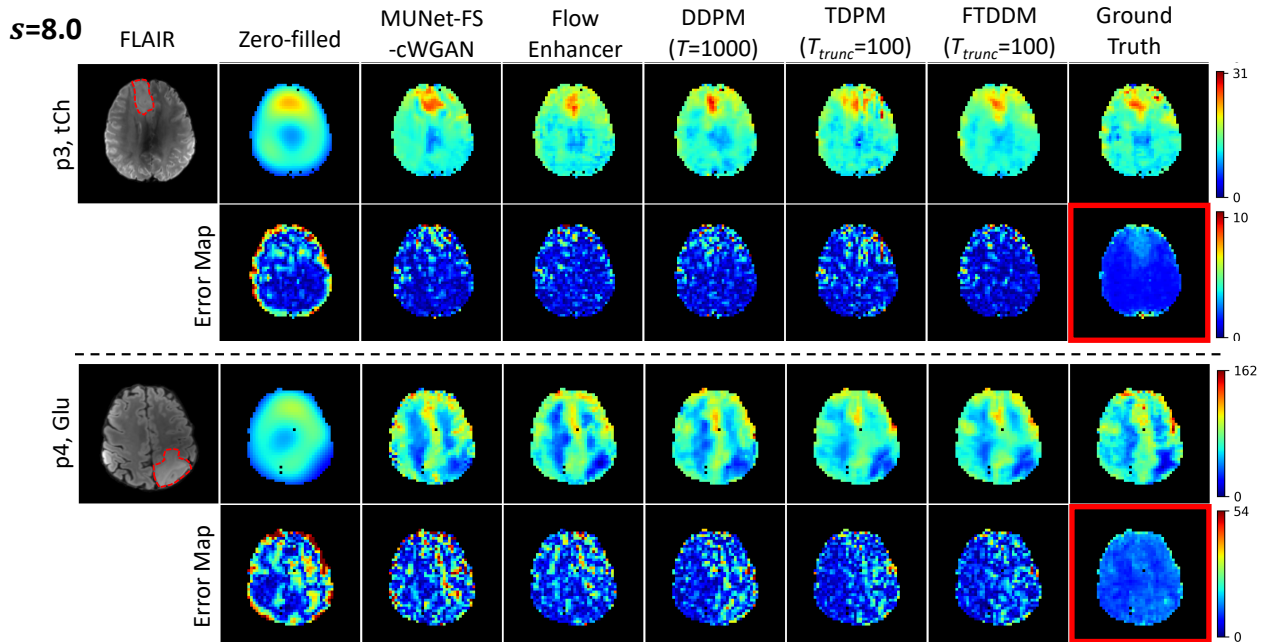


Fig. 4. Qualitative comparisons of FTDDM against other methods at upscaling factor  $s=8.0$ . The two examples are: a tCh image from patient p3 and a Glu image from patient p4.

features, resulting in a significant deterioration in model performance. Surprisingly, the removal of T1 and FLAIR MRIs only marginally affects the performance, which suggests that the correlation between MRIs and MRSI is somewhat trivial for our in vivo dataset. The network does not seem to integrate substantial anatomical information from MRIs into the super-resolution MRSI. This implies that the traditional super-resolution MRSI methods that rely heavily on MRI-based regularization (Jain

*et al.*, 2017; Lam and Liang, 2014; Kasten *et al.*, 2016) might fail on in vivo MRSI dataset. Lastly, the guidance loss helps to boost the model performance. Wilcoxon signed-rank tests indicate that all differences in Table 2 and Table 3 are statistically significant with P values  $< 0.001$ .



**Table 2. Ablation studies on the diffusion model and the normalizing flow model in FTDDM.**

| Method                               | NRMSE↓             | PSNR↑           | SSIM↑              |
|--------------------------------------|--------------------|-----------------|--------------------|
| Flow only ( $T_{trunc}=0$ )          | 0.320±0.155        | 28.3±2.3        | 0.910±0.030        |
| Diffusion only ( $T_{trunc}=1000$ )  | 0.282±0.134        | 29.4±2.9        | 0.929±0.027        |
| Diffusion + Flow ( $T_{trunc}=100$ ) | <b>0.276±0.134</b> | <b>29.6±2.9</b> | <b>0.932±0.027</b> |

