

Dual-Prior Guided Null-Space Learning with Mixture-of-Splines for Arbitrary Medical Slice Super-Resolution

Haofei Song[✉], Siyuan Xu[✉], Xintian Mao[✉], Shaojie Guo[✉],
Qingli Li[✉], and Yan Wang[★][✉]

Shanghai Key Laboratory of Multidimensional Information Processing,
East China Normal University, Shanghai 200241, China
{hfsong, syxu, 52265904010, 52275904013}@stu.ecnu.edu.cn,
qlii@cs.ecnu.edu.cn, ywang@cee.ecnu.edu.cn

Abstract. Arbitrary slice super-resolution reconstructs isotropic volumes from anisotropic clinical acquisitions by synthesizing intermediate slices at arbitrary scales. However, treating this ill-posed inverse problem as unconstrained residual-based regression risks hallucinating anatomically implausible structures or altering the originally observed data. To address both concerns, this paper presents the **Dual-Prior Null-Space Learning** (DP-NSL) framework, which reformulates the task as a constrained recovery process guided by two complementary priors. A Measurement-Consistent Projection (MCP) enforces a *Deterministic Observation Prior*: the reconstruction undergoes an exact orthogonal projection that reproduces every acquired slice with zero error, confining all learned details to the unobservable null space. Within this null space, a Mixture-of-Splines (MoS) module imposes a *Geometric Continuity Prior* by dynamically mixing B-spline experts of different analytic orders, allowing each anatomical region to be modeled with a content-aware level of continuity. To promote spatial coherence, a Local Spatial Consistency Decoder (LSCD) further injects local inductive bias. Experiments on three CT and one MRI benchmark show that DP-NSL outperforms existing approaches while strictly preserving measurement consistency. Code is available at <https://github.com/DeepMed-Lab-ECNU/Medical-Image-Reconstruction>.

Keywords: Medical image super-resolution · Arbitrary scale · Null-space decomposition · B-spline

1 Introduction

Modern volumetric imaging pipelines, including computed tomography (CT) and magnetic resonance imaging (MRI), routinely produce anisotropic volumes whose through-plane spacing far exceeds the in-plane resolution. This anisotropy,

[★] Corresponding author.




Fig. 1: (a) Residual-based methods predict full-resolution outputs without measurement constraints, risking unintended alteration of acquired slices and anatomical hallucinations. (b) The proposed DP-NSL framework structures the reconstruction through two complementary priors: a *Deterministic Observation Prior* that anchors the output to the observed slices, and a *Geometric Continuity Prior* that models unobserved regions with multi-order continuity.

driven by clinical constraints on acquisition time and radiation dose, compromises downstream tasks such as 3D visualization, medical analysis and diagnosis. Arbitrary slice super-resolution, the task of synthesizing intermediate slices at arbitrary scales, has therefore become a necessary post-processing step.

Classical interpolation methods (*e.g.*, linear, cubic) offer deterministic stability but fail to recover high-frequency anatomical details, yielding overly smooth reconstructions. Recent implicit neural representations (INRs) [5] model the volume as a continuous function and achieve superior visual quality. However, medical arbitrary slice SR is an ill-posed inverse problem. As shown in Fig. 1(a), treating it as unconstrained regression risks hallucinating anatomically implausible structures or altering the originally observed slices. Black-box networks alone cannot provide the guarantees needed in clinical settings; instead, the reconstruction should be reframed as a constrained process guided by strong inductive biases. The challenge in medical slice SR therefore extends beyond resolution enhancement: the solution must also satisfy **measurement fidelity** and **anatomical plausibility**.

To this end, this work presents the **Dual-Prior Null-Space Learning** (DP-NSL) framework, which replaces unconstrained prediction with dual-prior constrained reconstruction: a *Deterministic Observation Prior* that guarantees fidelity to acquired observations, and a *Geometric Continuity Prior* that regularizes the synthesis of unobserved anatomical details.

To enforce the *Deterministic Observation Prior*, the framework introduces a Measurement-Consistent Projection (MCP) grounded in null-range-space decomposition [28]. By projecting the solution onto the observation hyperplane, MCP partitions the reconstruction into a deterministic range-space anchor and an unobserved null-space component. This projection mathematically guarantees consistency with the original acquisition: the observed data remain unaltered, while a Null-Space Estimator (NSE) recovers only the missing structures.

Within this null space, synthesizing missing details requires structural regularization to avoid hallucinations. Biological tissues exhibit heterogeneous smoothness: homogeneous organs such as the liver parenchyma present smooth intensity transitions, whereas structural interfaces like bone-soft tissue boundaries demand sharp representations [3, 8]. To impose the *Geometric Continuity Prior*, this work proposes a Mixture-of-Splines (MoS) module. Instead of learning implicit black-box interpolations, MoS explicitly treats continuous B-spline bases [20] as geometric priors. By dynamically mixing spline experts with different mathematical orders of continuity, the framework adaptively allocates the optimal prior stiffness for each localized anatomical region.

To complement this macroscopic geometric prior, a multi-scale Local Spatial Consistency Decoder (LSCD) is introduced. Unlike standard point-wise MLPs that process coordinates independently, LSCD provides a local inductive bias by aggregating neighboring contexts through depthwise convolutions, promoting spatial coherence in the synthesized null-space details.

The main contributions are summarized as follows:

- This work formulates arbitrary medical slice SR as a constrained inverse problem and presents the DP-NSL framework, replacing unconstrained residual-based regression with dual-prior guided reconstruction.
- By integrating null-range-space decomposition, the proposed MCP provides a hard data-fidelity constraint that keeps every originally observed clinical slice unchanged throughout the reconstruction.
- MoS establishes a structural continuity prior within the null space by treating B-splines as explicit geometric priors and adaptively routing smoothness orders according to local anatomy; an LSCD supplies complementary local consistency.
- Experiments across CT and MRI datasets show that DP-NSL outperforms existing methods, achieving up to 1.07 dB PSNR improvement at in-scales while generalizing well to out-of-scales.

