DIAGNOSE BEFORE REFINING: GROUND-TRUTH-FREE SELF-DIAGNOSTIC REFINEMENT FOR IMPLICIT NEURAL REPRESENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Existing INR reconstruction methods typically terminate once optimization converges, even though non-negligible errors may remain at unobserved entries where ground truth is unavailable. This motivates us to investigate whether the remaining errors can be identified and further reduced without access to their ground truth. By the analysis of the post-reconstruction errors, we find that they are often concentrated and exhibit different optimization behaviors, indicating that uniform refinement is insufficient. Motivated by these observations, we propose SDR, a ground-truth-free self-diagnostic refinement framework that exploits out-of-fold (O-O-F) reconstruction to diagnose and adaptively refine the remaining errors. SDR uses cross-validation uncertainty to guide gated global optimization, and subsequently transfers structurally matched signed O-O-F residuals to correct local errors that are no longer effectively reduced by training. We further provide theoretical characterization for residual correction, establishing a sufficient condition for guaranteed error reduction and deriving an estimation-error bound that explicitly characterizes the roles of structural similarity, fold-to-full discrepancy, and residual uncertainty. Extensive experiments across diverse signal modalities, reconstruction methods, and multiple seeds demonstrate consistent and robust superior performance over the basic method, with mechanism, ablation, and sensitivity analysis further validating the consistent effectiveness of the proposed framework.

1 INTRODUCTION AND RELATED WORK

Multidimensional visual signals, such as color images, videos, and multispectral/hyperspectral images, are widely used in medical imaging (Liu et al., 2026; 2025), remote sensing (Wang et al., 2023a), data compression (Liu et al., 2023), and other applications (Qin et al., 2022a; Wang et al., 2021a; Qin et al., 2022b; Wang et al., 2022; 2020; Zhang et al., 2020a). Low-rank tensor models have long been employed for missing-data reconstruction by exploiting correlations across multiple modes. Classical decompositions such as CP Carroll & Chang (1970), Tucker De Lathauwer et al. (2000), and T-SVD (Wang et al., 2023b; Wu et al., 2026; Kong et al., 2025), together with tensor-network models including TT Oseledets (2011) and TR (Zhao et al., 2016; Qin et al., 2025), provide compact low-rank representations for multidimensional data and have been widely studied for visual signal completion (Zhang et al., 2020b; Hou et al., 2021; 2020; Qin et al., 2024). More recently, implicit neural representations (INRs) model signals as continuous coordinate-based functions and have shown strong representation capability in various inverse problems (Luo et al., 2023; 2022; Xu et al., 2026; Saragadam et al., 2023). Despite their different formulations, existing reconstruction methods mainly focus on improving the reconstruction itself and typically terminate once optimization converges. Non-negligible errors may nevertheless remain at unobserved entries, motivating recent efforts to assess or further refine reconstruction without access to ground truth.

Several studies have explored reconstruction assessment and refinement without ground truth. Noise2Self (Batson & Royer, 2019) exploits held-out observations to construct self-supervised supervision for denoising, showing that useful reconstruction cues can be obtained without ground truth. For deep image prior (DIP), prior studies have analyzed the degradation that may occur during prolonged optimization (Shi et al., 2022) and developed GT-free early-stopping strategies to identify favorable optimization states (Wang et al., 2021b). Beyond assessment and stopping, fur-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

ther refinement has also been investigated. I-INR (Haider et al., 2026) introduces a plug-and-play iterative refinement strategy for INR architectures, while RAR (Chen et al., 2026) combines restoration with quality assessment to decide whether additional restoration is needed. Residual-learning approaches further improve an initial reconstruction by explicitly estimating its residual correction (Jiao et al., 2017). These studies demonstrate that reconstruction quality can still be assessed or improved without direct access to ground truth. However, they mainly focus on global reconstruction states, stopping decisions, or predefined refinement procedures, while the spatial distribution, different behavior, and refinement potential of the remaining post-reconstruction errors remain less explored.

To better understand what remains after reconstruction, we analyze post-reconstruction errors across different signals and methods. We find that they are often concentrated and exhibit different optimization behaviors: some remain responsive to further optimization, whereas others become locally stable. This suggests that post-reconstruction refinement should first diagnose where the remaining errors lie and how they behave, rather than treating them uniformly. Based on this observation, we propose SDR, a ground-truth-free Self-Diagnostic Refinement framework. SDR uses out-of-fold (O-O-F) reconstructions to localize unreliable regions through cross-validation uncertainty and to gate further global optimization according to its estimated benefit. After gated optimization, signed O-O-F residuals are used as a residual memory and structurally transferred to correct persistent local errors. Thus, O-O-F reconstruction supports both self-diagnosis and subsequent refinement without requiring ground truth at missing entries.Overall, the key contributions of this paper can be summarized as:

- We systematically investigate post-reconstruction errors from a self-diagnostic perspective, characterizing their non-uniform spatial distribution and different optimization behaviors, and showing how these properties can guide targeted post-reconstruction refinement.
- We propose SDR, a ground-truth-free self-diagnostic refinement framework that exploits O-O-F reconstruction for error localization and adaptive optimization gating, and further leverages structural residual transfer to correct persistent errors, enabling post-reconstruction refinement without access to the ground truth of missing entries.
- We further provide theoretical analysis of SDR, establishing a sufficient condition for guaranteed error reduction and deriving an estimation-error bound for residual transfer. Extensive experiments across different signal modalities, reconstruction methods, sampling ratios and multiple seeds validate the effectiveness and robustness of the proposed framework.

2 Cross-Validated Error Diagnosis

2.1 Problem Setup

Let $X^\star$ denote an unknown complete multidimensional visual signal, such as an RGB image, video, or multispectral/hyperspectral image, with entries indexed by p. Let M be the observation mask, where $M(p) = 1$ for observed entries and $M(p) = 0$ otherwise. Given the observed values Y on Ω, an INR f_θ parameterized by θ maps each coordinate p to its corresponding signal value and is optimized over the available observations. We define

$$Y = M \odot X^\star, \qquad \widehat{X}(p) = f_{\widehat{\theta}}(p), \qquad e(p) = X^\star(p) - \widehat{X}(p), \tag{1}$$

where $\widehat{\theta}$ denotes the parameters and $\odot$ is the Hadamard product. The observed and missing sets are $\Omega = \{p : M(p) = 1\}$ and $\Omega^c = \{p : M(p) = 0\}$ respectively. The INR parameters are optimized using only the observations in Ω, together with some regularization imposed by the underlying reconstruction method.

Although e can be evaluated for diagnostic analysis when the complete signal is available, it is unavailable on Ω^c during actual reconstruction. Thus, convergence does not reveal where the remaining errors occur or whether they can be further reduced. We therefore first examine their distribution of post-reconstruction errors, and then investigate whether dominant error regions can be localized and refined without ground truth.

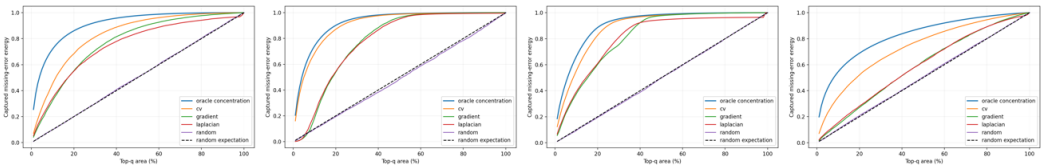


Figure 1: Error capture on four representative signals: Airplane (20%), News (20%), Toy (20%), and Washington DC (5%), where the percentages in parentheses denote the sampling ratios.

Table 1: Error capture ratios of the four cases in Fig. 1 under oracle ranking and different GT-free localization methods at varying top-q proportions.

Case	Oracle			CV			Gradient			Laplacian		
	Top 10%	Top 20%	Top 40%	Top 10%	Top 20%	Top 40%	Top 10%	Top 20%	Top 40%	Top 10%	Top 20%	Top 40%
Case 1	72.0	86.0	95.8	**45.4**	**68.4**	**88.8**	32.4	55.3	81.5	34.1	55.0	77.9
Case 2	74.0	90.1	98.2	**68.4**	**87.1**	**97.6**	18.0	53.7	88.4	21.1	54.3	86.6
Case 3	66.3	88.5	97.5	**55.3**	**82.6**	**96.8**	37.6	60.1	92.3	39.2	60.9	92.3
Case 4	54.5	68.8	83.7	**35.1**	**52.0**	**72.6**	14.0	27.1	51.5	15.4	28.1	51.5

2.2 Are Post-Reconstruction Errors Uniformly Distributed?

We first examine the distribution of post-reconstruction errors. Based on $e(p)$, we measure the fraction of total squared error captured by the top-$q\%$ entries ranked by error magnitude:

$$R_{\mathrm{err}}(q) = \frac{\sum_{p \in \mathcal{S}_{\mathrm{err}}(q)} |e(p)|^2}{\sum_p |e(p)|^2}, \tag{2}$$

where $\mathcal{S}_{\mathrm{err}}(q)$ contains the top-$q\%$ fraction of entries with the largest $|e(p)|$.

As shown by blue curve in Figure 1, the cumulative error generally increases much faster than the selected-entry ratio across different signals and reconstruction methods, indicating that post-reconstruction errors are often concentrated in a relatively small subset of entries. This motivates focusing refinement on error-dominant regions rather than uniformly treating the entire signal. The key question is then whether such regions can be localized without access to ground truth.

2.3 GT-Free Error Probing via Cross-Validation

To localize potential reconstruction errors without access to missing-entry ground truth, we construct out-of-fold (O-O-F) reconstructions from the observed set Ω. Specifically, Ω is partitioned into K disjoint folds, and each fold is held out in turn:

$$\Omega = \bigcup_{k=1}^{K} \Omega^{(k)}, \qquad \Omega^{(k)} \cap \Omega^{(l)} = \emptyset \ (k \neq l), \qquad \widehat{X}^{(-k)}(p) = f_{\widehat{\theta}}^{(-k)}(p), \tag{3}$$

We then define the cross-validation uncertainty as

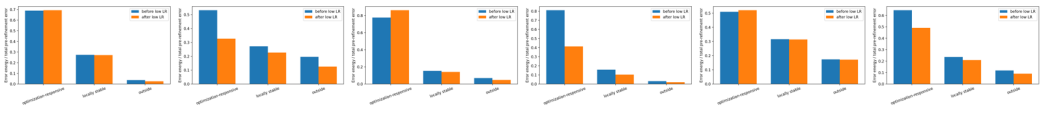
$$U_{\mathrm{CV}}(p) = \mathrm{Var}_{k=1,\ldots,K}\left[\widehat{X}^{(-k)}(p)\right]. \tag{4}$$

A large $U_{\mathrm{CV}}(p)$ indicates low prediction stability and hence potentially unreliable reconstruction. The prediction at the observation entries also provide O-O-F residuals, which will later be used for residual correction.

2.4 Can Reconstruction Errors Be Localized without GT?

To localize error-dominant regions without ground truth, we compare the CV uncertainty U_{CV} with two local variation indicators,the gradient magnitude and the Laplacian magnitude.

$$U_{\mathrm{Grad}}(p) = \|\nabla \widehat{X}(p)\|, \qquad U_{\mathrm{Lap}}(p) = |\Delta \widehat{X}(p)|. \tag{5}$$



(a) (b) (c) (d) (e) (f)

Figure 2: Proportions of different errors before and after optimization. image(a) and (b) are signals (News(5%),Flowers(20%)) under the method Rep-TRFD, image(c) and (d) are signals (Akiyo(5%),Toy(20%)) under the method DRO-TFF, image(e) and (f) are signals (House(20%),Toy(20%)) under the method HLRTF.

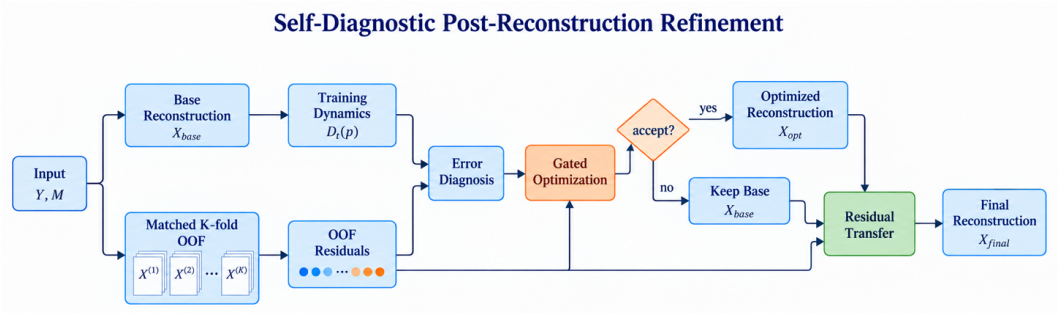


Figure 3: Overview of the proposed Self-Diagnostic Refinement (SDR) framework.