**Table 3. Ablation studies on design elements. ✓ or × indicates whether a certain design element is included or not.**


| Mask | MRI | $L_{guide}$ | NRMSE↓             | PSNR↑           | SSIM↑              |
|------|-----|-------------|--------------------|-----------------|--------------------|
| ×    | ✓   | ✓           | 0.419±0.221        | 25.9±2.7        | 0.866±0.047        |
| ✓    | ×   | ✓           | 0.361±0.170        | 27.2±2.5        | 0.885±0.036        |
| ✓    | ✓   | ×           | 0.358±0.169        | 27.3±2.4        | 0.884±0.035        |
| ✓    | ✓   | ✓           | <b>0.356±0.168</b> | <b>27.3±2.5</b> | <b>0.888±0.036</b> |

### 3.3.4. Selection of $T_{trunc}$

We studied the performance of different models across various sampling steps, specifically, 100, 50 and 20 (Fig.5). Compared to the unaccelerated DDPM (the cyan line), the resampling method (the blue line) significantly degrades the performance as fewer sampling steps are used ( $T_{respace}=50$  and  $T_{respace}=20$ ). Using the ODE solver DPM-Solver++, the rate of this degradation is considerably slower. The truncation method TDPM outperforms DPM-Solver++, and our FTDDM achieves the best performance among these acceleration methods. At around  $T_{trunc}=50$ , FTDDM achieves a similar level of performance as the unaccelerated DDPM, which is equivalent to an approximately of 20-fold acceleration. Therefore, although we used  $T_{trunc}=100$  throughout this paper to pursue a better performance, our model can certainly be operated at  $T_{trunc}=50$  to attain greater acceleration. At  $T_{trunc} > 50$ , our FTDDM is able to outperform DDPM because the Denoising UNet can focus on denoising a less noisy image  $\mathbf{x}_{T_{trunc}}$ , which is easier to learn than denoising a pure Gaussian noise image  $\mathbf{z}$  as in DDPM. This is consistent with the observation that less training iteration is required for the diffusion part to converge when  $T_{trunc}$  is smaller (Zheng *et al.*, 2022).

### 3.3.5. Combining with DPM-Solver++

Observing that the remaining sampling steps in the Truncated Denoising Diffusion model can be further reduced with DPM-Solver++, we studied how combining DPM-Solver++ with FT-


**Fig. 5. Model performance of DDPM, DDPM with resampling, DPM-Solver++, TDPM and FTDDM at different numbers of sampling steps ( $T_{respace}$  for DDPM respace,  $T_{solver}$  for DPM-Solver++,  $T_{trunc}$  for TDPM and FTDDM).****Table 4. Study of combining DPM-Solver++ with FTDDM to further reduce the number of sampling steps. SSIM scores are shown for each valid combination of  $T_{trunc}$  in FTDDM and  $T_{solver}$  in DPM-Solver++.**

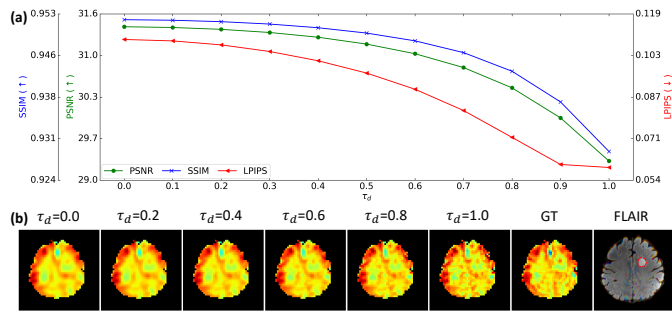
|                 | no DPM-solver++ | $T_{solver}=50$ | $T_{solver}=20$ | $T_{solver}=10$ |
|-----------------|-----------------|-----------------|-----------------|-----------------|
| $T_{trunc}=100$ | 0.932±0.027     | 0.928±0.028     | 0.928±0.028     | 0.927±0.028     |
| $T_{trunc}=50$  | 0.929±0.027     | -               | 0.922±0.029     | 0.922±0.029     |
| $T_{trunc}=20$  | 0.919±0.029     | -               | -               | 0.912±0.031     |

DDM impacts the model performance, as shown in Table 4. When  $T_{trunc}=100$ , the number of sampling steps (or NFE) can be further reduced to 50, 20 or 10 by using DPM-Solver++ at  $T_{solver}=50, 20$  or 10. When  $T_{trunc}=50$ , the number of sampling steps can be further reduced to 20 or 10 by using  $T_{solver}=20$  or 10. For  $T_{trunc}=20$ , it can be reduced to 10. It can be observed that introducing DPM-Solver++ always slightly worsens the performance compared to using FTDDM alone. However, achieving further reduction in sampling steps with fewer  $T_{solver}$  only marginally hurts the performance, better than using DPM-Solver++ alone as shown in Fig.5. Before applying DPM-Solver++ for further reduction, a larger  $T_{trunc}$ , such as 100, is recommended to ensure a higher upper limit of the performance.