2 Related Work

2.1 Arbitrary-Scale Image Super-Resolution

Arbitrary-scale super-resolution (ASSR) aims to reconstruct high-resolution images at any upsampling factor from a single trained model. Meta-SR [11] first explored this direction by dynamically predicting upsampling convolution weights conditioned on the scale factor. LIIF [5] then established the dominant paradigm by mapping continuous coordinates and 2D latent features to pixel intensities via a shared MLP. Subsequent works improve the implicit decoding stage along different axes: LTE [16] introduces Fourier-based texture estimation, CiaoSR [4] learns ensemble weights through implicit attention, HIIF [13] employs hierarchical positional encoding, and LMF [9] decouples latent modulation from per-pixel rendering for efficiency. However, LIIF-based methods render each coordinate through a point-wise MLP, limiting their ability to exploit local correlations.

Neural-operator-based approaches [18, 19, 29] offer an alternative by learning resolution-independent spectral transformations via Fourier neural operators.

A complementary direction estimates parameters of explicit interpolation bases rather than relying on black-box decoders. BTC [20] predicts B-spline coefficients, knots, and dilations from encoded features, whose compact support and non-negativity naturally suppress the ringing artifacts that affect Fourier-based representations near sharp edges. BF-STVSR [14] further validates this principle for temporal interpolation in spatiotemporal video SR. These results indicate that principled interpolation families, combined with data-driven parameter estimation, can achieve favorable accuracy–efficiency trade-offs, and they motivate the spline-based design adopted here.

2.2 Medical Volumetric Super-Resolution

Medical volumetric SR recovers high-resolution 3D images from sub-optimal clinical acquisitions. Distinct from natural images, medical data impose strict anatomical plausibility requirements. Many existing 3D SR approaches are designed for fixed, integer scale factors. To address the severe anisotropy in routine CT and MRI scans, methods like SAINT [21] and I3Net [22] decompose slice synthesis into specific inter-slice and intra-slice processes. Others [25] propose cross-view texture transfer, using multi-reference non-local attention to improve interpolations.

To enable arbitrary upsampling, ArSSR [30] extends implicit neural representations to medical volumes. Subsequent methods build upon this foundation to improve reconstruction quality. For example, SAINR [27] enhances feature querying with spatial attention, CycleINR [7] introduces cycle-consistency to reduce over-smoothing, and DC²SR [32] adopts dual-consistency guidance with curriculum learning.

However, existing INR-based arbitrary-scale methods still rely on coordinate-wise MLPs that overlook local spatial correlations. Without explicit data-fidelity constraints, they also risk altering pixel values at originally observed slices. The proposed framework addresses both limitations through null-space constrained learning and a locally consistent convolutional decoder.

2.3 Null-Range-Space Decomposition in Inverse Problems

Null-Range-Space Decomposition (NRSD) provides a mathematically rigorous framework for solving linear inverse problems by separating the solution into a measurement-consistent component (range space) and an unconstrained component (null space). DDNM [28] introduced this idea for zero-shot image restoration, refining null-space details with a pretrained diffusion model while projecting each iterate onto the data-consistency manifold. NPN [12] extends the framework with learned nonlinear null-space projections and task-specific regularization, providing convergence guarantees within the plug-and-play paradigm.

Chen *et al.* [6] demonstrate that bilinear interpolation serves as a valid generalized inverse for spectral demosaicing, enabling unsupervised NRSD without ground-truth supervision.

NRSD has been effective in natural image restoration and spectral imaging, but has not yet been applied to arbitrary medical slice SR. This work adopts NRSD as the theoretical basis of the DP-NSL framework, formulating arbitrary slice SR through a Measurement-Consistent Projection that combines the data-fidelity guarantee of NRSD with a geometry-aware Mixture-of-Splines prior for null-space synthesis.

3 Method

3.1 Problem Formulation and Overall Framework

Given an anisotropic 3D volume $\mathbf{V}_{LR} \in \mathbb{R}^{S \times H \times W}$ acquired with coarse through-plane resolution, the goal of arbitrary slice super-resolution is to reconstruct a high-resolution volume $\mathbf{V}_{SR} \in \mathbb{R}^{\hat{S} \times H \times W}$ at any upsampling factor R , under the supervision of the ground-truth volume $\mathbf{V}_{HR} \in \mathbb{R}^{\hat{S} \times H \times W}$. Let $\mathcal{D} : \mathbb{R}^{\hat{S} \times H \times W} \rightarrow \mathbb{R}^{S \times H \times W}$ denote the downsampling operator modeling the anisotropic slice acquisition. The degradation model is formulated as $\mathbf{V}_{LR} = \mathcal{D}(\mathbf{V}_{HR}) + \epsilon$, where ϵ represents acquisition noise.

Recovering arbitrary intermediate slices from severely downsampled observations is an ill-posed inverse problem. The unconstrained solution manifold is highly degenerate, and relying on networks to predict target slices without constraints risks hallucinations. The proposed DP-NSL framework therefore formulates the reconstruction as a dual-prior constrained process.

As illustrated in Fig. 2, the forward pipeline is driven by two parallel streams. A deterministic anchor $\mathbf{V}_{\mathcal{R}} = \mathcal{U}(\mathbf{V}_{LR})$ preserves the observed slices, where $\mathcal{U}$ is an upsampling pseudo-inverse. A prior-guided Null-Space Estimator (NSE) f_{θ} predicts the unobservable details $\hat{\mathbf{V}}_{\mathcal{N}} = f_{\theta}(\mathbf{V}_{LR})$. A Measurement-Consistent Projection (MCP) $\Pi_{\mathcal{N}}$ is then applied as a fixed projection, mapping this prediction to the null space to obtain the measurement-safe component $\mathbf{V}_{\mathcal{N}}$. The final super-resolved volume is formed by aggregating these two orthogonal components: $\mathbf{V}_{SR} = \mathbf{V}_{\mathcal{R}} + \mathbf{V}_{\mathcal{N}}$.

To regularize the prediction of $\hat{\mathbf{V}}_{\mathcal{N}}$, the NSE incorporates structural priors. It comprises three stages: a 3D encoder to extract spatial features $\mathbf{F}_{LR}$, a Mixture-of-Splines (MoS) module as an upsampler that enforces a continuous geometric prior, and a Local Spatial Consistency Decoder (LSCD) that provides local inductive bias for spatial coherence.