For any indicator U, let $\mathcal{S}_U(q)$ denote the top-$q\%$ fraction of entries ranked by $U(p)$. Its error capture ratio is

$$R_U(q) = \frac{\sum_{p\in\mathcal{S}_U(q)} |e(p)|^2}{\sum_p |e(p)|^2}. \tag{6}$$

As shown in Figure 1 and Table 1, We can find that, overall, CV uncertainty consistently captures more post-reconstruction error than the gradient-based and Laplacian-based indicators, and remains closer to the oracle ranking across various representative signals and methods. We therefore use U_{CV} for GT-free error localization in the following analysis.

2.5 What Remains after Further Optimization?

Having localized the error-dominant regions, we first evaluate whether they can be reduced by further optimization. Let $\widehat{X}^{(t)}$ denote the reconstruction at step t. We measure the local optimization variation by $D_t(p) = \|\widehat{X}^{(t+\Delta t)}(p) - \widehat{X}^{(t)}(p)\|_2$ and normalize it to $\widetilde{D}_t(p)$. Combined with the normalized CV uncertainty $\widetilde{U}_{\mathrm{CV}}(p)$, we define

$$w_{\mathrm{opt}}(p) = \widetilde{U}_{\mathrm{CV}}(p)\widetilde{D}_t(p), \quad w_{\mathrm{sta}}(p) = \widetilde{U}_{\mathrm{CV}}(p)(1-\widetilde{D}_t(p)), \quad w_{\mathrm{out}}(p) = 1-\widetilde{U}_{\mathrm{CV}}(p), \tag{7}$$

which are optimization-responsive, locally stable, and outside-CV components respectively.

For diagnostic analysis, let $E_t(p) = \|X^\star(p) - \widehat{X}^{(t)}(p)\|_2^2$, and $g \in \{\mathrm{opt}, \mathrm{sta}, \mathrm{out}\}$. We use the total pre-optimization error as a common denominator, and then report

$$R_g^{\mathrm{before}} = \frac{\sum_p w_g(p)E_t(p)}{\sum_h \mathcal{E}_h^{(t)}}, \qquad R_g^{\mathrm{after}} = \frac{\sum_p w_g(p)E_{t+\Delta t}(p)}{\sum_h \mathcal{E}_h^{(t)}}. \tag{8}$$

We plotted the relative proportions of the various errors before and after optimization, as shown in Figure 2, high-CV errors exhibit different optimization behaviors: the optimization-responsive component may decrease, remain nearly unchanged, or even increase after further training, while the locally stable component are relatively stable. Thus, additional optimization should be gated rather than applied directly, and the remaining stable residuals require a complementary correction mechanism.

3 SELF-DIAGNOSTIC ADAPTIVE REFINEMENT FRAMEWORK

Based on the above observations, we develop SDR, a ground-truth free self-diagnostic adaptive refinement framework, as illustrated in Figure 3. CV uncertainty guides gated global optimization, and the structured O-O-F residuals are transferred to correct the remaining stable errors. Together, the two stages provide complementary post-reconstruction refinement.

3.1 CV-GUIDED GATED OPTIMIZATION

The analysis above suggests that further optimization can be beneficial, but may also lead to overfitting. We therefore use O-O-F predictions to determine whether additional optimization should be applied and when it should stop.

Let $\widehat{X}_t^{(-k)}$ denote the k-th O-O-F reconstruction after t additional optimization steps, with $t = 0$ being the original converged state. For the held-out entries in $\Omega^{(k)}$, we define $E_k(t) = \sum_{s\in\Omega^{(k)}} \|\widehat{X}_t^{(-k)}(s) - Y(s)\|_2^2$, and aggregate all folds as

$$E_{\mathrm{OOF}}(t) = \frac{\sum_{k=1}^{K} E_k(t)}{\sum_{k=1}^{K} |\Omega^{(k)}|}. \tag{9}$$

The preferred continuation step and its relative O-O-F gain are

$$t^\star = \arg\min_{t\in\mathcal{T}} E_{\mathrm{OOF}}(t), \qquad G_{\mathrm{OOF}} = 1 - \frac{E_{\mathrm{OOF}}(t^\star)}{E_{\mathrm{OOF}}(0)}. \tag{10}$$

Additional optimization is accepted only if $G_{\mathrm{OOF}} \geq \tau$; the full reconstruction is then optimized for $t^\star$ steps, otherwise the original converged result is rejected. Thus, O-O-F evaluation acts as a gate that enables further optimization only when out-of-fold improvement is observed.

3.2 STRUCTURAL O-O-F RESIDUAL TRANSFER

Since some high-CV errors may persist after gated optimization, we therefore use the signed O-O-F residuals to correct these remaining errors. We assume that locations with similar local structures tend to exhibit similar systematic residuals. For $s \in \Omega^{(k)}$, we define

$$d^{\mathrm{OOF}}(s) = Y(s) - \widehat{X}^{(-k)}(s), \qquad z^{(-k)}(s) = \phi(\widehat{X}^{(-k)}, s), \tag{11}$$

where $\phi(X, p) = \mathrm{vec}(P_h([\bar{X}, \nabla_x\bar{X}, \nabla_y\bar{X}], p))$ is an $h \times h$ local structural descriptor and $\bar{X}$ denotes a compact channel-averaged representation. The pairs $(z^{(-k)}(s), \widetilde{d}^{\mathrm{OOF}}(s))$ from all folds form the residual memory, where $\widetilde{d}^{\mathrm{OOF}}$ denotes the clipped residual.

For a query location p, we compute $z^q(p) = \phi(\widehat{X}^{\mathrm{opt}}, p)$ and retrieve its K_{NN} most similar memory entries. Their residuals are aggregated as

$$C(p) = \sum_{s\in\mathcal{N}_{K_{\mathrm{NN}}}(p)} \alpha_{ps}\widetilde{d}^{\mathrm{OOF}}(s), \qquad \alpha_{ps} = \frac{\exp(\mathrm{sim}(p,s)/\tau_r)}{\sum_{r\in\mathcal{N}_{K_{\mathrm{NN}}}(p)} \exp(\mathrm{sim}(p,r)/\tau_r)}. \tag{12}$$

The final correction is applied only to the detected high-CV regions:

$$\widehat{X}^{\mathrm{corr}}(p) = \widehat{X}^{\mathrm{opt}}(p) + \beta M_{\mathrm{CV}}(p)C(p). \tag{13}$$

where M_{CV} denotes the mask of the detected high-CV regions. Thus, the module estimates the remaining structured residual rather than reconstructing the signal from scratch.

3.3 ANALYSIS OF STRUCTURAL O-O-F RESIDUAL TRANSFER

We next analyze when the transferred residual improves the reconstruction and how structurally matched O-O-F residuals can provide an effective correction.

We first characterize when the transferred residual is guaranteed to improve the reconstruction. Let $e_{\mathrm{opt}}(p) = X^\star(p) - \widehat{X}^{\mathrm{opt}}(p)$ denote the remaining reconstruction error after gated optimization, and

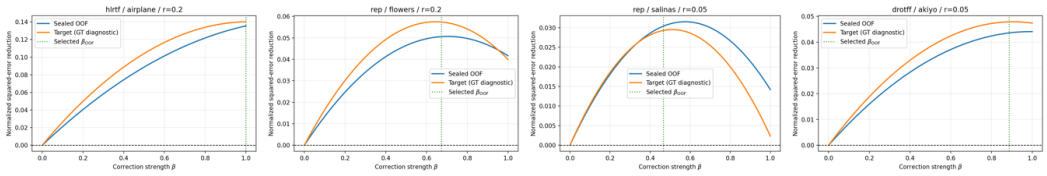


Figure 4: O-O-F and target-region error reduction curves under different correction strengths β on representative signals,Airplane(20%), Akiyo(5%), Flowers(20%), Salinas(5%).

Table 2: Validation of the residual-correction condition on representative cases.

Case	Rep-TRFD					DRO-TFF					HLRTF				
	A	B	$2A/B$	Selected β	Valid?	A	B	$2A/B$	Selected β	Valid?	A	B	$2A/B$	Selected β	Valid?
Airplane (20%)	1.14	1.22	1.86	0.87	✓	3.34	2.78	2.40	1.00	✓	1.36	1.03	2.64	1.00	✓
Akiyo (5%)	0.27	0.61	0.88	0.54	✓	1.64	1.65	1.98	0.88	✓	1.10	1.25	1.76	1.00	✓
Flowers (20%)	0.05	0.07	1.42	0.67	✓	0.07	0.10	1.40	0.67	✓	0.28	0.32	1.75	0.90	✓
Salinas (5%)	0.04	0.07	1.14	0.46	✓	5.52	5.64	1.95	0.96	✓	1.55	1.46	2.12	1.00	✓

let $\widetilde{C}(p) = M_{\mathrm{CV}}(p)C(p)$ denote the transferred correction restricted to the high-CV regions. The corrected reconstruction is then $\widehat{X}_\beta(p) = \widehat{X}^{\mathrm{opt}}(p) + \beta\widetilde{C}(p)$. Define $A = \sum_p \langle e_{\mathrm{opt}}(p), \widetilde{C}(p)\rangle$ and $B = \sum_p \|\widetilde{C}(p)\|_2^2$.

Theorem 3.1 (Effectiveness of residual correction)**.** *If $A > 0$ and $B > 0$, any $0 < \beta < 2A/B$ strictly decreases the squared reconstruction error:*

$$\sum_p \|X^\star(p) - \widehat{X}_\beta(p)\|_2^2 < \sum_p \|X^\star(p) - \widehat{X}^{\mathrm{opt}}(p)\|_2^2. \tag{14}$$

All proofs of the theoretical results presented in this section are provided in Appendix A. Theorem 3.1 shows that a transferred residual positively aligned with the remaining error is guaranteed to improve the reconstruction when its correction strength is properly bounded.

We next analyze how structurally matched O-O-F residuals can provide an effective estimate of the remaining error. We assume that the residual contains a transferable systematic component, $e_{\mathrm{opt}}(p) = b(z(p)) + \varepsilon(p)$, where $b(z)$ denotes the systematic residual and ε denotes an unpredictable component satisfying $\mathbb{E}\big[\|\varepsilon(p)\|_2^2 \mid z(p)\big] \le \sigma^2$. We further assume that b is L-Lipschitz in the descriptor space, i.e., $\|b(z_1) - b(z_2)\|_2 \le L\|z_1 - z_2\|_2$.

Then for an O-O-F memory entry $s \in \Omega^{(k)}$, we model its observable residual as $d^{\mathrm{OOF}}(s) = b(z(s)) + \Delta_k(z(s)) + \varepsilon(s)$, where Δ_k represents the discrepancy between the k-th fold reconstruction and the full reconstruction, assumed to satisfy $\|\Delta_k(z)\|_2 \le \delta$. The structural residual estimator at a query location p is

$$\widehat{b}(p) = \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} d^{\mathrm{OOF}}(s), \qquad \alpha_{ps} \ge 0, \qquad \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} = 1. \tag{15}$$

We further define the weighted neighborhood radius as

$$r_K(p) = \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \|z(s) - z(p)\|_2, \tag{16}$$

which measures the structural discrepancy between the query and its retrieved neighbors.

Theorem 3.2 (Error bound of structural O-O-F residual transfer)**.** *Under the assumptions above, the structural residual estimator satisfies*

$$\mathbb{E}\left[\left\|\widehat{b}(p) - b(z(p))\right\|_2^2\right] \le \left(L r_K(p) + \delta + \sigma\right)^2. \tag{17}$$

The bound consists of three factors: the structural retrieval error $L r_K(p)$, the fold-to-full discrepancy δ, and the unpredictable residual variance. Thus, better structural matches and effective neighbor aggregation improve the accuracy of residual transfer. When the transferred residual accurately estimates the systematic component, it is more likely to align positively with the remaining reconstruction error and hence satisfy the improvement condition in Theorem 3.1.