### 3.3.6. Justification of network conditioning

We subsequently studied the model performance under various combinations of actual upscaling factors and conditioned upscaling factors, as shown in Fig.6. The model performance, measured with SSIM, consistently reaches its optimum when the conditioned upscaling factor matches with the actual upscaling factor. This indicates that leveraging Conditional Instance Normalization to modulate the Denoising UNet empowers the network to super-resolve different low resolutions at their appropriate scales.

**Fig. 6. Model performance (measured in SSIM) under different combinations of actual upscaling factors (horizontal axis) and conditioned upscaling factors (color bars). For conciseness, only 5 adjacent values of conditions are shown for each actual upscaling factor.**



**Fig. 7. Sharpness adjustment with  $\tau_d$ .** (a) Network performance in PSNR, SSIM and LPIPS at various levels of  $\tau_d$ . Lower LPIPS (↓) is better. (b) Visual sharpness of a super-resolution Ins image at various  $\tau_d$  compared to ground truth (GT). The tumor is delineated by the red dashed line in FLAIR.

### 3.3.7. Sharpness adjustability

Our method provides the capability to adjust image sharpness within the same network. The visual sharpness of the super-resolution images can be flexibly controlled by the temperature parameter  $\tau_d$ , which comes with a trade-off between image visual quality and image fidelity. As shown in Fig.7(a), using smaller values of  $\tau_d$  gives higher fidelity (indicated by higher PSNR and SSIM) but worse visual quality (indicated by higher LPIPS scores). Conversely, a larger  $\tau_d$  improves the visual quality but downgrades the fidelity. Fig.7(b) shows the corresponding images. Smaller  $\tau_d$  values, such as 0.0 and 0.2, produces blurry images with reduced tumor contrast compared to the ground truth. Setting  $\tau_d$  between 0.8 and 1.0 achieves sharpness close to the ground truth. According to the curves in Fig.7(a), increasing  $\tau_d$  from 0.9 to 1.0 yields marginal enhancement in visual quality (LPIPS) but at a significant cost to fidelity (PSNR and SSIM). Therefore, we adopted  $\tau_d=0.9$  for the entirety of this work.

### 3.3.8. Image quality assessment from neuroradiologists

Fig.8 shows neuroradiologists’ ratings to demonstrate the clinical benefits of the super-resolution images given by FTDDM. We presented 2 experienced neuroradiologists with pairs of low-resolution and super-resolution images in a blind way (not labeled) from 20 slices, evenly selected from 10 patients. Each slice contains images of 7 metabolites, summing to a total evaluation of 140 image pairs. These images spanned five up-scaling factors,  $s=8.0, 6.4, 5.3, 4.6$  and  $4.0$ , for a thorough evaluation. Conspicuity of each metabolite (tCr, tCh, NAA, Glu, Ins, Gln and Gly) was judged relative to the ground truth and scored as: 10 (poor), 30 (bad), 50 (fair), 70 (good) and 90 (excellent) (Chow et al., 2016). Using all metabolites, an overall sharpness and the overall conspicuities of gray matter, white matter, ventricles and lesion were judged relative to the ground truth and scored using the same scoring system. For the slices in which ventricles or lesion were not detected by the neuroradiologists, an “NA” was given. We observe that the scores for the low-resolution images ranged between 10-70 for Neuroradiologist 1 and predominantly stood at 50 for Neuroradiologist 2. After being super-resolved by FTDDM, Neuroradiologist 1 rated over 87% of the images as 90 (excellent), and Neuroradiologist

2 rated over 96% of the images as 90, across all evaluation metrics. This reveals a significant enhancement compared to the low-resolution images. Fig.9 shows the low-resolution, super-resolution and ground truth images of 7 metabolites from a slice containing a lesion, which is an example of what was presented to the neuroradiologists for image quality assessment. It is obvious that the lesion is less conspicuous in the  $8\times 8$  zero-filled low-resolution images, and FTDDM successfully reconstructed the lesion contrast in tCr, tCho, Ins, Gln and Gly. This affirms why all super-resolution images were scored 90 (excellent) for lesion conspicuity. Moreover, the super-resolution images exhibit clear contrast between white matter and gray matter. This is especially observable in tCr, NAA and Glu maps, where concentrations are higher in gray matter. These metabolic characteristics are consistent with the ground truth images and the biological facts (Pouwels and Frahm, 1998).