3.2 Deterministic Observation Prior via Orthogonal Projection

The first principle in medical image reconstruction is *data fidelity*: the super-resolved volume must strictly agree with the clinically acquired slices. However, existing methods [27,30] directly optimize an unconstrained mapping $\mathbb{R}^{S \times H \times W} \rightarrow$




Fig. 2: Overview of the proposed Dual-Prior Null-Space Learning (DP-NSL) framework. The reconstruction is decomposed into a deterministic range-space anchor $\mathbf{V}_R$, which preserves observed data, and a learnable null-space component $\mathbf{V}_N$. Within the Null-Space Estimator (NSE), the Mixture-of-Splines (MoS) module enforces spatially adaptive geometric continuity, while a Local Spatial Consistency Decoder (LSCD) injects local inductive bias.

$\mathbb{R}^{\hat{S} \times H \times W}$. Without mathematical constraints, the network is prone to perturbing the pixel values at the originally observed positions.

To enforce the proposed *Deterministic Observation Prior*, this unconstrained process is reformulated into an exact orthogonal projection constrained to the observation subspace via Null-Range-Space Decomposition [28]. Given the degradation operator $\mathcal{D}$, the constrained reconstruction requires $\mathcal{D}(\mathbf{V}_{SR}) = \mathbf{V}_{LR}$. Let $\mathcal{U} : \mathbb{R}^{S \times H \times W} \rightarrow \mathbb{R}^{\hat{S} \times H \times W}$ be a pseudo-inverse satisfying the consistency condition $\mathcal{D}\mathcal{U} = \mathbf{I}$. For any target-resolution volume $\mathbf{V}_{SR}$, applying the identity matrix yields an exact orthogonal decomposition:

$$\mathbf{V}_{SR} = \mathcal{U}\mathcal{D}(\mathbf{V}_{SR}) + (\mathbf{I} - \mathcal{U}\mathcal{D})\mathbf{V}_{SR} = \mathcal{U}(\mathbf{V}_{LR}) + (\mathbf{I} - \mathcal{U}\mathcal{D})\mathbf{V}_{SR}. \quad (1)$$

In this formulation, the first term $\mathcal{U}(\mathbf{V}_{LR})$ is deterministically measurable from the low-resolution input. Conversely, the second term resides in the null space of the degradation operator and involves the unobserved ideal full-resolution volume.

To resolve this missing information, the unobserved volume $\mathbf{V}_{SR}$ in the null-space term is approximated by the structural output of the NSE, denoted as $\hat{\mathbf{V}}_N = f_\theta(\mathbf{V}_{LR})$. Letting the Measurement-Consistent Projection (MCP) operator be defined as $\Pi_N \triangleq \mathbf{I} - \mathcal{U}\mathcal{D}$, a valid constrained solution manifold is established:

$$\mathbf{V}_{SR} = \underbrace{\mathcal{U}(\mathbf{V}_{LR})}_{\mathbf{V}_R: \text{deterministic anchor}} + \underbrace{\Pi_N(\hat{\mathbf{V}}_N)}_{\mathbf{V}_N: \text{null-space component}}. \quad (2)$$

Because $\mathcal{D} \Pi_{\mathcal{N}} = \mathcal{D} - \mathcal{D} \mathcal{U} \mathcal{D} = \mathcal{D} - \mathcal{D} = \mathbf{0}$, any signal in the range of $\Pi_{\mathcal{N}}$ is mathematically invisible to the forward operator $\mathcal{D}$. The range-space component $\mathbf{V}_{\mathcal{R}}$ is computed from observations and is the immutable anchor. The null-space component $\mathbf{V}_{\mathcal{N}} = (\mathbf{I} - \mathcal{U} \mathcal{D}) \hat{\mathbf{V}}_{\mathcal{N}}$ is confined to the null subspace of $\mathcal{D}$, within which the network synthesizes unobserved anatomical details.

This derivation provides a hard data-fidelity constraint. Regardless of how the network predicts $\hat{\mathbf{V}}_{\mathcal{N}}$, MCP confines all synthetic details to unobserved regions.

3.3 Prior-Guided Null-Space Refinement

While MCP locks the reconstruction to observed data, the unconstrained raw prediction $\hat{\mathbf{V}}_{\mathcal{N}}$ remains vulnerable to anatomical implausibility if learned blindly. To prevent this, the mapping from extracted features $\mathbf{F}_{LR}$ to the high-frequency unobserved details $\hat{\mathbf{V}}_{\mathcal{N}}$ is regularized by two complementary biases: a macroscopic geometric continuity prior and a microscopic local consistency prior.

Spatially Adaptive Geometric Prior via MoS Medical volumes exhibit heterogeneous smoothness, requiring high-order continuity for homogeneous soft tissues and sharp transitions at anatomical boundaries. Classical INR-based methods map coordinate domains via shared MLPs, imposing a uniform global prior that compromises either smooth regions or sharp edges. To address this, the proposed Mixture-of-Splines replaces unconstrained MLP regression with spatially adaptive geometric interpolation to assign coordinate-specific analytic continuity.

To model arbitrary-scale slice SR, let $\mathbf{F}_{LR} \in \mathbb{R}^{C \times S \times H \times W}$ denote low-resolution feature from a 3D encoder and $\mathbf{q} \in \mathbb{R}^3$ denote a target query coordinate (Coord in Fig. 2) in the high-resolution space. The target high-resolution feature mapping is defined as $\mathbf{F}_{SR}^{\mathbf{q}} = \text{MoS}(\mathbf{F}_{LR}, \mathbf{q})$.

To enforce structural consistency in the continuous spatial field, spatially adaptive B-spline bases are introduced as geometric priors rather than generic interpolators. A B-spline of order p inherently guarantees C^{p-1} analytic continuity [26]. For a specific query $\mathbf{q}$, rather than relying on an MLP, the continuous mapping is evaluated by a geometric spline expert $\mathcal{U}_p$. The upsampled feature evaluated at $\mathbf{q}$ is formulated as a separable tensor product:

$$\mathcal{U}_p(\mathbf{F}_{LR}; \mathbf{q}) = \mathbf{c} \odot (\mathbf{b}_x \otimes \mathbf{b}_y \otimes \mathbf{b}_z). \quad (3)$$

Here, $\otimes$ denotes the standard tensor product constructing the M^3 -dimensional separable 3D basis, and $\odot$ is element-wise multiplication. The 1D B-spline basis vector along spatial axis $d \in \{x, y, z\}$ is computed as $\mathbf{b}_d = \beta^p((\Delta q_d - \mathbf{k}_d) \odot \mathbf{s}_d)$. The relative projection distance $\Delta \mathbf{q} = \mathbf{q} - \mathbf{q}_{LR} + \mathbf{o}$ incorporates a learned subvoxel offset $\mathbf{o} \in \mathbb{R}^3$ for grid misalignment compensation, where $\mathbf{q}_{LR}$ is the nearest counterpart in the LR grid. The geometric variables are projected from the immediate feature neighborhood surrounding $\mathbf{q}_{LR}$ through shallow convolution layers: structurally-aware scaling coefficients $\mathbf{c} \in \mathbb{R}^{M^3}$, adaptive knots $\mathbf{k}_d \in \mathbb{R}^M$,

and continuous dilations $\mathbf{s}_d \in \mathbb{R}^M$ projected from a step size (Cell in Fig. 2) of the HR grid.