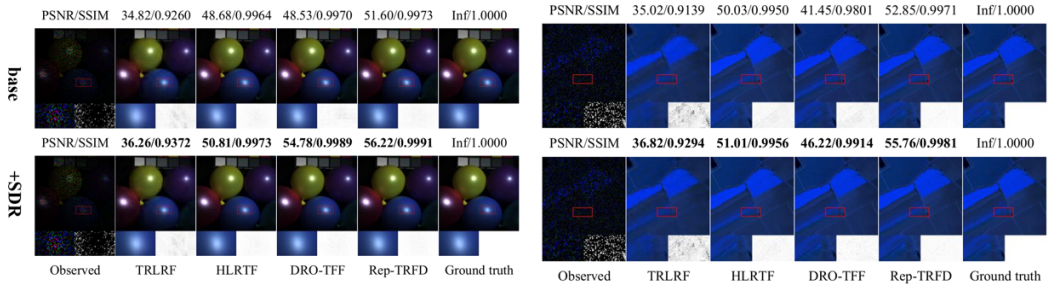


Figure 5: Representative visual comparison of different reconstruction methods before and after applying SDR on MSI Balloons(20%) and HSI Salinas(20%)

Table 3: Multi-seed reconstruction results on representative cases. All values are reported as mean $\pm$ standard deviation of the PSNR in 5 seeds.

Case	Rep-TRFD			DRO-TFF			HLRTF			TRLRF		
	Base	+SDR	Gain	Base	+SDR	Gain	Base	+SDR	Gain	Base	+SDR	Gain
Peppers (10%)	26.00 ± 0.11	26.92 ± 0.14	$+0.91 \pm 0.08$	23.53 ± 0.23	24.48 ± 0.12	$+0.95 \pm 0.15$	23.87 ± 0.07	24.61 ± 0.07	$+0.73 \pm 0.05$	15.62 ± 0.33	16.69 ± 0.32	$+1.07 \pm 0.02$
Akiyo (20%)	38.99 ± 0.06	40.10 ± 0.02	$+1.10 \pm 0.07$	35.28 ± 0.15	36.13 ± 0.08	$+0.85 \pm 0.08$	36.67 ± 0.25	37.62 ± 0.16	$+0.95 \pm 0.10$	29.23 ± 0.02	30.53 ± 0.04	$+1.29 \pm 0.04$
Balloons (20%)	51.81 ± 0.44	56.18 ± 0.17	$+4.37 \pm 0.42$	48.32 ± 0.27	54.55 ± 0.17	$+6.23 \pm 0.21$	48.48 ± 0.21	50.82 ± 0.17	$+2.34 \pm 0.23$	35.50 ± 0.56	37.00 ± 0.60	$+1.49 \pm 0.05$
Salinas (20%)	53.05 ± 0.12	55.86 ± 0.09	$+2.80 \pm 0.13$	40.87 ± 0.55	45.76 ± 0.47	$+4.88 \pm 0.15$	49.77 ± 0.23	50.87 ± 0.22	$+1.09 \pm 0.11$	35.93 ± 0.51	37.69 ± 0.49	$+1.76 \pm 0.03$

3.4 Validation of Residual Transfer

We empirically validate the correction condition in Theorem 3.1 and the transferability of O-O-F behavior to the missing region. Table 2 reports A, B, and $\beta_c = 2A/B$ on the missing entries, where positive A confirms that the transferred residual is aligned with the remaining reconstruction error. We further change β and compare the relative error reduction on O-O-F and missing entries. As shown in Figure 4, the two curves exhibit similar trends, indicating that O-O-F correction behavior provides a useful proxy for selecting a suitable correction strength on the unknown missing region.

4 Experiments

In this section, we evaluate the proposed self-diagnosis adaptive refinement framework on missing-data reconstruction of RGB images, videos, multispectral images and hyperspectral images under different sampling ratios and seeds. The experiments include overall performance comparison, mechanism validation, ablation studies, and parameter sensitivity analysis.

4.1 Experimental Settings

The datasets are taken from USC-SIPI, YUV sequences, CAVE, and Remote Sensing Scenes for their extensive usage. Unless otherwise specified, the O-O-F gain threshold is set to 1%, the top 40% entries are selected as high-CV regions, and an 3×3 patch with 16 neighbors are used for residual transfer. All compared methods follow their original optimization settings. Experiments are conducted on an Intel(R) Core(TM) i9-14900KF CPU, 64 GB RAM, and NVIDIA GeForce RTX 4090 D.

4.2 Multi-Seed Main Results

We evaluate the proposed self-diagnosis adaptive refinement framework on several representative INR-based reconstruction methods, Rep-TRFDXu et al. (2026), DRO-TFFLi et al. (2025) and HLRTFLuo et al. (2022), we also apply it to a traditional method TRLRFYuan et al. (2019) to examine its applicability beyond INRs. All images are resized or cropped to consistent input sizes, while 100 consecutive frames are used for video reconstruction. Reconstruction quality is evaluated using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM), which are widely used for signal reconstruction assessment. Existing iterative refinement methods are not directly comparable to our post-reconstruction refinement and ground-truth-free setting due to differences in problem formulation and training requirements; we discuss this issue in detail in Appendix B.

Table 4: Validation of the O-O-F-guided optimization gate on representative cases. ΔPSNR denotes the change after additional optimization.

Case	Rep-TRFD			DRO-TFF			HLRTF		
	Accept?	ΔPSNR	Gate Correct?	Accept?	ΔPSNR	Gate Correct?	Accept?	ΔPSNR	Gate Correct?
Peppers (10%)	Yes	+0.44	✓	Yes	+0.26	✓	Yes	+0.25	✓
Akiyo (3%)	Yes	+0.08	✓	No	-0.18	✓	No	-0.06	✓
Toy (20%)	Yes	+3.30	✓	Yes	+2.71	✓	Yes	+1.02	✓
Botswana (3%)	No	-0.01	✓	Yes	+0.49	✓	No	-0.01	✓

Table 5: Comparison of different residual-transfer strategies on representative cases. Structural denotes the proposed structural residual transfer. All values are reported in ΔPSNR (dB).

Case	Rep-TRFD				DRO-TFF				HLRTF				TRLRF			
	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse
Airplane (20%)	**+0.53**	-0.11	-0.13	-1.44	**+0.96**	+0.25	-0.04	-2.19	**+0.58**	-0.05	-0.12	-1.30	**+0.69**	+0.31	-0.11	-1.29
Akiyo (20%)	**+0.21**	-0.06	-0.02	-0.38	**+0.42**	-0.13	-0.07	-0.87	**+0.80**	-0.12	-0.14	-0.64	**+1.30**	+0.13	-0.10	-1.98
Toy (5%)	**+0.73**	-0.13	-0.14	-1.82	**+0.77**	+0.19	-0.07	-1.94	**+0.65**	+0.43	-0.08	-1.43	**+0.15**	+0.07	-0.13	-0.72
Salinas (20%)	**+0.16**	-0.01	-0.02	-0.36	**+3.90**	+3.76	+2.64	-4.56	**+1.00**	+0.56	-0.13	-2.32	**+1.80**	+1.06	-0.09	-2.37

We thus evaluate the results on each method with four representative data, which are RGB image, video, MSI, and HSI. And the results are averaged over 5 seeds with standard deviations reported. As shown in Table 3, SDR consistently improves all evaluated reconstruction methods across the four signal types, demonstrating its broad effectiveness and stability, and the largest improvement is more than 6 dB in PSNR. Representative visual results on MSI and HSI are shown in Figure 5, where the improvements brought by SDR are consistently observed across different base methods.

4.3 Further Analysis

In this section, we further evaluate the proposed framework through gated optimization, residual transfer, ablation, and parameter sensitivity analyses. Representative results are reported here, with complete results provided in Appendix C.

Validation of Gated Optimization We first validate whether the CV-guided gate can reliably determine when additional optimization is beneficial. For each reconstruction method, we evaluate gated optimization on four representative datasets by measuring the PSNR change after continued optimization and checking whether the gating decision is consistent with the actual performance change. As shown in Table 4, the gate correctly rejects all representative cases in which further optimization leads to performance degradation. It is worth noting that this includes not only cases with relatively noticeable PSNR drops, but also cases where the degradation is as small as approximately 0.01 dB, indicating that the gate is sensitive to even marginally harmful updates. At the same time, beneficial continued optimization is retained when further training improves reconstruction quality. These results demonstrate that the proposed gate can effectively distinguish useful optimization from unnecessary or detrimental updates, thereby providing a reliable basis for adaptive post-reconstruction refinement.

Residual-Memory Controls We further validate whether the improvement of residual transfer indeed comes from structurally matched O-O-F residuals. We compare the proposed structural residual transfer with three alternatives: spatial nearest-neighbor transfer, random residual matching, and reversed residual transfer. For each reconstruction method, we evaluate four representative cases and report the resulting change of PSNR. As shown in Table 5, we can find that the proposed structural residual transfer consistently achieves the largest PSNR improvement among representative cases. Spatial nearest-neighbor transfer can sometimes yield positive gains, but these are generally smaller than those obtained by structural matching, and performance degradation is still observed in some cases. Random residual matching leads to slight PSNR decreases in most cases, and reversed residual transfer substantially decreases reconstruction quality. These results indicate that effective residual correction depends on both the correct residual direction and structure-aware correspondence, rather than arbitrary local perturbations.

Ablation Experiment We conduct an ablation experiment to examine the contribution of each refinement component, including optimization-only refinement, gated optimization, residual transfer,

Table 6: Ablation study and computational cost on representative cases. Runtime is reported in seconds. O-O-F denotes the summed runtime over all fold reconstructions, and Total denotes the aggregate computational cost.

Case	Method	PSNR						Runtime (s)				
		Base	Opt.	Gated Opt.	Resid.	w/o Gate	Full	Base	Base+Opt.	O-O-F	Transfer	Total
Balloons (20%)	Rep-TRFD	51.60	55.85	55.85	52.17	56.22	56.22	25.17	53.83	526.62	3.29	586.51
	DRO-TFF	48.53	52.36	52.36	49.14	54.78	54.78	18.34	46.26	486.77	3.39	538.99
	HLRTF	48.68	48.96	48.96	50.49	50.81	50.81	7.07	18.45	155.16	3.28	179.38
	TRLRF	34.82	NA	NA	36.26	NA	36.26	76.02	NA	870.74	2.18	951.41
Salinas (20%)	Rep-TRFD	52.85	55.60	55.60	52.88	55.76	55.76	48.69	95.83	957.34	18.30	1077.71
	DRO-TFF	41.45	42.31	42.31	44.36	46.22	46.22	30.98	77.10	760.33	16.03	859.57
	HLRTF	50.03	50.01	50.01	50.97	51.01	51.01	8.46	25.02	250.77	15.68	298.23
	TRLRF	35.02	NA	NA	36.82	NA	36.82	365.34	NA	3944.01	18.69	4328.05

Table 7: Parameter sensitivity analysis on representative cases. All values are reported in ΔPSNR (dB).

Case	Method	Top-$q\%$ Region			Patch Size			K_{NN}		
		15%	20%	40%	3×3	5×5	7×7	8	16	32
Peppers (20%)	Rep-TRFD	+0.53	+0.59	**+0.68**	+0.68	**+0.70**	+0.60	+0.65	**+0.68**	+0.66
	DRO-TFF	+0.56	+0.65	**+0.85**	+0.85	**+0.95**	+0.92	+0.82	**+0.85**	+0.81
	HLRTF	+0.46	+0.54	**+0.75**	+0.75	**+0.80**	+0.75	+0.70	**+0.75**	+0.72
	TRLRF	+0.36	+0.46	**+0.77**	**+0.77**	+0.71	+0.68	+0.73	**+0.77**	+0.75
Carphone (20%)	Rep-TRFD	+0.43	+0.51	**+0.87**	**+0.87**	+0.83	+0.75	+0.86	**+0.87**	+0.85
	DRO-TFF	+0.22	+0.27	**+0.43**	**+0.43**	+0.41	+0.37	+0.43	**+0.44**	+0.43
	HLRTF	+0.28	+0.33	**+0.47**	**+0.47**	+0.45	+0.41	**+0.47**	**+0.47**	+0.46
	TRLRF	+0.86	+1.03	**+1.45**	+1.45	+1.70	**+1.83**	**+1.69**	+1.45	+1.21

refinement without the gate, and the full framework. We compare four methods on two representative data. For TRLRF (Yuan et al., 2019), which employs an ADMM-based solver iterated until convergence, further optimization is not applicable in the same sense as for INR-based methods; we therefore exclude it from the gated optimization analysis. As shown in Table 6, gated optimization generally preserves or improves the base reconstruction, whereas structural residual transfer provides further improvement; their combination achieves the best overall performance. In some cases, the gain of the full framework even exceeds the sum of the gains from the two individual stages, further demonstrating their complementary effects. Table 6 also reports the runtime of each component. The O-O-F runtime denotes the aggregate cost over all folds. Since the fold reconstructions are independent, they can be executed in parallel, so the practical wall-clock time can be substantially reduced, whereas the reported total reflects the aggregate computational cost.