### 3.3.9. Determining the lowest input resolution

Fig.10 presents neuroradiologists’ ratings for images at three extremely low resolutions to determine the “breakpoint” of our method. Low-resolution images at  $2\times 2, 4\times 4$  and  $6\times 6$  predominantly scored 10 across all metrics, indicating very poor quality. Super-resolution results of the  $6\times 6$  images mostly reach scores of 70 and 90, which is a significant enhancement. For  $4\times 4$ , the super-resolution images were scored 50 or 70, indicating borderline acceptability. However, when it comes to  $2\times 2$ , many super-resolution images were scored 50 or below. Supplementary Fig.S1 provides examples of these low-resolution and super-resolution images, highlighting that super-resolution images of  $2\times 2$  do not align well with ground truth. Notably, tumor contrast begins to emerge in the super-resolution images of  $6\times 6$ . As a result, our confidence in the super-resolution results is retained only for the low-resolutions higher than  $6\times 6\times 39$  (equivalent to approximately  $36\times 36\times 3.4\text{mm}^3$ ); below this, our method might fail. For future applications, to achieve good super-resolution image quality, it is recommended that low-resolution MRSI data is acquired at resolutions finer than  $36\times 36\times 3.4\text{mm}^3$  (assuming similar acquisition techniques).

### 3.3.10. Clinical values

Lastly, it is worth noting that our 3D high-resolution MRSI datasets were acquired on a 7T MRI scanner using spatial-spectral encoding, which make it possible to acquire  $3.4\times 3.4\times 3.4\text{mm}^3$  resolution data in only 15 minutes. Using a standard phase encoding scheme on a clinical 3T MRI system (TR=450 ms and ellipsoidal k-space encoding) would require approximately 8 hours of scan time to achieve a similar resolution, which is clinically unfeasible. The proposed FTDDM, trained with our dataset, is able to transform lower resolution and faster scans into high-resolution metabolic maps. For example, an MRSI acquisition at  $8\times 8\times 39$  (or  $27\times 27\times 3.4\text{mm}^3$ ) can in theory be completed in approximately 2 minutes, which is a significant reduction in scan time from the 15 minutes requirement for a  $64\times 64\times 39$  (or  $3.4\times 3.4\times 3.4\text{mm}^3$ ) scan using our acquisition protocol. This reduction in scan time elevates MRSI’s potential to be integrated in a routine clinical neuroimaging protocol. Furthermore, this time reduction also




Fig. 8. Image quality assessment based on neuroradiologists' ratings. Image quality of low-resolution (LR) images and super-resolution (SR) images given by FTDDM was blind-reviewed by Neurologist 1 and 2. The horizontal axis shows the evaluation metrics. Each colored bar denotes the count of images that received a specific score for the corresponding metric. The scoring system is: NA (not applicable), 10 (poor), 30 (bad), 50 (fair), 70 (good) and 90 (excellent).

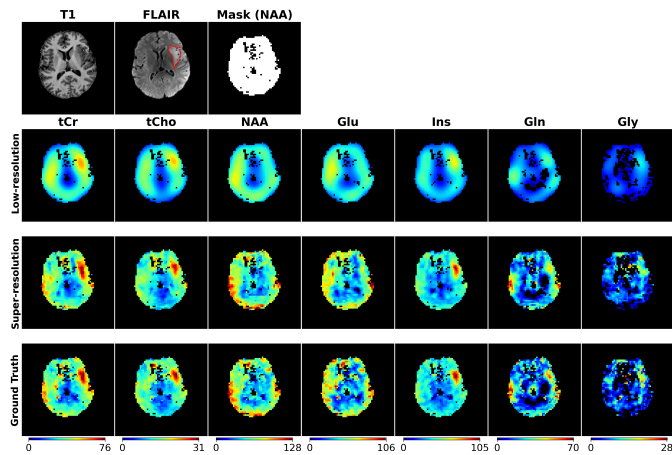


Fig. 9. Visualization of low-resolution (zero-filled from  $8 \times 8$ ), FTDDM super-resolution and ground truth images of 7 metabolites from a slice containing tumor. The first row shows T1 MRI, FLAIR MRI and an example of quality-filtering mask for NAA. The tumor is delineated by the red dashed line in FLAIR.

markedly decreases effects of motion, an important consideration when performing MRSI in clinical populations.