Instead of rigidly committing to a single uniform continuous order, MoS dynamically allocates spatially adaptive prior stiffness. A bank of K spline experts $\{\mathcal{U}_{p_k}\}_{k=1}^K$ is instantiated with distinct analytic orders p_k . Simultaneously, a lightweight routing network calculates the mixing weights $\boldsymbol{\pi}(\mathbf{q}) \in \mathbb{R}^K$ from the locally interpolated features $\tilde{\mathbf{F}}_{LR}^{\mathbf{q}}$:

$$\boldsymbol{\pi}(\mathbf{q}) = \text{Softmax}(\mathbf{R}_{\pi}(\tilde{\mathbf{F}}_{LR}^{\mathbf{q}})). \quad (4)$$

The final high-resolution feature at $\mathbf{q}$ is a weighted mixture of the geometric priors:

$$\mathbf{F}_{SR}^{\mathbf{q}} = \sum_{k=1}^K \pi_k(\mathbf{q}) \mathcal{U}_{p_k}(\mathbf{F}_{LR}; \mathbf{q}). \quad (5)$$

MoS thus produces a content-aware continuous spatial representation with locally adapted geometric priors.

Local Consistency Prior via LSCD By evaluating $\mathbf{F}_{SR}^{\mathbf{q}}$ over the target high-resolution coordinate, a high-resolution dense feature volume $\mathbf{F}_{SR} \in \mathbb{R}^{C \times S \times H \times W}$ is finally formed. Typical point-wise MLPs process each coordinate independently when mapping this spatial field to the raw unobserved details $\hat{\mathbf{V}}_{\mathcal{N}}$, which can induce micro-level local inconsistencies. To address this and supplement the macroscopic MoS prior, a multi-scale Local Spatial Consistency Decoder is introduced to inject a local inductive bias.

Based on an Inception-style split-transform-merge mechanism [23, 24], $\mathbf{F}_{SR}$ is partitioned along the channel dimension into parallel sub-groups: $\mathbf{f}_{\text{id}}$, $\mathbf{f}_{\text{conv}}$, $\mathbf{f}_1$, $\mathbf{f}_2$, and $\mathbf{f}_3$. The identity group $\mathbf{f}_{\text{id}}$ skips computation to preserve the high-frequency base signal. The $\mathbf{f}_{\text{conv}}$ branch operates as a standard 3D convolution, while the remaining sibling groups are processed through depthwise 3D convolutions with monotonically increasing kernel sizes to capture multi-scale spatial context:

$$\mathbf{F}_{\text{out}} = [\mathbf{f}_{\text{id}}, \mathcal{W}_3(\mathbf{f}_{\text{conv}}), \mathcal{W}_3^{\text{DW}}(\mathbf{f}_1), \mathcal{W}_5^{\text{DW}}(\mathbf{f}_2), \mathcal{W}_7^{\text{DW}}(\mathbf{f}_3)], \quad (6)$$

where $[\cdot]$ denotes concatenation along the channel dimension, $\mathcal{W}_k$ represents a $k \times k \times k$ standard convolution, and $\mathcal{W}_k^{\text{DW}}$ denotes a depthwise spatial convolution. After cascading multiple such blocks, a final projection collapses the refined multi-scale feature maps to the raw unobservable signal $\hat{\mathbf{V}}_{\mathcal{N}}$, subsequently bounded by the hard MCP constraint.

4 Experiments

4.1 Experimental Setup

Datasets. Evaluation is conducted on both CT and MRI modalities to test cross-modality generalization. For CT, 3D volumes are collected from the Medical Segmentation Decathlon (MSD) [2]: Colon (190 volumes), Liver (201 volumes), and Hepatic Vessels (433 volumes), all exhibiting clinically anisotropic

Table 1: Quantitative comparison at in-scales ($\times 2$, $\times 3$, $\times 4$) on four datasets. Best results in **bold**, second best underlined.

Method	Colon			Liver		
	$\times 2$	$\times 3$	$\times 4$	$\times 2$	$\times 3$	$\times 4$
EDSR3D	<u>41.84/0.9809</u>	37.71/0.9600	<u>35.37/0.9421</u>	40.62/0.9733	37.02/0.9507	<u>34.86/0.9307</u>
MetaSR	41.26/0.9790	37.22/0.9567	34.81/0.9361	40.76/0.9734	36.61/0.9461	34.25/0.9220
LTE	41.72/0.9806	37.62/0.9597	35.15/0.9399	41.40/0.9766	<u>37.15/0.9512</u>	34.71/0.9281
HIIF	<u>41.83/0.9810</u>	<u>37.75/0.9604</u>	35.26/0.9411	<u>41.42/0.9768</u>	<u>37.15/0.9512</u>	34.69/0.9280
ArSSR	40.67/0.9771	36.79/0.9534	34.24/0.9293	40.90/0.9737	36.69/0.9452	34.18/0.9190
SAINR	41.48/0.9798	<u>37.78/0.9603</u>	35.30/0.9408	41.32/0.9763	37.14/0.9509	34.68/0.9273
CycleINR	41.35/0.9794	37.38/0.9574	34.87/0.9360	40.88/0.9732	36.70/0.9445	34.33/0.9194
DC2SR	41.29/0.9794	37.40/0.9580	34.94/0.9374	41.18/0.9763	36.96/0.9501	34.50/0.9261
DP-NSL	42.55/0.9832	38.43/0.9648	35.95/0.9472	42.49/0.9801	38.06/0.9572	35.47/0.9351
Method	Hepatic Vessels			IXI		
	$\times 2$	$\times 3$	$\times 4$	$\times 2$	$\times 3$	$\times 4$
EDSR3D	<u>43.25/0.9847</u>	<u>38.97/0.9671</u>	<u>36.71/0.9520</u>	47.08/0.9902	43.45/0.9788	<u>41.47/0.9679</u>
MetaSR	42.04/0.9813	37.87/0.9605	35.63/0.9428	46.67/0.9893	42.87/0.9759	40.75/0.9625
LTE	42.64/0.9832	38.43/0.9640	36.13/0.9472	47.11/ <u>0.9903</u>	43.40/0.9786	41.33/0.9669
HIIF	43.17/0.9846	38.93/0.9670	36.68/0.9518	<u>47.12/0.9903</u>	43.44/0.9787	41.38/0.9672
ArSSR	41.89/0.9812	37.85/0.9605	35.57/0.9420	46.51/0.9888	42.70/0.9748	40.45/0.9600
SAINR	42.68/0.9833	38.79/0.9661	36.52/0.9504	46.47/0.9887	43.20/0.9775	41.20/0.9658
CycleINR	42.33/0.9826	38.33/0.9636	36.05/0.9464	46.75/0.9895	43.16/0.9773	41.04/0.9646
DC2SR	42.64/0.9832	38.54/0.9646	36.23/0.9478	47.09/0.9902	<u>43.52/0.9791</u>	41.41/0.9675
DP-NSL	43.40/0.9854	39.23/0.9690	36.93/0.9543	47.54/0.9912	43.93/0.9810	41.92/0.9710