Parameter Sensitivity We further evaluate the sensitivity of three key parameters: the proportion q of high-CV regions, the patch size h, and the number of retrieved residual entries K. We compare the Top-10%, Top-20%, and Top-40% region, the 3×3, 5×5, and 7×7 patch size, with K_{NN} = 8, 16, 32. For each reconstruction method, we evaluate two representative cases and report the resulting change of PSNR. As shown in Table 7, the performance remains relatively stable over a reasonable range of parameter settings. For the consideration of both reconstruction performance and computational cost, we set $q = [40]$, $h = [3]$, and $K = [16]$ in all experiments.

5 Conclusion

In this work, we presented SDR, a ground-truth-free self-diagnostic refinement framework for post-reconstruction error reduction. We showed that post-reconstruction errors are often concentrated, can be localized through cross-validation uncertainty, and exhibit different optimization behaviors. We used SDR which is composed of gated global optimization and structural O-O-F residual transfer to further reduce these errors without access to missing-entry ground truth. Extensive experiments across different signal types, reconstruction methods, and missing ratios demonstrated consistent improvements and validate the effectiveness and robustness of the proposed framework. Future work includes extending the framework for denoising task.

AI USE STATEMENT

In this work, we used generative AI tools to assist with manuscript drafting and language polishing, literature search and related-work discovery, and feedback on the presentation of the methodology, experiments, and theoretical analysis. AI tools were also used to assist in checking and refining the presentation of mathematical claims and their derivations. All research ideas, methodological decisions, experimental results, and theoretical claims were proposed and verified by the authors. All cited literature was manually checked against the original sources, and all AI-assisted text was reviewed and revised by the authors. We take full responsibility for the final content of this work, including all text, claims, results, and artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide the detailed formulation and implementation procedure of SDR in Section 2 and Section 3, including O-O-F construction, gated optimization, and structural residual transfer. Experimental settings, datasets, preprocessing procedures, evaluation metrics, and hyperparameter choices are described in Section 4.1 and Appendix C. Complete proofs of the theoretical results, together with their assumptions, are provided in Appendix A. Additional experimental results, including multi-seed evaluations and extended analyses, are also included in the Appendix C. These details are provided to facilitate reproduction of both the proposed method and the reported results.

REFERENCES

Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International conference on machine learning*, pp. 524–533. PMLR, 2019.

J Douglas Carroll and Jih-Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of "eckart-young" decomposition. *Psychometrika*, 35(3):283–319, 1970.

I-Hsiang Chen, Isma Hadji, Enrique Sanchez, Adrian Bulat, Sy-Yen Kuo, Radu Timofte, Georgios Tzimiropoulos, and Brais Martinez. Restore, assess, repeat: A unified framework for iterative image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15506–15515, June 2026.

Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. An introduction to independent component analysis. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 14(3):123–149, 2000.

Ali Haider, Muhammad Salman Ali, Maryam Qamar, Tahir Khalil, Soo Ye Kim, Jihyong Oh, Enzo Tartaglione, and Sung-Ho Bae. I-inr: iterative implicit neural representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 4520–4528, 2026.

Jingyao Hou, Feng Zhang, Yao Wang, and Jianjun Wang. Low-tubal-rank tensor recovery from one-bit measurements. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3302–3306. IEEE, 2020.

Jingyao Hou, Feng Zhang, Haiquan Qiu, Jianjun Wang, Yao Wang, and Deyu Meng. Robust low-tubal-rank tensor recovery from binary measurements. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4355–4373, 2021.

Jianbo Jiao, Wei-Chih Tu, Shengfeng He, and Rynson W. H. Lau. Formresnet: Formatted residual learning for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

Weichao Kong, Qingrong Feng, Qianyu Shu, Jianjun Wang, Tingwen Huang, and Bin Zhang. Generalized subspace coupling approach for robust low-tubal-rank tensor completion. *IEEE Transactions on Neural Networks and Learning Systems*, 2025.

Yanyi Li, Xi Zhang, Yisi Luo, and Deyu Meng. Deep rank-one tensor functional factorization for multi-dimensional data recovery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 18539–18547, 2025.

Chunyan Liu, Dianlin Hu, Jiangjun Peng, Hong Wang, Qianyu Shu, and Jianjun Wang. Convergence-guaranteed spectral ct reconstruction via internal and external prior mining. *IEEE Transactions on Computational Imaging*, 12:46–59, 2025.

Chunyan Liu, Tongle Wu, Hong Wang, Dianlin Hu, Bin Zhang, and Jianjun Wang. Low-complexity reconstruction of low-dose spectral ct via double low-rank tensor factorization with adaptive transforms. *Medical Image Analysis*, pp. 104119, 2026.

Xinling Liu, Jingyao Hou, Jiangjun Peng, Hailin Wang, Deyu Meng, and Jianjun Wang. Tensor compressive sensing fused low-rankness and local-smoothness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 8879–8887, 2023.

Yisi Luo, Xi-Le Zhao, Deyu Meng, and Tai-Xiang Jiang. Hlrtf: Hierarchical low-rank tensor factorization for inverse problems in multi-dimensional imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19303–19312, 2022.

Yisi Luo, Xile Zhao, Zhemin Li, Michael K Ng, and Deyu Meng. Low-rank tensor function representation for multi-dimensional data recovery. *IEEE transactions on pattern analysis and machine intelligence*, 46(5):3351–3369, 2023.

Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

Wenjin Qin, Hailin Wang, Weijun Ma, and Jianjun Wang. Robust high-order tensor recovery via nonconvex low-rank approximation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3633–3637. IEEE, 2022a.

Wenjin Qin, Hailin Wang, Feng Zhang, Jianjun Wang, Xin Luo, and Tingwen Huang. Low-rank high-order tensor completion with applications in visual data. *IEEE Transactions on Image Processing*, 31:2433–2448, 2022b.

Wenjin Qin, Hailin Wang, Feng Zhang, Weijun Ma, Jianjun Wang, and Tingwen Huang. Nonconvex robust high-order tensor completion using randomized low-rank approximation. *IEEE Transactions on Image Processing*, 33:2835–2850, 2024.

Wenjin Qin, Hailin Wang, Hao Shu, Feng Zhang, Jianjun Wang, Xiangyong Cao, Xi-Le Zhao, and Gemine Vivone. Hyperspectral anomaly detection fused unified nonconvex tensor ring factors regularization. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–21, 2025.

Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18507–18516, 2023.

Zenglin Shi, Pascal Mettes, Subhransu Maji, and Cees GM Snoek. On measuring and controlling the spectral bias of the deep image prior. *International Journal of Computer Vision*, 130(4):885–908, 2022.

Hailin Wang, Feng Zhang, Jianjun Wang, and Yao Wang. Estimating structural missing values via low-tubal-rank tensor completion. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3297–3301. IEEE, 2020.

Hailin Wang, Feng Zhang, Jianjun Wang, Tingwen Huang, Jianwen Huang, and Xinling Liu. Generalized nonconvex approach for low-tubal-rank tensor recovery. *IEEE Transactions on Neural Networks and Learning Systems*, 33(8):3305–3319, 2021a.

Hailin Wang, Jiangjun Peng, Xiangyong Cao, Jianjun Wang, Qibin Zhao, and Deyu Meng. Hyperspectral image denoising via nonlocal spectral sparse subspace representation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:5189–5203, 2023a.

Hailin Wang, Jiangjun Peng, Wenjin Qin, Jianjun Wang, and Deyu Meng. Guaranteed tensor recovery fused low-rankness and smoothness. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10990–11007, 2023b.

Hengkang Wang, Taihui Li, Zhong Zhuang, Tiancong Chen, Hengyue Liang, and Ju Sun. Early stopping for deep image prior. *arXiv preprint arXiv:2112.06074*, 2021b.

Jianjun Wang, Jingyao Hou, and Yonina C Eldar. Tensor robust principal component analysis from multilevel quantized observations. *IEEE transactions on information theory*, 69(1):383–406, 2022.

Liang Wu, Jianjun Wang, Wei-Shi Zheng, and Guangming Shi. Double nonconvex tensor robust kernel principal component analysis and its visual applications. *IEEE Transactions on Image Processing*, 2026.

Yangyang Xu, Junbo Ke, You-Wei Wen, and Chao Wang. Reparameterized tensor ring functional decomposition for multi-dimensional data recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26389–26398, 2026.

Longhao Yuan, Chao Li, Danilo Mandic, Jianting Cao, and Qibin Zhao. Tensor ring decomposition with rank minimization on latent space: An efficient approach for tensor completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 9151–9158, 2019.

Feng Zhang, Jingyao Hou, Jianjun Wang, and Wendong Wang. Uniqueness guarantee of solutions of tensor tubal-rank minimization problem. *IEEE Signal Processing Letters*, 27:540–544, 2020a.

Feng Zhang, Jianjun Wang, Wendong Wang, and Chen Xu. Low-tubal-rank plus sparse tensor recovery with prior subspace information. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3492–3507, 2020b.

Qibin Zhao, Guoxu Zhou, Shengli Xie, Liqing Zhang, and Andrzej Cichocki. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*, 2016.

The Appendix provides additional technical detail and experimental results.In Appendix A, we present all the proofs of the theorems in the main paper. In Appendix B, we analyze the differences from related work and explain why we don't do comparison experiments with other methods. In Appendix C, we present the comprehensive experiments on robust results, error diagnostic, correction analysis, mechanism validation, and ablation experiments.

A Proofs of Theoretical Results

A.1 Proof of Theorem 3.1

Proof. Let $\mathcal{D}$ denote the region where the residual correction is applied, i.e., the support of M_{CV}. Define the squared reconstruction loss over $\mathcal{D}$ as

$$\mathcal{L}_{\mathcal{D}}(\beta) = \sum_{p\in\mathcal{D}} \left\| X^{\star}(p) - \widehat{X}_{\beta}(p) \right\|_2^2, \tag{18}$$

where $\widehat{X}_{\beta}(p) = \widehat{X}^{\mathrm{opt}}(p) + \beta\widetilde{C}(p)$. For $\beta = 0$, we have

$$\mathcal{L}_{\mathcal{D}}(0) = \sum_{p\in\mathcal{D}} \|e_{\mathrm{opt}}(p)\|_2^2. \tag{19}$$

Since

$$X^{\star}(p) - \widehat{X}_{\beta}(p) = e_{\mathrm{opt}}(p) - \beta\widetilde{C}(p),$$

the error reduction introduced by the correction is

$$\begin{aligned}
\mathcal{L}_{\mathcal{D}}(0) - \mathcal{L}_{\mathcal{D}}(\beta) &= \sum_{p\in\mathcal{D}} \left(\|e_{\mathrm{opt}}(p)\|_2^2 - \|e_{\mathrm{opt}}(p) - \beta\widetilde{C}(p)\|_2^2 \right) && (20)\\
&= 2\beta \sum_{p\in\mathcal{D}} \left\langle e_{\mathrm{opt}}(p), \widetilde{C}(p) \right\rangle - \beta^2 \sum_{p\in\mathcal{D}} \|\widetilde{C}(p)\|_2^2 && (21)\\
&= 2\beta A - \beta^2 B && (22)\\
&= \beta(2A - \beta B). && (23)
\end{aligned}$$

Since $A > 0$, $B > 0$, and $0 < \beta < 2A/B$, we have $\beta(2A - \beta B) > 0$. Therefore,

$$\mathcal{L}_{\mathcal{D}}(\beta) < \mathcal{L}_{\mathcal{D}}(0).$$

Since $\widetilde{C}(p) = 0$ outside $\mathcal{D}$, the reconstruction remains unchanged there. Hence, the total squared reconstruction error is also strictly decreased. □

Remark 1. *For a signal with N entries and fixed dynamic range $I_{\max}$, PSNR is defined as*

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{I_{\max}^2}{\mathrm{MSE}}\right), \qquad \mathrm{MSE} = \frac{1}{N}\sum_{p} \left\| X^{\star}(p) - \widehat{X}(p) \right\|_2^2. \tag{24}$$

Therefore, for a fixed signal size and dynamic range, PSNR is monotonically decreasing when the squared reconstruction error increasing. Thus, any β satisfying the condition in Theorem 3.1 guarantees a strictly higher PSNR than the uncorrected reconstruction.