#### 4. Conclusion

In this work, we introduce a novel Flow-based Truncated Denoising Diffusion Model (FTDDM) for super-resolution MRSI. Our method truncates the diffusion chain and hence accelerates

the sampling process of the diffusion model. The noisy image at the truncation point is estimated directly from Gaussian noise via a normalizing flow-based network. The diffusion network is conditioned on the upscaling factor, facilitating multi-scale super-resolution. Additionally, a temperature parameter is incorporated into our diffusion model to allow sharpness adjustment. Experimental results on our self-developed in vivo  $^1\text{H}$ -MRSI dataset indicate that FTDDM outperforms other deep learning models and achieves over 9-fold acceleration compared to the conventional diffusion model without loss of image quality. Neuroradiologists' assessments confirm the excellent image quality given by the proposed method from clinical perspectives and assist in identifying the lowest operable resolution. To the best of our knowledge, we are the first to develop deep learning-based super-resolution models and conduct comprehensive evaluations on in vivo MRSI dataset. Future work to improve the generalization and robustness of the proposed method can involve the inclusion of more datasets acquired from other MRI systems (Nassirpour et al., 2018; Dong et al., 2024) and/or from patients with different tumor types. The proposed model can also be extended to deuterium metabolic imaging (DMI) (De Feyter et al., 2018; Dong et al., 2020).

#### CRediT authorship contribution statement

**Siyuan Dong:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Visualization, Writing - original draft. **Zhuotong Cai:** Conceptual-



**Fig. 10.** Image quality assessment for three extremely low resolutions  $2\times 2$ ,  $4\times 4$  and  $6\times 6$ . Due to limited space, only a subset of most representative evaluation metrics is displayed. For each metric shown on the horizontal axis, ratings for low-resolution (LR) images at  $2\times 2$ ,  $4\times 4$  and  $6\times 6$  (indicated by the number in the parenthesis) and their corresponding super-resolution (SR) images are presented.

ization, Writing - review & editing. **Gilbert Hangel:** Data curation, Writing - review & editing. **Wolfgang Bogner:** Data curation. **Georg Widhalm:** Data curation. **Yaqing Huang:** Writing - review & editing. **Qinghao Liang:** Writing - review & editing. **Chenyu You:** Writing - review & editing. **Chathura Kumaragamage:** Writing - review & editing. **Robert K. Fulbright:** Validation. **Amit Mahajan:** Validation. **Amin Karbasi:** Writing - review & editing. **John A. Onofrey:** Conceptualization, Writing - review & editing. **Robin A. de Graaf:** Conceptualization, Validation, Writing - review & editing. **James S. Duncan:** Conceptualization, Writing - review & editing, Supervision, Funding acquisition.

## Data availability

The code will be available online. Data will be made available on request.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: James S. Duncan is one of the Editors-in-Chief for Medical Image Analysis.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used GPT-4 from OpenAI in order to improve readability and language. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

## Acknowledgments

This work is part of Siyuan Dong's PhD research (Dong, 2024), supported by the National Institutes of Health under R01EB025840, R01CA206180 and R01NS035193. The data acquisition was supported by the Austrian Science Fund grants KLI 646, P 30701 and P 34198.

## Supplementary material

Supplementary material was prepared and provided as a separate file along with the manuscript.