spacing. For MRI, 185 T1-weighted brain scans acquired at Hammersmith Hospital are selected from the IXI dataset [1]. Each dataset is split into training and test subsets following [21, 22]. Low-resolution inputs are synthesized by directly downsampling the high-resolution volume by a factor R along the slice axis while preserving in-plane resolution.

Implementation Details. The encoder adopts a 3D EDSR [17]. The MoS module instantiates three B-spline experts (orders $p \in \{2, 3, 4\}$). All models are optimized with Adam [15] ($\beta_1 = 0.9$, $\beta_2 = 0.999$) and an ℓ_1 reconstruction loss on NVIDIA RTX 3090 GPUs. The initial learning rate is 1×10^{-4} , halved every 200 epochs over a total of 1000 epochs, with a batch size of 8. During training, each volume is center-cropped to 256×256 in the axial plane, and a training sample consists of four consecutive LR slices randomly extracted along the slice axis.

Compared Methods. The proposed method is compared against approaches from three categories: (i) fixed-scale volumetric SR: EDSR3D [17], which trains a specific model for each factor; (ii) generic ASSR: MetaSR [11], LTE [16], and HIIF [13]; and (iii) medical ASSR: ArSSR [30], SAINR [27], CycleINR [7], and DC²SR [32]. All arbitrary-scale methods adopt the same encoder EDSR3D and are jointly trained on $R \in \{2, 3, 4\}$. Evaluation covers in-scales ($R \in \{2, 3, 4\}$) and out-of-scales ($R \in \{5, 6, 7\}$), which tests generalization to factors never encountered during training. The reconstruction performance is evaluated using the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index (SSIM).

Table 2: Quantitative comparison at out-of-scales ($\times 5$, $\times 6$, $\times 7$) on four datasets. All methods are trained only on $\times 2$, $\times 3$, $\times 4$. Best in **bold**, second best underlined.

Method	Colon			Liver		
	$\times 5$	$\times 6$	$\times 7$	$\times 5$	$\times 6$	$\times 7$
MetaSR	33.10/0.9145	32.07/0.9012	31.11/0.8879	32.71/0.9006	31.68/0.8821	31.02/0.8723
LTE	33.44/0.9204	<u>32.33/0.9066</u>	<u>31.34/0.8931</u>	<u>33.11/0.9075</u>	<u>32.13/0.8928</u>	<u>31.39/0.8812</u>
HIIF	31.59/0.8537	30.72/0.8461	30.89/0.8716	31.14/0.8202	30.57/0.8213	30.85/0.8564
ArSSR	32.78/0.9114	31.66/0.8959	30.74/0.8830	32.70/0.8993	31.61/0.8831	30.92/0.8717
SAINR	<u>33.55/0.9223</u>	<u>32.33/0.9070</u>	31.27/0.8929	33.07/ <u>0.9077</u>	31.96/0.8923	31.22/0.8810
CycleINR	32.88/0.9119	31.80/0.8972	30.87/0.8847	32.72/0.8976	31.53/0.8824	30.92/0.8718
DC2SR	33.16/0.9168	31.91/0.9001	30.93/0.8870	32.88/0.9051	31.90/0.8888	31.16/0.8772
DP-NSL	34.05/0.9286	32.80/0.9141	31.76/0.9008	33.72/0.9167	32.55/0.9011	31.68/0.8885
Method	Hepatic Vessels			IXI		
	$\times 5$	$\times 6$	$\times 7$	$\times 5$	$\times 6$	$\times 7$
MetaSR	33.97/0.9240	32.95/0.9121	32.17/0.9016	38.97/0.9461	37.75/0.9318	36.84/0.9192
LTE	34.44/0.9303	33.28/0.9178	32.47/0.9070	39.52/0.9516	<u>38.27/0.9383</u>	<u>37.32/0.9262</u>
HIIF	32.12/0.8522	31.27/0.8408	32.04/0.8809	38.91/0.9448	37.53/0.9241	37.05/0.9197
ArSSR	33.98/0.9257	32.79/0.9115	32.03/0.9014	38.81/0.9446	37.51/0.9291	36.61/0.9166
SAINR	<u>34.76/0.9345</u>	<u>33.46/0.9203</u>	<u>32.58/0.9095</u>	<u>39.55/0.9525</u>	<u>38.22/0.9387</u>	<u>37.26/0.9268</u>
CycleINR	33.93/0.9242	32.81/0.9107	32.10/0.9019	38.98/0.9460	37.75/0.9321	36.89/0.9204
DC2SR	34.36/0.9299	32.97/0.9139	32.20/0.9044	39.53/0.9524	38.10/0.9372	37.19/0.9258
DP-NSL	35.20/0.9393	33.98/0.9260	33.04/0.9146	40.17/0.9578	38.80/0.9442	37.69/0.9303

Table 3: Results ($\times 5$) on the RPLHR-CT real dataset.