A.2 Proof of Theorem 3.2

Proof. For an O-O-F memory entry $s \in \Omega^{(k)}$, we have

$$d^{\mathrm{OOF}}(s) = b(z(s)) + \Delta_k(z(s)) + \varepsilon(s).$$

Using

$$\widehat{b}(p) = \sum_{s\in\mathcal{N}_K(p)} \alpha_{ps} d^{\mathrm{OOF}}(s), \qquad \sum_{s\in\mathcal{N}_K(p)} \alpha_{ps} = 1,$$

we obtain

$$\widehat{b}(p) - b(z(p)) = \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \big(b(z(s)) - b(z(p)) \big) + \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \Delta_k(z(s)) + \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \varepsilon(s). \tag{25}$$

Conditioning on the descriptors and the retrieved neighborhood, Minkowski's inequality gives,

$$\left(\mathbb{E} \left[\left\| \widehat{b}(p) - b(z(p)) \right\|_2^2 \right] \right)^{1/2} \leq \left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \big(b(z(s)) - b(z(p)) \big) \right\|_2 + \left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \Delta_k(z(s)) \right\|_2 + \left(\mathbb{E} \left[\left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \varepsilon(s) \right\|_2^2 \right] \right)^{1/2}. \tag{26}$$

For the first term, the L-Lipschitz continuity of b shows that

$$\left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \big(b(z(s)) - b(z(p)) \big) \right\|_2 \leq \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \| b(z(s)) - b(z(p)) \|_2 \tag{27}$$

$$\leq L \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \| z(s) - z(p) \|_2 \tag{28}$$

$$= L r_K(p). \tag{29}$$

Similarly, since $\|\Delta_k(z)\|_2 \leq \delta$, we have

$$\left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \Delta_k(z(s)) \right\|_2 \leq \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \| \Delta_k(z(s)) \|_2 \tag{30}$$

$$\leq \delta. \tag{31}$$

For the unpredictable residual term, Minkowski's inequality gives

$$\left(\mathbb{E} \left[\left\| \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \varepsilon(s) \right\|_2^2 \right] \right)^{1/2} \leq \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \left(\mathbb{E}[\|\varepsilon(s)\|_2^2] \right)^{1/2} \tag{32}$$

$$\leq \sigma \sum_{s \in \mathcal{N}_K(p)} \alpha_{ps} \tag{33}$$

$$= \sigma. \tag{34}$$

Combining the above bounds, we have

$$\left(\mathbb{E} \left[\left\| \widehat{b}(p) - b(z(p)) \right\|_2^2 \right] \right)^{1/2} \leq L r_K(p) + \delta + \sigma.$$

Squaring both sides, we have

$$\mathbb{E} \left[\left\| \widehat{b}(p) - b(z(p)) \right\|_2^2 \right] \leq \left(L r_K(p) + \delta + \sigma \right)^2,$$

which completes the proof. □

Remark 2. *The bound in Theorem 3.2 reveals three sources of estimation error. First, $Lr_K(p)$ characterizes the structural mismatch between the query and its retrieved O-O-F neighbors. Here, L measures the sensitivity of the systematic residual to changes in the descriptor, while $r_K(p)$ is directly determined by the retrieved neighborhood. Thus, structure-aware retrieval that selects closer neighbors reduces this term, and $r_K(p)$ can be directly computed from the descriptor distances.*

Second, δ measures the discrepancy between the O-O-F reconstruction and the full reconstruction. Although the underlying systematic residuals are not directly observable, this discrepancy can be reduced by using balanced folds and consistent reconstruction settings, and can be empirically approximated by the prediction discrepancy between O-O-F and full reconstructions on observed entries.

Finally, σ represents the unpredictable component of the remaining residual and therefore characterizes the intrinsic uncertainty that cannot be explained by the structural descriptor. Consequently, the bound becomes tighter when structurally similar neighbors are retrieved, the O-O-F and full reconstructions are consistent, and the unpredictable residual component is small.

B Comparability with Related Methods

B.1 Closely Related Refinement Methods

Several recent works share a similar goal of improving an initial reconstruction through iterative refinement. However, their problem settings and training requirements differ substantially from ours. For example, I-INR (Haider et al., 2026) introduces dedicated iterative modules and trains the resulting model toward clean target signals, rather than refining an already converged reconstruction without access to missing-entry ground truth. Similarly, RAR (Chen et al., 2026) jointly trains restoration and quality assessment within an end-to-end image restoration framework, where refinement is performed through a predefined restoration process for image degradations.

In contrast, SDR does not require modifying or retraining the underlying reconstruction model. It operates after the base reconstruction has converged and determines, using only the available observations, whether further optimization is beneficial and how the remaining errors should be corrected. Therefore, directly applying these methods as experimental baselines would change the underlying reconstruction pipeline, supervision requirement, or task formulation, making a controlled comparison with SDR non-trivial.

B.2 Methods with Related Principles

Other methods are conceptually related to individual components of SDR, such as held-out self-supervision (Batson & Royer, 2019), ground-truth-free model assessment (Shi et al., 2022), early stopping (Wang et al., 2021b), or residual-based refinement (Jiao et al., 2017). These ideas provide important inspiration for evaluating or improving reconstructions without relying entirely on clean targets. However, they are generally developed for specific tasks or objectives, such as denoising, restoration quality assessment, or optimization stopping, rather than for diagnosing and refining post-reconstruction errors across different reconstruction models and signal modalities.

SDR instead treats ground-truth-free diagnosis and refinement as a unified post-reconstruction problem: O-O-F uncertainty is used to localize potentially erroneous regions and gate continued optimization, while structurally matched O-O-F residuals provide correction cues for the remaining errors. Consequently, methods sharing only individual principles do not constitute direct experimental baselines under the same problem setting.

C Comprehensive Experimental Results

As discussed in the preceding section, existing iterative refinement methods are different from our setting in supervision, training requirements, or task formulation, and therefore do not provide directly controlled baselines for post-reconstruction refinement. We instead evaluate SDR as a unified refinement framework applied to different reconstruction methods, signal modalities, and sampling ratios.

For each case, the base reconstruction is first optimized using its original setting until pre-set iterations. SDR is then applied without access to the ground truth of missing entries: O-O-F reconstructions are constructed from the available observations, CV uncertainty is used for gated global optimization, and structurally matched O-O-F residuals are subsequently transferred to correct the remaining errors. Unless otherwise specified, we use $q = 40\%$, a 3×3 structural patch, $K_{\mathrm{NN}} = 16$, and a gating threshold of $\tau = 1\%$ throughout all experiments. All remaining settings of the underlying reconstruction methods follow their original implementations and are kept unchanged before and after applying SDR.

C.1 Robustness and Visual Results

In this subsection, we provide a comprehensive evaluation of SDR across all considered signals and sampling ratios. The PSNR of the base reconstruction, the result after applying SDR, and the corresponding improvement (ΔPSNR), where all results are reported as mean$\pm$standard deviation over 5 seeds are shown in Table C.1. From the table we can find that SDR consistently improves the base reconstructions across the evaluated signal modalities, methods, and sampling ratios. Although the magnitude of improvement varies with the reconstruction setting, positive gains are observed throughout the evaluation, with the relatively small variations across repeated runs demonstrate the robustness of the refinement.

We further present visual comparisons between the base and SDR-refined reconstructions on different signal modalities, together with their PSNR and SSIM values in Figures C.1–C.3. From the figures, we can find that the refined results consistently achieve improved quantitative quality and preserve the overall structures of the original reconstructions, with improvements in local reconstruction quality observed across different cases. These results are consistent with the quantitative evaluation and further demonstrate the effectiveness and robustness of SDR under different reconstruction conditions.

C.2 Diagnostic and Correction Analysis

In this subsection, we first extend the analysis of post-reconstruction error concentration and GT-free localization to a broader range of signals and reconstruction settings. As shown in Figure C.4, post-reconstruction errors are often concentrated in a relatively small subset of entries, although the degree of concentration varies across different signals and sampling ratios. Overall, CV uncertainty provides the strongest error-localization performance among the evaluated GT-free indicators. In a few cases, the capture curves of different indicators largely overlap, indicating comparable localization performance among these indicators. These results support the use of CV uncertainty as a generally effective indicator for locating error-dominant regions without access to the ground truth of missing-entries.

We further examine whether the correction behavior observed on O-O-F entries is consistent with that on the missing region. Figure C.5 compares the relative error reduction on the O-O-F and target regions under different correction strengths β, together with the selected β. Across most of the evaluated cases, the two curves exhibit similar variation trends, indicating that O-O-F behavior provides a useful proxy for selecting the residual correction strength. Moreover, the selected β consistently leads to effective correction on the target region, further supporting the transferability of O-O-F residual information for post-reconstruction refinement.

C.3 Mechanism Validation and Ablation Study

We first evaluate the effectiveness of the O-O-F-based optimization gate. As shown in Table C.2, the gating decision is correct in most evaluated cases, successfully accepting beneficial continued optimization and rejecting cases that would lead to performance degradation. The few mismatched cases correspond only to marginal PSNR changes on the order of several hundredths of a decibel, indicating that the gate remains reliable for practically meaningful optimization decisions.

We next examine whether structural matching is necessary for residual transfer. We compares the proposed structural KNN retrieval with spatial nearest-neighbor, random, and reversed residual transfer in Table C.3. Structural matching achieves the best performance in most cases and yields positive improvements throughout the evaluation. Spatial nearest-neighbor transfer sometimes pro-

vides a slightly larger gain, but its behavior is inconsistent, with both positive and negative effects across different cases. Random and reversed transfer are generally less effective or even detrimental. These results indicate that spatial proximity alone is not a reliable criterion for residual selection, whereas structural similarity provides a more stable basis for residual transfer.

Finally, We presents the ablation results of the two refinement components in Table C.4. Both gated optimization and structural residual transfer independently improve the base reconstruction, while their combination consistently provides further improvement. The gain of the complete framework can be larger than, smaller than, or close to the sum of the two individual gains depending on the reconstruction setting, reflecting interactions between the two stages rather than a simple additive effect. Overall, the results confirm that the two components are complementary and jointly contribute to the effectiveness of SDR.

Table C.1: Complete robustness results across different signals and sampling ratios. Base, +SDR, and Gain denote the original reconstruction, the SDR-refined reconstruction, and the corresponding PSNR improvement, respectively. All results are reported as mean±standard deviation of PSNR over 5 seeds.