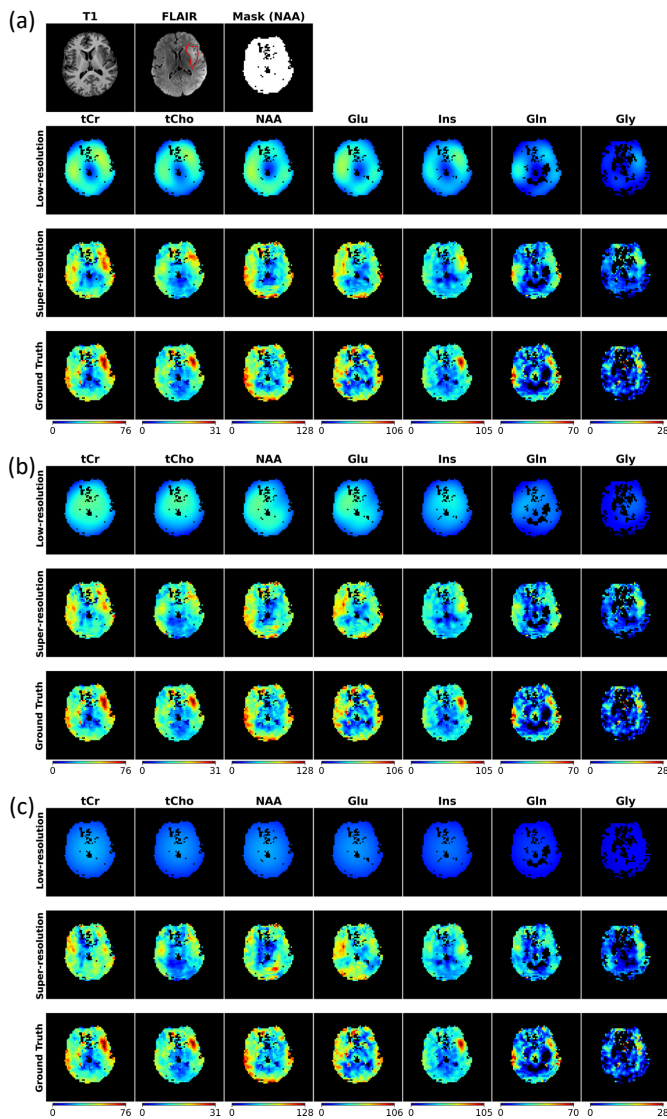
## References

- Ardizzone, L., Kruse, J., Rother, C., Köthe, U., 2018. Analyzing inverse problems with invertible neural networks, in: International Conference on Learning Representations.
- Bogner, W., Otazo, R., Henning, A., 2021. Accelerated mr spectroscopic imaging—a review of current and emerging techniques. *NMR in Biomedicine* 34, e4314.
- Chen, Z., Dong, S., Sun, S., Chen, X., Liu, Y., Chen, T., 2024a. Mri reconstruction based on generative models. US Patent App. 17/891,702.
- Chen, Z., Dong, S., Sun, S., Chen, X., Liu, Y., Chen, T., 2024b. Systems and methods for processing medical images with invertible neural networks. US Patent App. 17/891,668.
- Chow, L.S., Rajagopal, H., Paramesran, R., Initiative, A.D.N., et al., 2016. Correlation between subjective and objective assessment of magnetic resonance (mr) images. *Magnetic resonance imaging* 34, 820–831.
- Chung, H., Ye, J.C., 2022. Score-based diffusion models for accelerated mri. *Medical image analysis* 80, 102479.
- De Feyter, H.M., Behar, K.L., Corbin, Z.A., Fulbright, R.K., Brown, P.B., McIntyre, S., Nixon, T.W., Rothman, D.L., de Graaf, R.A., 2018. Deuterium metabolic imaging (dmi) for mri-based 3d mapping of metabolism in vivo. *Science advances* 4, eaat7314.
- De Graaf, R.A., 2019. *In vivo NMR spectroscopy: principles and techniques*. John Wiley & Sons.
- Dhariwal, P., Nichol, A., 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* 34, 8780–8794.
- Dong, S., 2024. *Data-Driven Methods for Super-Resolution Magnetic Resonance Spectroscopic Imaging*. Ph.D. thesis. Yale University.