Metric	EDSR3D	MetaSR	LTE	HIIF	ArSSR	SAINR	CycleINR	DC2SR	DP-NSL
PSNR	29.63	30.92	<u>31.20</u>	21.83	28.06	30.66	29.07	27.44	34.91
SSIM	0.8342	0.8444	<u>0.8505</u>	0.7187	0.8093	0.8384	0.8205	0.7926	0.8925

4.2 Comparison with State-of-the-Art

In-Scales. Table 1 reports PSNR and SSIM at in-scales across all datasets. The proposed DP-NSL consistently achieves the highest performance. On Colon, improvements over the second-best method reach 0.71 dB at $\times 2$ and 0.65 dB at $\times 3$; on Liver, the margin widens to 1.07 dB at $\times 2$. Hepatic Vessels dataset benefits from the spatially adaptive MoS prior, yielding 0.26 dB at $\times 3$. On the IXI brain MRI benchmark, a gain of 0.42 dB at $\times 2$ and 0.45 dB at $\times 4$ shows that the dual-prior formulation transfers across modalities. EDSR3D, trained separately at each scale, remains competitive at individual factors but cannot generalize to other scales. INR-based medical methods such as SAINR and DC²SR perform favorably but lack explicit data-fidelity constraints, leaving observed slices vulnerable to subtle alteration.

Out-of-Scales. Table 2 evaluates generalization at unseen factors $\times 5$, $\times 6$, and $\times 7$. DP-NSL maintains a consistent lead over all baselines across all datasets. On Hepatic Vessels, improvements exceed 0.44 dB at every unseen factor, with the gap widening at the most challenging $\times 7$. On IXI, the margin is up to 0.62 dB at $\times 5$. HIIF, despite strong in-scales results, suffers severe degradation at unseen scales. In contrast, DP-NSL shows a stable geometric prior: the inherent conti-




Fig. 3: Visualization on the Liver dataset (sagittal view, $\times 3$). The sagittal plane directly reveals through-plane reconstruction quality.

Table 4: Downstream segmentation results (Dice / PSNR) on KiTS19.

Scale	MetaSR	LTE	HIIF	ArSSR	SAINR	CycleINR	DC2SR	DP-NSL
$\times 3$	0.8562/39.77	0.8566/40.03	0.8570/40.31	0.8549/39.95	0.8569/39.68	0.8561/40.15	0.8569/40.04	0.8581/40.28
$\times 6$	0.7820/34.88	0.7891/35.08	0.7044/31.50	0.7634/34.86	0.7922/35.08	0.7820/34.98	0.7794/35.05	0.8067/35.13

nuity of B-spline bases holds regardless of sampling density, allowing principled extrapolation to unseen upsampling factors.

Real-Data Validation. Further experiments on the paired real dataset RPLHR-CT [31] test practical validity of the method. As shown in Table 3, DP-NSL achieves 34.91 dB PSNR at $\times 5$.

Qualitative Comparison. Fig. 3 presents sagittal-view comparisons on the Liver dataset at $\times 3$. It directly exposes through-plane reconstruction quality. Other methods produce visibly blurred organ boundaries and fail to separate intervertebral spaces clearly. By contrast, the proposed DP-NSL recovers smooth, anatomically coherent contours. Fig. 4 shows axial-view comparisons on the IXI brain MRI dataset at $\times 5$. Cortical folding patterns and white-matter boundaries pose particular challenges due to their convoluted geometry. While other methods blur sulcal regions, DP-NSL reconstructs these complex structures with less error, demonstrating that the locally adaptive geometric prior generalizes to highly textured brain MRI.

Downstream Tasks. Table 4 reports downstream segmentation results on KiTS19 [10]. DP-NSL achieves the highest Dice score at both $\times 3$ and $\times 6$, indicating that it preserves the anatomical details for downstream analysis.

4.3 Ablation Studies

Ablation experiments validate the contribution of each proposed component.

Component Analysis. Table 5 examines the contribution of each core module: MCP, MoS, and LSCD, with PSNR and SSIM. Without the proposed modules, the baseline degenerates to direct estimation with trilinear upsampling and EDSR, performing unconstrained synthesis. Activating MCP alone improves PSNR by 0.31 dB at $\times 2$, showing that a measurement consistency anchor already provides a noticeable boost without architectural changes. Replacing the trilinear upsampler with MoS evaluates the macroscopic geometric prior, improving




Fig. 4: Visualization on the IXI brain MRI dataset (axial view, $\times 5$) with error maps.

Table 5: Component ablation on the Colon dataset with PSNR and SSIM.

MCP	MoS	LSCD	$\times 2$ (mean $\pm$ std)	$\times 3$ (mean $\pm$ std)	$\times 4$ (mean $\pm$ std)	$\times 5$ (mean $\pm$ std)
$\times$	$\times$	$\times$	41.82 $\pm$ 0.02	37.76 $\pm$ 0.03	35.30 $\pm$ 0.02	33.64 $\pm$ 0.01
$\checkmark$	$\times$	$\times$	42.13 $\pm$ 0.06	38.03 $\pm$ 0.06	35.52 $\pm$ 0.06	33.80 $\pm$ 0.04
$\times$	$\checkmark$	$\times$	42.45 $\pm$ 0.02	38.32 $\pm$ 0.03	35.82 $\pm$ 0.04	33.96 $\pm$ 0.08
$\times$	$\times$	$\checkmark$	42.06 $\pm$ 0.01	37.95 $\pm$ 0.02	35.43 $\pm$ 0.01	33.67 $\pm$ 0.02
$\checkmark$	$\checkmark$	$\times$	42.52 $\pm$ 0.01	38.40 $\pm$ 0.01	35.89 $\pm$ 0.01	34.08 $\pm$ 0.02
$\checkmark$	$\checkmark$	$\checkmark$	42.54 $\pm$ 0.03	38.42 $\pm$ 0.02	35.93 $\pm$ 0.03	34.12 $\pm$ 0.03
$\times$	$\times$	$\times$	0.9810 $\pm$ 0.0001	0.9605 $\pm$ 0.0002	0.9412 $\pm$ 0.0002	0.9232 $\pm$ 0.0004
$\checkmark$	$\times$	$\times$	0.9821 $\pm$ 0.0002	0.9628 $\pm$ 0.0003	0.9440 $\pm$ 0.0004	0.9261 $\pm$ 0.0004
$\times$	$\checkmark$	$\times$	0.9829 $\pm$ 0.0001	0.9641 $\pm$ 0.0001	0.9461 $\pm$ 0.0003	0.9269 $\pm$ 0.0010
$\times$	$\times$	$\checkmark$	0.9817 $\pm$ 0.0001	0.9618 $\pm$ 0.0003	0.9420 $\pm$ 0.0003	0.9227 $\pm$ 0.0004
$\checkmark$	$\checkmark$	$\times$	0.9831 $\pm$ 0.0001	0.9648 $\pm$ 0.0002	0.9470 $\pm$ 0.0003	0.9290 $\pm$ 0.0004
$\checkmark$	$\checkmark$	$\checkmark$	0.9832 $\pm$ 0.0001	0.9648 $\pm$ 0.0002	0.9470 $\pm$ 0.0003	0.9290 $\pm$ 0.0003