	Rep-TRFD			DRO-TFF			HLRTF			TRLRF		
Setting	Base	+SDR	Gain	Base	+SDR	Gain	Base	+SDR	Gain	Base	+SDR	Gain
Airplane (0.1)	26.35 ± 0.183	27.09 ± 0.090	$+0.74 \pm 0.098$	22.98 ± 0.287	24.97 ± 0.052	$+1.98 \pm 0.238$	25.31 ± 0.223	25.78 ± 0.186	$+0.46 \pm 0.070$	17.59 ± 0.153	18.72 ± 0.163	$+1.13 \pm 0.012$
Airplane (0.2)	30.46 ± 0.124	31.32 ± 0.116	$+0.86 \pm 0.073$	27.81 ± 0.275	29.19 ± 0.057	$+1.38 \pm 0.290$	28.62 ± 0.207	29.17 ± 0.198	$+0.54 \pm 0.031$	23.48 ± 0.098	24.17 ± 0.113	$+0.68 \pm 0.042$
Peppers (0.1)	26.00 ± 0.119	26.92 ± 0.141	$+0.91 \pm 0.084$	23.53 ± 0.238	24.48 ± 0.126	$+0.95 \pm 0.155$	23.87 ± 0.078	24.61 ± 0.072	$+0.73 \pm 0.054$	15.62 ± 0.339	16.69 ± 0.326	$+1.07 \pm 0.023$
Peppers (0.2)	29.82 ± 0.234	30.80 ± 0.120	$+0.97 \pm 0.203$	28.52 ± 0.187	29.44 ± 0.142	$+0.91 \pm 0.097$	27.94 ± 0.151	28.70 ± 0.135	$+0.75 \pm 0.048$	21.62 ± 0.120	22.41 ± 0.108	$+0.79 \pm 0.023$
House (0.1)	25.36 ± 0.037	25.76 ± 0.076	$+0.40 \pm 0.047$	22.79 ± 0.173	23.57 ± 0.133	$+0.78 \pm 0.071$	24.14 ± 0.071	24.48 ± 0.047	$+0.33 \pm 0.046$	17.67 ± 0.098	18.75 ± 0.091	$+1.08 \pm 0.015$
House (0.2)	28.67 ± 0.066	29.01 ± 0.064	$+0.34 \pm 0.046$	26.75 ± 0.092	27.43 ± 0.074	$+0.67 \pm 0.062$	27.31 ± 0.085	27.59 ± 0.098	$+0.27 \pm 0.056$	22.89 ± 0.073	23.44 ± 0.081	$+0.54 \pm 0.019$
Sailboat (0.1)	25.03 ± 0.078	25.38 ± 0.077	$+0.34 \pm 0.053$	23.38 ± 0.080	23.96 ± 0.042	$+0.58 \pm 0.084$	23.34 ± 0.205	23.67 ± 0.135	$+0.32 \pm 0.080$	16.62 ± 0.066	17.61 ± 0.077	$+0.98 \pm 0.021$
Sailboat (0.2)	28.47 ± 0.113	28.89 ± 0.131	$+0.41 \pm 0.074$	27.11 ± 0.049	27.63 ± 0.032	$+0.52 \pm 0.025$	26.77 ± 0.190	27.08 ± 0.161	$+0.30 \pm 0.039$	21.76 ± 0.352	22.37 ± 0.391	$+0.60 \pm 0.048$
News (0.1)	32.47 ± 0.072	32.75 ± 0.067	$+0.28 \pm 0.023$	29.09 ± 0.073	29.20 ± 0.075	$+0.11 \pm 0.003$	30.03 ± 0.400	30.26 ± 0.398	$+0.20 \pm 0.013$	24.62 ± 0.048	24.93 ± 0.059	$+0.31 \pm 0.018$
News (0.2)	34.72 ± 0.030	35.20 ± 0.040	$+0.47 \pm 0.013$	31.74 ± 0.072	32.17 ± 0.051	$+0.42 \pm 0.039$	33.05 ± 0.332	33.36 ± 0.322	$+0.28 \pm 0.023$	25.81 ± 0.398	26.42 ± 0.419	$+0.61 \pm 0.035$
Carphone (0.1)	30.75 ± 0.061	31.63 ± 0.049	$+0.87 \pm 0.019$	28.98 ± 0.022	29.25 ± 0.027	$+0.26 \pm 0.021$	28.99 ± 0.132	29.39 ± 0.134	$+0.39 \pm 0.018$	25.00 ± 0.039	25.76 ± 0.073	$+0.76 \pm 0.045$
Carphone (0.2)	32.49 ± 0.023	33.61 ± 0.023	$+1.11 \pm 0.016$	30.89 ± 0.049	31.56 ± 0.024	$+0.66 \pm 0.055$	31.71 ± 0.322	32.31 ± 0.221	$+0.60 \pm 0.106$	25.74 ± 0.027	27.12 ± 0.094	$+1.37 \pm 0.085$
Foreman (0.1)	27.66 ± 0.059	28.52 ± 0.040	$+0.85 \pm 0.025$	25.83 ± 0.056	26.24 ± 0.047	$+0.40 \pm 0.010$	24.71 ± 0.257	25.00 ± 0.261	$+0.28 \pm 0.010$	21.94 ± 0.048	22.88 ± 0.025	$+0.93 \pm 0.032$
Foreman (0.2)	29.42 ± 0.018	30.58 ± 0.033	$+1.16 \pm 0.025$	27.45 ± 0.078	28.21 ± 0.059	$+0.75 \pm 0.033$	28.17 ± 0.075	28.65 ± 0.066	$+0.47 \pm 0.027$	22.61 ± 0.030	24.02 ± 0.133	$+1.41 \pm 0.125$
Grandma (0.1)	36.99 ± 0.054	37.37 ± 0.045	$+0.38 \pm 0.046$	33.86 ± 0.020	34.13 ± 0.057	$+0.27 \pm 0.051$	34.69 ± 0.157	35.17 ± 0.107	$+0.47 \pm 0.053$	28.57 ± 0.080	28.95 ± 0.102	$+0.38 \pm 0.033$
Grandma (0.2)	39.34 ± 0.035	40.04 ± 0.033	$+0.67 \pm 0.051$	36.06 ± 0.075	36.62 ± 0.029	$+0.55 \pm 0.067$	36.99 ± 0.405	37.70 ± 0.303	$+0.70 \pm 0.118$	30.34 ± 0.392	30.93 ± 0.435	$+0.59 \pm 0.046$
Akiyo (0.1)	36.36 ± 0.122	36.92 ± 0.111	$+0.56 \pm 0.025$	32.85 ± 0.049	33.06 ± 0.058	$+0.21 \pm 0.019$	33.28 ± 0.502	33.92 ± 0.445	$+0.64 \pm 0.086$	28.44 ± 0.129	29.10 ± 0.172	$+0.66 \pm 0.159$
Akiyo (0.2)	38.99 ± 0.067	40.10 ± 0.020	$+1.10 \pm 0.07$	35.28 ± 0.156	36.13 ± 0.088	$+0.85 \pm 0.081$	36.67 ± 0.251	37.62 ± 0.162	$+0.95 \pm 0.102$	29.23 ± 0.027	30.53 ± 0.045	$+1.29 \pm 0.042$
Toy (0.1)	43.91 ± 0.202	45.65 ± 0.123	$+1.73 \pm 0.204$	40.94 ± 0.205	43.46 ± 0.152	$+2.51 \pm 0.029$	40.76 ± 0.383	42.08 ± 0.139	$+1.32 \pm 0.436$	31.08 ± 0.140	31.48 ± 0.187	$+0.40 \pm 0.051$
Toy (0.2)	48.40 ± 0.229	52.08 ± 0.145	$+3.67 \pm 0.112$	43.60 ± 0.309	48.00 ± 0.165	$+4.40 \pm 0.254$	44.72 ± 0.961	46.99 ± 0.203	$+2.27 \pm 0.791$	32.18 ± 0.160	32.90 ± 0.185	$+0.71 \pm 0.027$
Balloons (0.1)	49.50 ± 0.558	52.19 ± 0.456	$+2.69 \pm 0.491$	46.57 ± 0.332	50.79 ± 0.132	$+4.21 \pm 0.417$	44.77 ± 0.321	46.10 ± 0.425	$+1.33 \pm 0.167$	32.91 ± 0.973	33.81 ± 0.872	$+0.90 \pm 0.103$
Balloons (0.2)	51.81 ± 0.445	53.18 ± 0.178	$+4.37 \pm 0.423$	48.32 ± 0.275	54.55 ± 0.174	$+6.23 \pm 0.211$	48.48 ± 0.213	50.82 ± 0.175	$+2.34 \pm 0.235$	35.50 ± 0.566	37.00 ± 0.602	$+1.49 \pm 0.053$
Face (0.1)	48.03 ± 0.163	49.30 ± 0.133	$+1.26 \pm 0.095$	46.28 ± 0.287	48.17 ± 0.161	$+1.89 \pm 0.250$	43.79 ± 0.221	45.10 ± 0.156	$+1.30 \pm 0.101$	33.56 ± 0.150	34.03 ± 0.178	$+0.46 \pm 0.038$
Face (0.2)	51.90 ± 0.121	54.47 ± 0.186	$+2.57 \pm 0.206$	49.04 ± 0.331	52.64 ± 0.111	$+3.59 \pm 0.258$	47.86 ± 0.151	49.53 ± 0.315	$+1.67 \pm 0.192$	38.46 ± 0.130	39.35 ± 0.153	$+0.88 \pm 0.048$
Flowers (0.1)	45.08 ± 0.176	46.18 ± 0.104	$+1.09 \pm 0.134$	44.43 ± 0.209	46.11 ± 0.073	$+1.67 \pm 0.189$	42.08 ± 0.081	42.50 ± 0.134	$+0.41 \pm 0.059$	31.00 ± 1.490	31.28 ± 1.496	$+0.27 \pm 0.019$
Flowers (0.2)	49.86 ± 0.168	51.95 ± 0.050	$+2.08 \pm 0.176$	47.64 ± 0.360	51.69 ± 0.112	$+4.05 \pm 0.273$	46.39 ± 0.069	47.23 ± 0.047	$+0.83 \pm 0.062$	33.18 ± 0.023	33.80 ± 0.032	$+0.61 \pm 0.018$
Washington DC (0.1)	45.92 ± 0.048	46.87 ± 0.040	$+0.94 \pm 0.055$	39.67 ± 0.270	42.33 ± 0.148	$+2.66 \pm 0.223$	42.03 ± 0.557	42.42 ± 0.532	$+0.39 \pm 0.041$	28.40 ± 0.388	28.62 ± 0.336	$+0.22 \pm 0.053$
Washington DC (0.2)	48.31 ± 0.069	49.64 ± 0.025	$+1.33 \pm 0.056$	40.61 ± 0.581	44.42 ± 0.075	$+3.80 \pm 0.608$	46.55 ± 0.403	47.21 ± 0.323	$+0.65 \pm 0.330$	30.11 ± 0.036	30.46 ± 0.048	$+0.34 \pm 0.022$
Botswana (0.1)	42.88 ± 0.036	43.22 ± 0.020	$+0.33 \pm 0.016$	37.58 ± 0.273	39.49 ± 0.211	$+1.90 \pm 0.125$	40.36 ± 0.259	40.69 ± 0.220	$+0.32 \pm 0.130$	28.79 ± 0.559	29.06 ± 0.511	$+0.27 \pm 0.072$
Botswana (0.2)	45.35 ± 0.036	46.08 ± 0.019	$+0.72 \pm 0.041$	39.09 ± 0.236	41.35 ± 0.039	$+2.25 \pm 0.238$	43.91 ± 0.367	44.13 ± 0.312	$+0.21 \pm 0.055$	30.21 ± 0.084	30.71 ± 0.101	$+0.49 \pm 0.024$
Salinas (0.1)	51.57 ± 0.057	53.42 ± 0.036	$+1.85 \pm 0.090$	40.47 ± 0.084	44.66 ± 0.105	$+4.18 \pm 0.069$	45.73 ± 0.195	46.56 ± 0.212	$+0.82 \pm 0.136$	33.67 ± 1.485	34.63 ± 1.500	$+0.96 \pm 0.129$
Salinas (0.2)	53.05 ± 0.129	55.86 ± 0.094	$+2.80 \pm 0.137$	40.87 ± 0.551	45.76 ± 0.479	$+4.88 \pm 0.157$	49.77 ± 0.239	50.87 ± 0.228	$+1.09 \pm 0.113$	35.93 ± 0.519	37.69 ± 0.491	$+1.76 \pm 0.037$

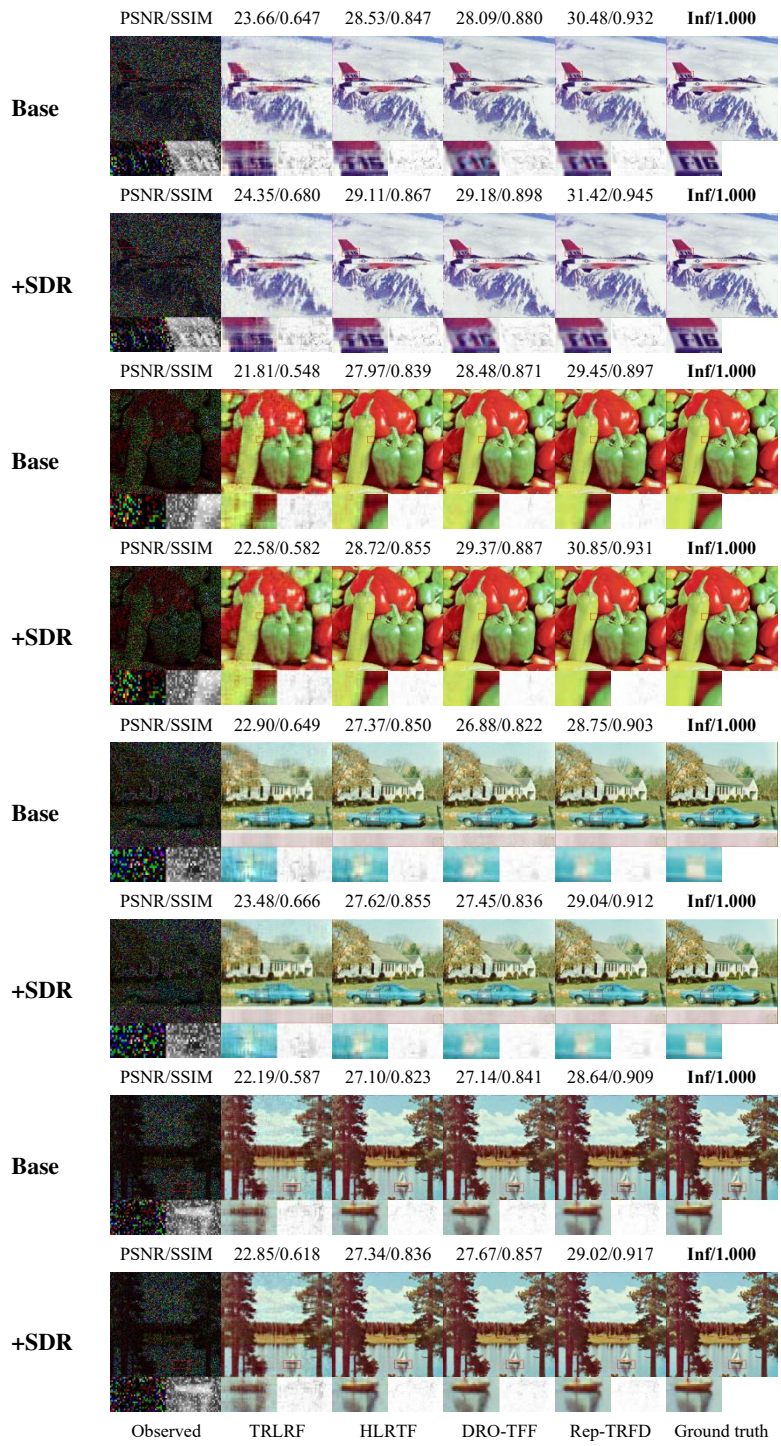


Figure C.1: Visual comparison on RGB images Airplane(20%), Peppers(20%), House(20%) and Sailboat(20%). PSNR and SSIM are reported for the base reconstructions and the corresponding SDR-refined results.