- Dong, S., Chen, E.Z., Zhao, L., Chen, X., Liu, Y., Chen, T., Sun, S., 2022a. Invertible sharpening network for mri reconstruction enhancement, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 582–592.
- Dong, S., De Feyter, H.M., Thomas, M.A., de Graaf, R.A., Duncan, J.S., 2020. A deep learning method for sensitivity enhancement in deuterium metabolic imaging (dmi), in: Proceedings of the 28th Annual Meeting of ISMRM.
- Dong, S., Feyter, H.M.D., Thomas, M.A., de Graaf, R.A., Duncan, J.S., 2023. Preserved edge convolutional neural network for sensitivity enhancement of deuterium metabolic imaging (dmi). *arXiv:2309.04100*.
- Dong, S., Hangel, G., Bogner, W., Trattnig, S., Rössler, K., Widhalm, G., De Feyter, H.M., De Graaf, R.A., Duncan, J.S., 2021. High-resolution magnetic resonance spectroscopic imaging using a multi-encoder attention u-net with structural and adversarial loss, in: 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), IEEE. pp. 2891–2895.
- Dong, S., Hangel, G., Bogner, W., Widhalm, G., Rössler, K., Trattnig, S., You, C., de Graaf, R., Onofrey, J.A., Duncan, J.S., 2022b. Multi-scale super-resolution magnetic resonance spectroscopic imaging with adjustable sharpness, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 410–420.
- Dong, S., Hangel, G., Chen, E.Z., Sun, S., Bogner, W., Widhalm, G., You, C., Onofrey, J.A., de Graaf, R., Duncan, J.S., 2022c. Flow-based visual quality enhancer for super-resolution magnetic resonance spectroscopic imaging, in: MICCAI Workshop on Deep Generative Models, Springer. pp. 3–13.
- Dong, S., Shewarega, A., Chapiro, J., Cai, Z., Hyder, F., Coman, D., Duncan, J.S., 2024. High-resolution extracellular ph imaging of liver cancer with multiparametric mr using deep image prior. *NMR in Biomedicine*, e5145.
- Elfving, S., Uchibe, E., Doya, K., 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks* 107, 3–11.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Communications of the ACM* 63, 139–144.
- Hangel, G., Cadrien, C., Lazen, P., Furtner, J., Lipka, A., Hečková, E., Hingerl, L., Motyka, S., Gruber, S., Strasser, B., et al., 2020. High-resolution metabolic imaging of high-grade gliomas using 7t-crt-fid-mrsi. *NeuroImage: Clinical* 28, 102433.
- Hangel, G., Jain, S., Springer, E., Hečková, E., Strasser, B., Považan, M., Gruber, S., Widhalm, G., Kiesel, B., Furtner, J., et al., 2019. High-resolution metabolic mapping of gliomas via patch-based super-resolution magnetic resonance spectroscopic imaging at 7t. *Neuroimage* 191, 587–595.
- Hingerl, L., Strasser, B., Moser, P., Hangel, G., Motyka, S., Heckova, E., Gruber, S., Trattnig, S., Bogner, W., 2020. Clinical high-resolution 3d-mr spectroscopic imaging of the human brain at 7 t. *Investigative radiology* 55, 239–248.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33, 6840–6851.
- Huang, X., Belongie, S., 2017. Arbitrary style transfer in real-time with adaptive instance normalization, in: Proceedings of the IEEE international conference on computer vision, pp. 1501–1510.
- Iqbal, Z., Nguyen, D., Hangel, G., Motyka, S., Bogner, W., Jiang, S., 2019. Super-resolution 1h magnetic resonance spectroscopic imaging utilizing deep learning. *Frontiers in oncology* 9, 1010.
- Jain, S., Sima, D.M., Sanaei Nezhad, F., Hangel, G., Bogner, W., Williams, S., Van Huffel, S., Maes, F., Smeets, D., 2017. Patch-based super-resolution of mr spectroscopic images: application to multiple sclerosis. *Frontiers in neuroscience* 11, 13.
- Kasten, J., Klausner, A., Lazeyras, F., Van De Ville, D., 2016. Magnetic resonance spectroscopic imaging at superresolution: overview and perspectives. *Journal of Magnetic Resonance* 263, 193–208.
- Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Hachililoglu, I., Merhof, D., 2022. Diffusion models for medical image analysis: A comprehensive survey. *arXiv preprint arXiv:2211.07804*.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lam, F., Liang, Z.P., 2014. A subspace approach to high-resolution spectroscopic imaging. *Magnetic resonance in medicine* 71, 1349–1357.
- Li, H., Yang, Y., Chang, M., Chen, S., Feng, H., Xu, Z., Li, Q., Chen, Y., 2022a. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* 479, 47–59.
- Li, X., Strasser, B., Neuberger, U., Vollmuth, P., Bendszus, M., Wick, W., Dietrich, J., Batchelor, T.T., Cahill, D.P., Andronesi, O.C., 2022b. Deep learning super-resolution magnetic resonance spectroscopic imaging of brain metabolism and mutant isocitrate dehydrogenase glioma. *Neuro-oncology advances* 4, vda071.