PSNR by 0.63 dB at $\times 2$. Combining MoS with MCP can further improve performance. This supports the hypothesis that, while MoS provides strong detail generation, the null-space projection confines the generated details and prevents them from overriding known geometry. Finally, the LSCD adds the local inductive bias. The complete DP-NSL achieves the best performance by satisfying both data fidelity and spatial continuity constraints.

Efficacy of the Multi-Order Geometric Prior. Table 6 investigates the impact of B-spline order within the MoS. Single-order configurations perform comparably, yielding incremental improvements in reconstruction fidelity from order 1 to 3; however, order 4 offers no additional in-scales benefit and slightly degrades the out-of-scales results. Multi-order MoS variants that combine spline experts of different orders consistently outperform the single-order variants, supporting the premise that anatomical regions with heterogeneous smoothness require spatially varying geometric priors. The combination of orders $p \in \{2, 3, 4\}$ achieves the best in-scales accuracy, while including order 4 is marginally superior




Fig. 5: Visualization of MoS routing weights on the CT slice. Frequency maps show the spatial variation intensity of the slice. Expert n shows spatial activation of the B-spline Upsampler with order $p = n$.

Table 6: Ablation study on B-spline order combinations in the MoS module. $p = n$ denotes the expert of order n .

Orders of MoS	PSNR			
	$\times 2$	$\times 3$	$\times 4$	$\times 5$
$p = 1$	42.16	38.09	35.59	33.80
$p = 2$	42.22	38.12	35.62	33.82
$p = 3$	42.23	38.14	35.65	33.81
$p = 4$	42.23	38.14	35.64	33.73
$p \in \{2, 3, 4\}$	42.39	38.30	35.80	33.94
$p \in \{1, 2, 3\}$	42.40	38.29	35.79	33.91
$p \in \{1, 2, 3, 4\}$	42.32	38.23	35.75	33.97

at $\times 5$. The tri-order ensemble is therefore adopted as the default configuration, balancing boundary preservation with high-order smoothness.

Interplay of Range and Null-Space Operators. Table 7 dissects the sampling operator choices across three progressive paradigms evaluated via the Range Anchor $\mathcal{U}$ and MCP projection $\Pi_{\mathcal{N}}$. The Range Anchor $\mathcal{U}$ provides the deterministic low-frequency scaffold, while the MCP projection $\Pi_{\mathcal{N}}$ bounds the high-frequency synthesis within the null space. Without both components, the framework reduces to a pure black-box estimator that must synthesize the low-frequency topology from scratch, undermining stability. Introducing only a trilinear range anchor $\mathcal{U}$ converts the model to a residual-based architecture. This secures a reliable low-frequency scaffold, yet performance remains sub-optimal due to unconstrained high-frequency reconstruction. Adding the MCP projection $\Pi_{\mathcal{N}}$ completes the NRSD framework. Within this complete formulation, pairing a trilinear anchor with a zero-padding projection proves optimal. This confirms that well-paired orthogonal operators are needed to balance measurement fidelity with null-space expressiveness.

4.4 Analysis and Discussion

Expert Routing Interpretation. Fig. 5 visualizes per-expert routing weight maps on a representative CT slice to investigate whether the MoS router intrinsically learns semantically meaningful allocations. Two distinct spatial be-




Fig. 6: Joint visualization of MCP behavior. (a) and (b) show the decomposition into range-space anchor and null-space component respectively. (c) reports slice-wise MAE versus slice index.

Table 7: Ablation study on the null-range-space operators. “-” indicates the corresponding component is omitted.

Anchor $\mathcal{U}$	Projection $\Pi_{\mathcal{N}}$	PSNR			
		$\times 2$	$\times 3$	$\times 4$	$\times 5$
-	-	41.45	37.35	34.91	33.31
Trilinear	-	41.76	37.68	35.21	33.54
Zero-Pad	-	41.11	37.08	34.69	33.56
Zero-Pad	Zero-Pad	41.89	37.78	35.28	33.63
Trilinear	Trilinear	41.83	36.92	34.60	32.41
Trilinear	Zero-Pad	42.06	37.96	35.45	33.76

haviors emerge. The low-order expert is more active in homogeneous soft-tissue interiors, providing baseline approximations across smooth anatomical regions. Higher-order experts concentrate on complex structural sites such as fine vascular branches and sharp organ boundaries, capturing high-frequency structural details. This content-aware routing emerges from data-driven optimization alone, without explicit boundary supervision, indicating that the optimal geometric prior varies by region. Transitional zones show more balanced expert mixtures, where the router blends geometric priors to match the local structural complexity.

Measurement Consistency Verification. A central claim of the DP-NSL framework is that MCP guarantees exact data fidelity while confining all learned high-frequency content to the null space. Fig. 6 jointly visualizes this decomposition and its quantitative implications.

Fig. 6(a) depicts the range-space component $\mathbf{V}_{\mathcal{R}} = \mathcal{U}(\mathbf{V}_{LR})$, the deterministic anchor obtained via trilinear upsampling. This component reproduces the coarse anatomical layout and faithfully retains all originally observed slice intensities. Fig. 6(b) displays the corresponding null-space component $\mathbf{V}_{\mathcal{N}}$. Safeguarded by MCP, the NSE synthesizes high-frequency anatomical details without generating any superfluous content at the observed slice positions.