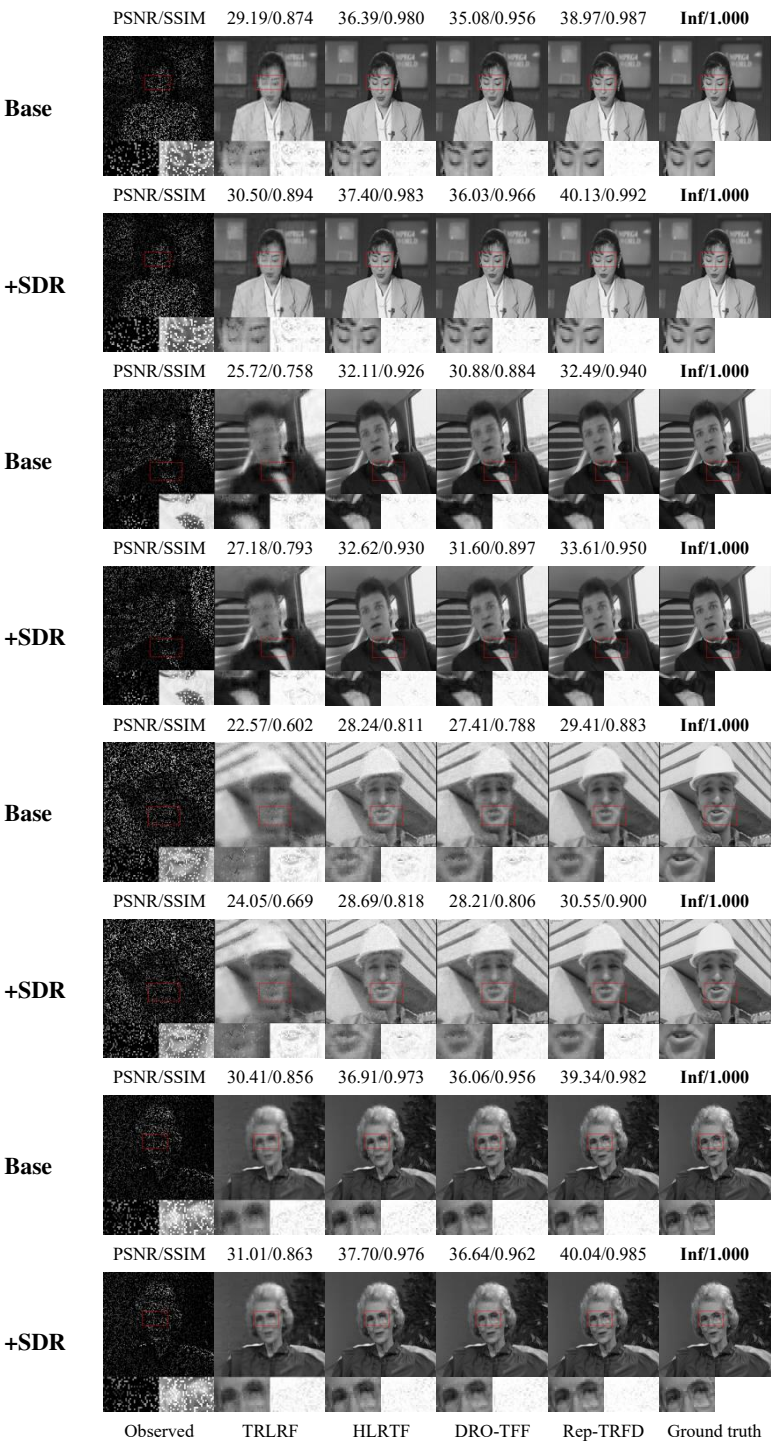


Figure C.2: Visual comparison on video data Akiyo(20%), Carphone(20%), Foreman(20%) and Grandma(20%). PSNR and SSIM are reported for the base reconstructions and the corresponding SDR-refined results.

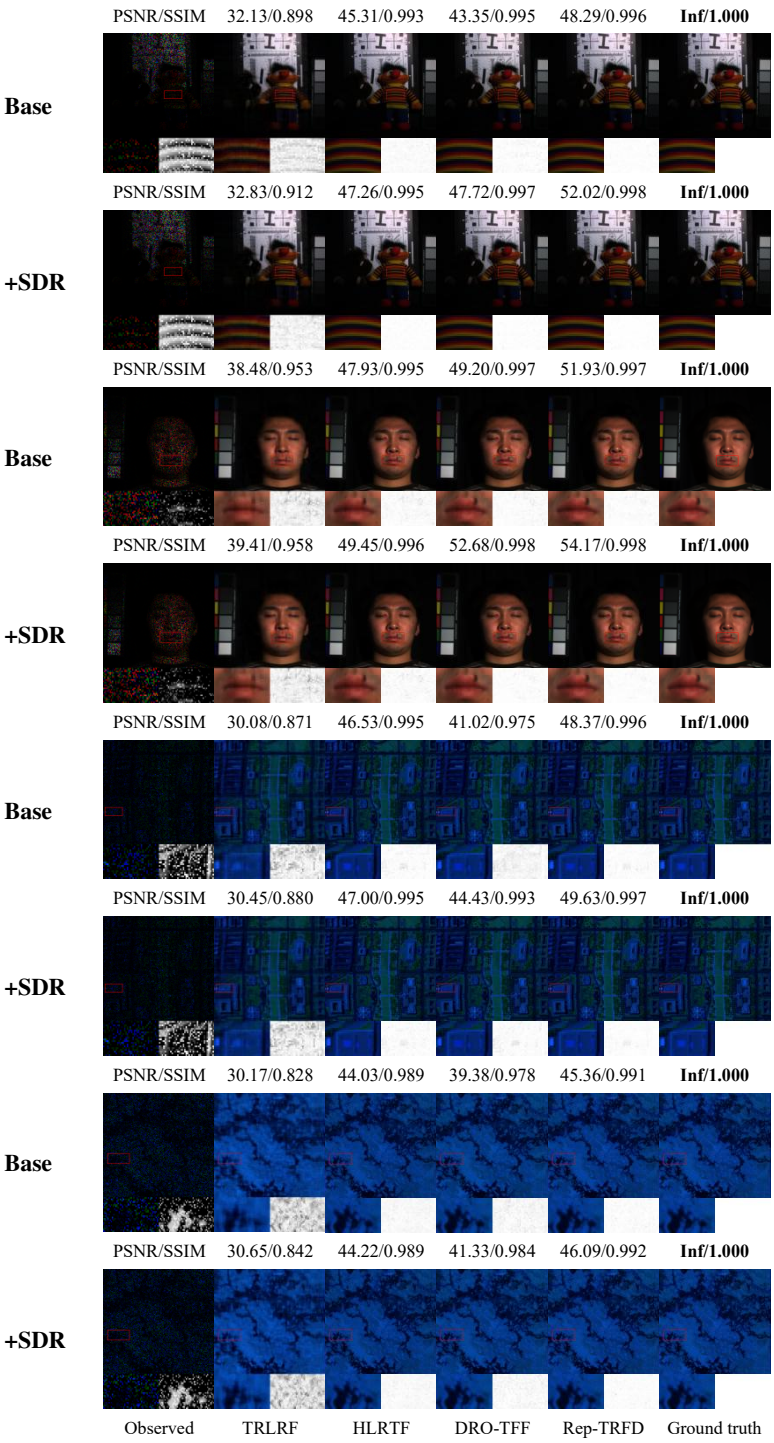


Figure C.3: Visual comparison on multispectral and hyperspectral images Toy(20%), Face(20%), Washington DC(20%) and Botswana(20%). PSNR and SSIM are reported for the base reconstructions and the corresponding SDR-refined results.

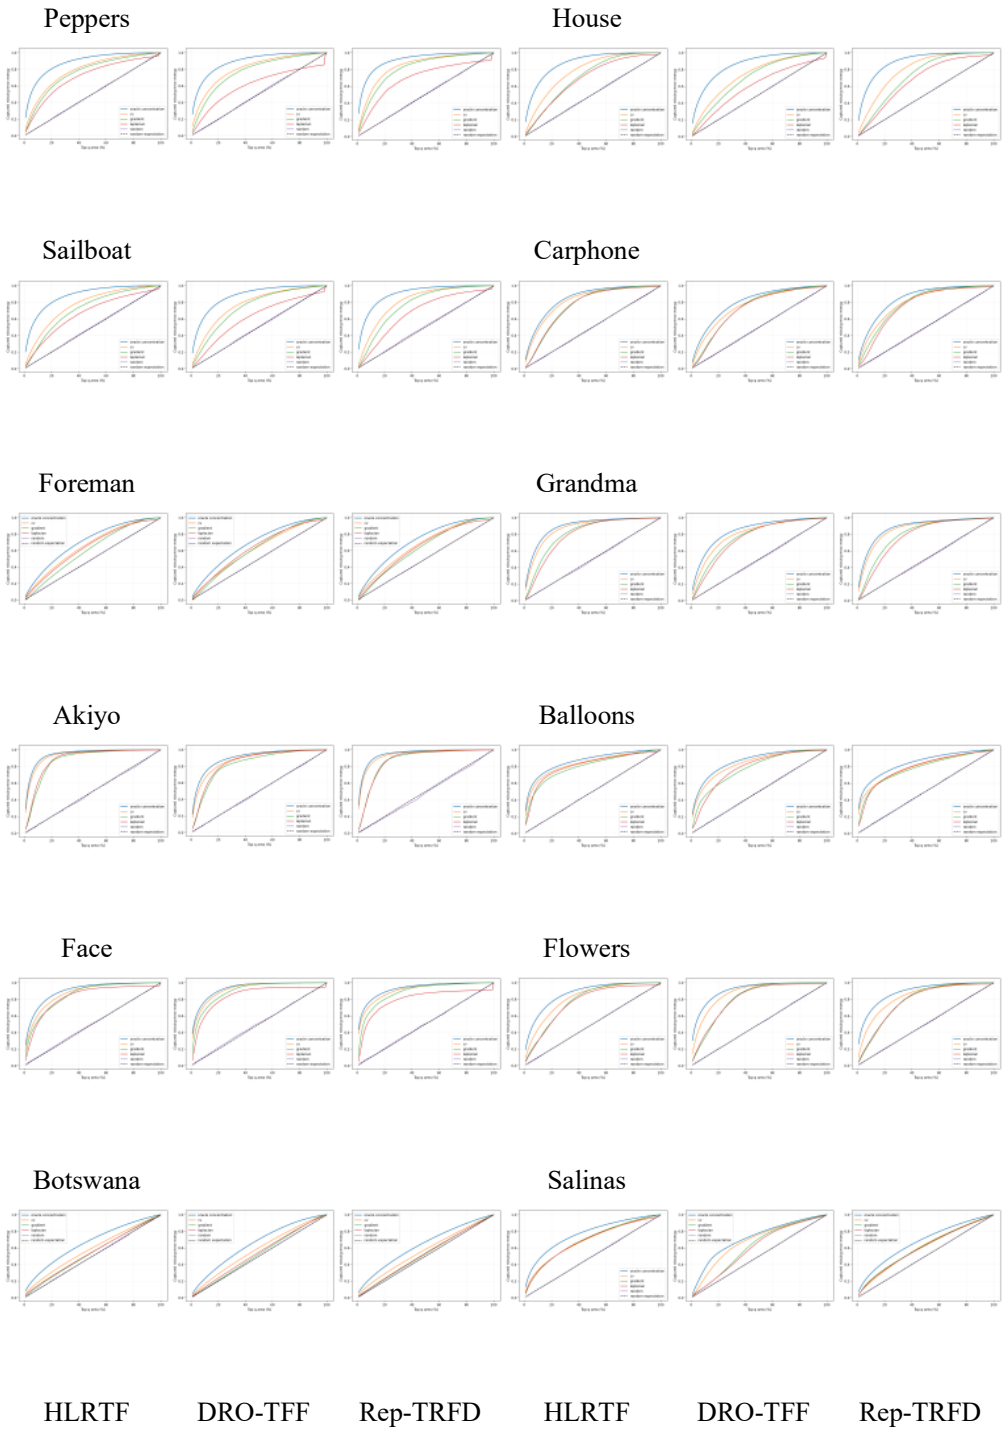


Figure C.4: Extended analysis of post-reconstruction error concentration and GT-free error localization across different signals with 20% sampling ratio and reconstruction settings. CV uncertainty generally provides stronger error capture than the gradient- and Laplacian-based indicators, while remaining comparable in cases where the different indicators exhibit similar behavior.