- Li, Y., Sixou, B., Peyrin, F., 2021. A review of the deep learning methods for medical images super resolution problems. *Irbm* 42, 120–133.
- Liang, J., Lugmayr, A., Zhang, K., Danelljan, M., Van Gool, L., Timofte, R., 2021. Hierarchical conditional flow: A unified framework for image super-resolution and image rescaling, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4076–4085.
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J., 2022. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*.
- Lugmayr, A., Danelljan, M., Van Gool, L., Timofte, R., 2020. Srflo: Learning the super-resolution space with normalizing flow, in: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16, Springer. pp. 715–732.
- Muckley, M.J., Riemenschneider, B., Radmanesh, A., Kim, S., Jeong, G., Ko, J., Jun, Y., Shin, H., Hwang, D., Mostapha, M., et al., 2021. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE transactions on medical imaging* 40, 2306–2317.
- Nassirpour, S., Chang, P., Henning, A., 2018. High and ultra-high resolution metabolite mapping of the human brain using 1h fid mrsi at 9.4 t. *Neuroimage* 168, 211–221.
- Nichol, A.Q., Dhariwal, P., 2021. Improved denoising diffusion probabilistic models, in: International Conference on Machine Learning, PMLR. pp. 8162–8171.
- Peng, C., Guo, P., Zhou, S.K., Patel, V.M., Chellappa, R., 2022. Towards performant and reliable undersampled mr reconstruction via diffusion model sampling, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 623–633.
- Pouwels, P.J., Frahm, J., 1998. Regional metabolite concentrations in human brain as determined by quantitative localized proton mrs. *Magnetic resonance in medicine* 39, 53–60.
- Provencher, S.W., 2014. Lcmodel & lcmgui user’s manual. *LCModel version 6*.
- Rezende, D., Mohamed, S., 2015. Variational inference with normalizing flows, in: International conference on machine learning, PMLR. pp. 1530–1538.
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M., 2022. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 4713–4726.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., et al., 2004. Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage* 23, S208–S219.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., 2015. Deep unsupervised learning using nonequilibrium thermodynamics, in: International conference on machine learning, PMLR. pp. 2256–2265.
- Song, Y., Ermon, S., 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* 32.
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B., 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Wang, Z., Chen, J., Hoi, S.C., 2020. Deep learning for image super-resolution: A survey. *IEEE transactions on pattern analysis and machine intelligence* 43, 3365–3387.
- Wang, Z., Simoncelli, E.P., Bovik, A.C., 2003. Multiscale structural similarity for image quality assessment, in: The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003, Ieee. pp. 1398–1402.
- Winkler, C., Worrall, D., Hoogeboom, E., Welling, M., 2019. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*.
- Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Shao, Y., Zhang, W., Cui, B., Yang, M.H., 2022. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*.
- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O., 2018. The unreasonable effectiveness of deep features as a perceptual metric, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 586–595.
- Zhao, H., Gallo, O., Frosio, I., Kautz, J., 2016. Loss functions for image restoration with neural networks. *IEEE Transactions on computational imaging* 3, 47–57.

- Zheng, H., He, P., Chen, W., Zhou, M., 2022. Truncated diffusion probabilistic models and diffusion-based adversarial auto-encoders, in: The Eleventh International Conference on Learning Representations.
- Zhou, B., Schlemper, J., Dey, N., Salehi, S.S.M., Sheth, K., Liu, C., Duncan, J.S., Sofka, M., 2022. Dual-domain self-supervised learning for accelerated non-cartesian mri reconstruction. *Medical Image Analysis* 81, 102538.
- Ziegs, T., Wright, A.M., Henning, A., 2023. Test–retest reproducibility of human brain multi-slice 1h fid-mrsi data at 9.4 t after optimization of lipid regularization, macromolecular model, and spline baseline stiffness. *Magnetic resonance in medicine* 89, 11–28.

## Supplementary Material



**Fig. S1.** Study of three extremely low resolutions. This is an example of the images provided to the neuroradiologists to give the scores in Fig.10. The low-resolutions at  $6\times 6$ ,  $4\times 4$  and  $2\times 2$  are shown in (a), (b) and (c). The tumor is delineated by the red dashed line in FLAIR.