Fig. 6(c) quantifies this behavior through per-slice Mean Absolute Error (MAE) along the slice axis. Because the observation constraint is enforced, the MAE drops to zero at the originally acquired positions ($s = 0, R, 2R, \dots$). This periodic zero-error pattern proves that $\mathcal{D}(\mathbf{V}_{SR}) = \mathbf{V}_{LR}$ holds identically. By contrast, baseline methods lacking MCP operate as unconstrained direct estimation and residual-based estimation: they exhibit non-negligible errors even at inherently known positions. Such unconstrained behavior subtly alters the original clinical data, presenting an unacceptable risk in diagnostic workflows.

5 Conclusion

This paper formulates arbitrary medical slice super-resolution as a constrained inverse problem and presents the DP-NSL framework, which grounds the reconstruction in two complementary priors rather than unconstrained residual-based regression. The *Deterministic Observation Prior*, enforced by MCP, mathematically guarantees that every originally acquired clinical slice remains unchanged, separating the observable range space from the learnable null space. The *Geometric Continuity Prior*, instantiated by MoS, regularizes the null-space synthesis with spatially adaptive B-spline bases whose analytic orders adapt to local anatomical complexity. Together with the LSCD for local spatial coherence, these two priors turn an ill-posed reconstruction problem into a well-constrained process reconciling data fidelity with anatomically plausible detail synthesis. Experiments across three CT and one MRI datasets confirm consistent improvements at both in-scale and out-of-scale factors. Future work can explore adapting the MCP formulation to more realistic degradation scenarios.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62471182), Science and Technology Commission of Shanghai Municipality Basic Research Program (Grant No. 25JD1401300), Shanghai Rising-Star Program (Grant No. 24QA2702100), and the Science and Technology Commission of Shanghai Municipality (Grant No. 22DZ2229004)

References

1. Ixi dataset. <https://brain-development.org/ixi-dataset/>
2. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature Communications* **13**(1), 4128 (2022)
3. Birkfellner, W.: *Applied Medical Image Processing: A Basic Course*. CRC Press (2014)
4. Cao, J., Wang, Q., Xian, Y., Van Gool, L., Timofte, R.: Ciaosr: Continuous implicit attention-in-attention network for arbitrary-scale image super-resolution. In: *CVPR* (2023)

5. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: CVPR (2021)
6. Chen, Y., Wang, Y., Zhang, H.: Unsupervised range-nullspace learning prior for multispectral images reconstruction. IEEE TIP (2025)
7. Fang, W., Tang, Y., Xu, M., Pan, J., Cai, J.: Cycleinr: Cycle implicit neural representation for arbitrary-scale volumetric super-resolution of medical data. In: CVPR (2024)
8. Gerig, G., Kubler, O., Kikinis, R., Jolesz, F.A.: Nonlinear anisotropic filtering of mri data. IEEE Trans. Med. Imaging **11**(2), 221–232 (1992)
9. He, Z., Jin, Z.: Latent modulated function for computational optimal continuous image representation. In: CVPR (2024)
10. Heller, N., Sathianathan, N., Kalapara, A., Walczak, E., Moore, K., Kaluzniak, H., Rosenberg, J., Blake, P., Rengel, Z., Oestreich, M., et al.: The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. arXiv preprint arXiv:1904.00445 (2019)
11. Hu, X., Mu, H., Zhang, X., Wang, Z., Tan, T., Sun, J.: Meta-sr: A magnification-arbitrary network for super-resolution. In: CVPR (2019)
12. Jacome, R., Gualdrón-Hurtado, R., Arguello, H.: Npn: Non-linear projections of the null-space for imaging inverse problems. In: NeurIPS (2025)
13. Jiang, Y., Kwan, H.M., Zhang, F., Bull, D.: Hiif: Hierarchical encoding based implicit image function for continuous super-resolution. In: CVPR (2025)
14. Kim, E., Kim, H., Yoo, J.: Bf-stvsr: B-splines and fourier – best friends for high fidelity spatial-temporal video super-resolution. In: CVPR (2025)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
16. Lee, J., Jin, K.H.: Local texture estimator for implicit representation function. In: CVPR (2022)
17. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: CVPRW (2017)
18. Liu, X., Tang, H.: Diffno: Diffusion fourier neural operator. In: CVPR. pp. 150–160 (2025)
19. Luo, X., Qian, X., Yoon, B.J.: Hierarchical neural operator transformer with learnable frequency-aware loss prior for arbitrary-scale super-resolution. In: ICML (2024)
20. Pak, B., Lee, J., Jin, K.H.: B-spline texture coefficients estimator for screen content image super-resolution. In: CVPR (2023)
21. Peng, C., Lin, W.A., Liao, H., Chellappa, R., Zhou, S.K.: Saint: spatially aware interpolation network for medical slice synthesis. In: CVPR (2020)
22. Song, H., Mao, X., Wang, Y., Shen, W., Li, Q.: I3net: Inter-intra-slice interpolation network for medical slice synthesis. IEEE TMI (2024)
23. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR (2015)
24. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)
25. Uhm, K.H., Cho, H., Hong, S.H., Jung, S.W.: An anisotropic cross-view texture transfer with multi-reference non-local attention for ct slice interpolation. IEEE Trans. Med. Imaging (2025)
26. Wang, P.S., Liu, Y., Yang, Y.Q., Tong, X.: Spline positional encoding for learning 3d implicit signed distance fields. In: IJCAI (2021)

27. Wang, X., Wang, S., Xiong, H., Xuan, K., Zhuang, Z., Liu, M., Shen, Z., Zhao, X., Zhang, L., Wang, Q.: Spatial attention-based implicit neural representation for arbitrary reduction of mri slice spacing. *Medical Image Analysis* **94**, 103158 (2024)
28. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. In: *ICLR* (2023)
29. Wei, M., Zhang, X.: Super-resolution neural operator. In: *CVPR*. pp. 18247–18256 (2023)
30. Wu, Q., Li, Y., Sun, Y., Zhou, Y., Wei, H., Yu, J., Zhang, Y.: An arbitrary scale super-resolution approach for 3d mr images via implicit neural representation. *IEEE JBHI* **27**(2), 1004–1015 (2022)
31. Yu, P., Zhang, H., Kang, H., Tang, W., Arnold, C.W., Zhang, R.: Rplhr-ct dataset and transformer baseline for volumetric super-resolution from ct scans. In: *MICCAI* (2022)
32. Zeng, C., Zhang, Z., Huang, W., Zhang, L.: Dc2-sr: A dual-consistency guided curriculum learning method for thick-slice fetal mri super-resolution. In: *ACM MM* (2025)