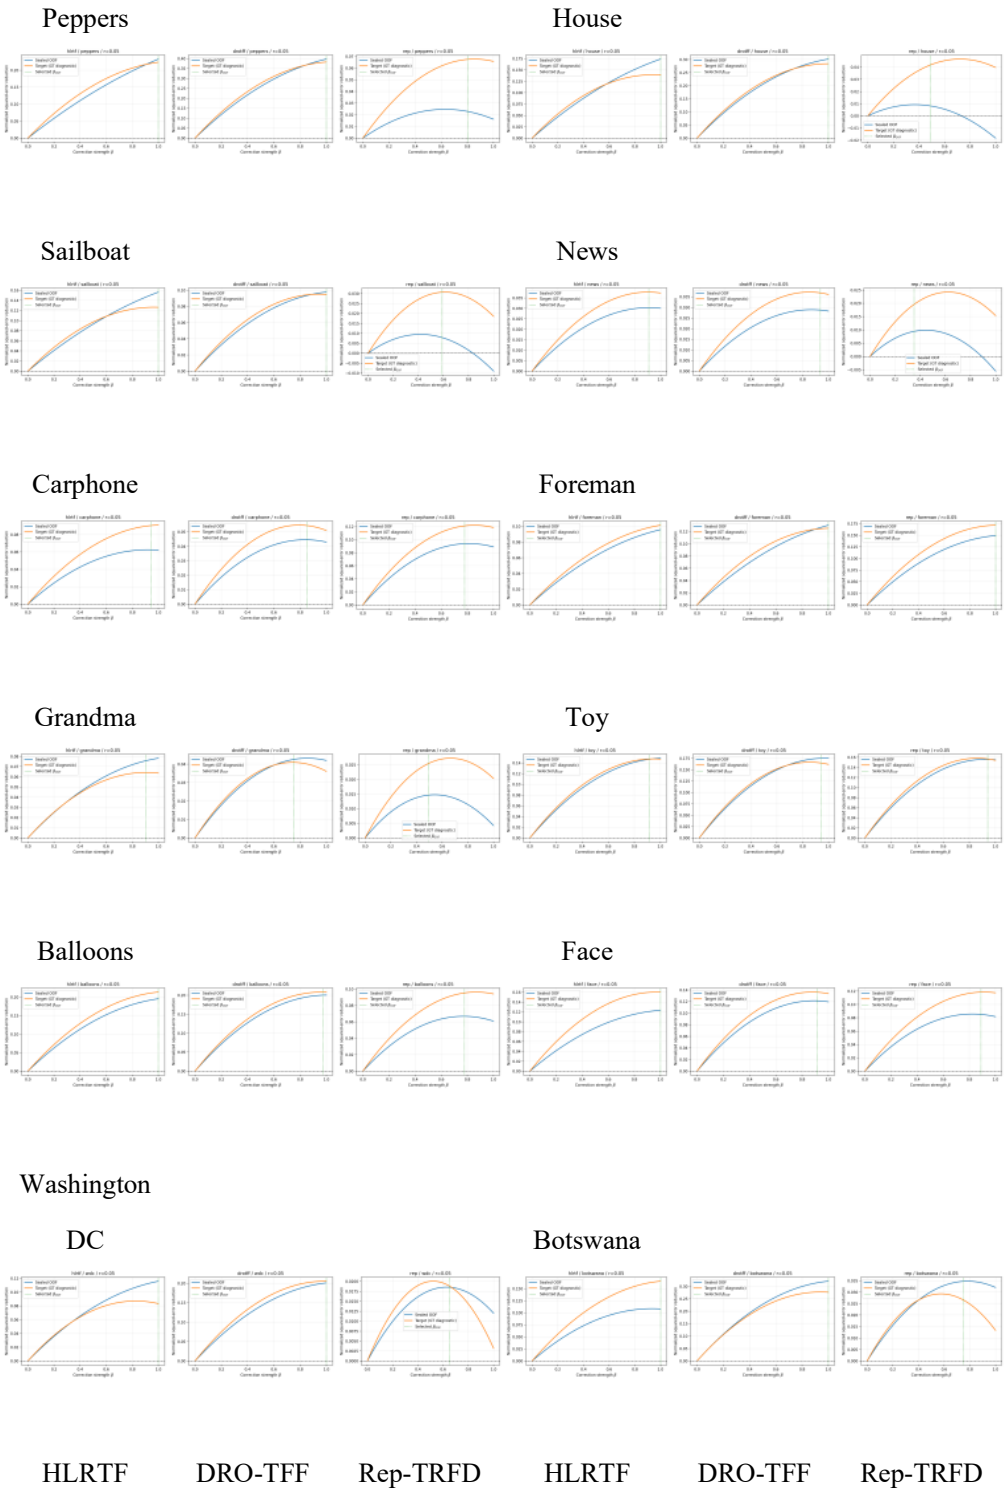


Figure C.5: Comparison of the relative error reduction on O-O-F and target regions under different correction strengths β on different signals with 5% sampling ratio. The two regions tend to exhibit similar variation trends across different cases, and the selected β consistently yields effective correction on the target region.

Table C.2: Complete validation of the O-O-F-based optimization gate across different signals and sampling ratios. "Accept?" denotes whether continued optimization is accepted by the gate, ΔPSNR denotes the actual PSNR change after continued optimization, and "Gate Correct?" indicates whether the gating decision agrees with the observed performance change.

	Rep-TRFD			DRO-TFF			HLRTF		
Setting	Accept?	ΔPSNR	Gate Correct?	Accept?	ΔPSNR	Gate Correct?	Accept?	ΔPSNR	Gate Correct?
Airplane (0.1)	Yes	+0.383	✓	Yes	-0.057	×	No	-0.049	✓
Airplane (0.2)	Yes	+0.400	✓	Yes	+0.126	✓	No	-0.046	✓
Peppers (0.1)	Yes	+0.446	✓	Yes	+0.264	✓	Yes	+0.252	✓
Peppers (0.2)	Yes	+0.592	✓	Yes	+0.033	✓	No	-0.036	✓
House (0.1)	Yes	+0.140	✓	Yes	+0.053	✓	Yes	+0.049	✓
House (0.2)	Yes	+0.064	✓	Yes	+0.104	✓	No	-0.031	✓
Sailboat (0.1)	Yes	+0.186	✓	Yes	+0.155	✓	No	-0.042	✓
Sailboat (0.2)	Yes	+0.159	✓	Yes	+0.077	✓	Yes	-0.040	×
News (0.1)	Yes	+0.189	✓	No	-0.054	✓	Yes	+0.052	✓
News (0.2)	Yes	+0.335	✓	Yes	+0.271	✓	Yes	+0.086	✓
Carphone (0.1)	Yes	+0.154	✓	Yes	+0.061	✓	Yes	+0.014	✓
Carphone (0.2)	Yes	+0.246	✓	Yes	+0.282	✓	Yes	+0.031	✓
Foreman (0.1)	Yes	+0.178	✓	Yes	+0.042	✓	Yes	+0.025	✓
Foreman (0.2)	Yes	+0.262	✓	Yes	+0.112	✓	Yes	+0.105	✓
Grandma (0.1)	Yes	+0.317	✓	Yes	+0.104	✓	Yes	+0.055	✓
Grandma (0.2)	Yes	+0.536	✓	Yes	+0.364	✓	Yes	+0.201	✓
Akiyo (0.1)	Yes	+0.394	✓	No	+0.090	×	Yes	+0.217	✓
Akiyo (0.2)	Yes	+0.941	✓	Yes	+0.526	✓	Yes	+0.213	✓
Toy (0.1)	Yes	+1.207	✓	Yes	+1.190	✓	Yes	+0.230	✓
Toy (0.2)	Yes	+3.300	✓	Yes	+2.719	✓	Yes	+1.021	✓
Balloons (0.1)	Yes	+2.322	✓	Yes	+2.555	✓	Yes	+0.046	✓
Balloons (0.2)	Yes	+4.252	✓	Yes	+3.830	✓	Yes	+0.275	✓
Face (0.1)	Yes	+0.602	✓	Yes	+0.765	✓	No	-0.045	✓
Face (0.2)	Yes	+1.512	✓	Yes	+2.403	✓	Yes	+0.210	✓
Flowers (0.1)	Yes	+0.532	✓	Yes	+1.338	✓	Yes	+0.072	✓
Flowers (0.2)	Yes	+1.685	✓	Yes	+3.735	✓	Yes	+0.214	✓
Washington DC (0.1)	Yes	+0.954	✓	Yes	+2.056	✓	Yes	+0.019	✓
Washington DC (0.2)	Yes	+1.141	✓	Yes	+2.158	✓	Yes	+0.131	✓
Botswana (0.1)	Yes	+0.310	✓	Yes	+0.957	✓	Yes	+0.336	✓
Botswana (0.2)	Yes	+0.686	✓	Yes	+0.885	✓	Yes	+0.070	✓
Salinas (0.1)	Yes	+1.728	✓	Yes	+0.546	✓	Yes	+0.001	✓
Salinas (0.2)	Yes	+2.753	✓	Yes	+0.862	✓	Yes	-0.023	×

Table C.3: Complete mechanism validation of residual transfer across different signals and sampling ratios. Original, Spatial, Random, and Reverse denote the proposed structural residual transfer, spatial nearest-neighbor transfer, random residual matching, and reversed residual transfer, respectively. All entries report the PSNR change after residual transfer, and the best result within each reconstruction method and setting is highlighted in bold.

	Rep-TRFD				DRO-TFF				HLRTF				TRLRF			
Setting	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse	Structural	Spatial	Random	Reverse
Airplane (0.05)	**+0.278**	-0.020	-0.058	-0.571	+1.250	**+1.318**	+0.307	-2.254	**+0.914**	+0.547	+0.043	-2.036	+0.981	**+1.011**	+0.139	-1.285
Airplane (0.2)	**+0.538**	-0.113	-0.139	-1.448	**+0.961**	+0.255	-0.047	-2.199	**+0.580**	-0.051	-0.127	-1.304	**+0.697**	+0.319	-0.114	-1.291
Peppers (0.05)	**+0.265**	+0.081	-0.048	-0.648	+0.961	**+1.017**	+0.021	-1.880	**+0.674**	+0.426	-0.052	-1.265	+0.711	**+0.777**	-0.005	-0.929
Peppers (0.2)	**+0.689**	+0.025	-0.138	-1.540	**+0.859**	+0.347	-0.096	-1.550	**+0.753**	+0.126	-0.129	-1.352	**+0.771**	+0.717	-0.109	-1.263
House (0.05)	**+0.147**	+0.019	-0.028	-0.282	+0.850	**+1.064**	+0.346	-1.862	**+0.464**	+0.259	-0.072	-1.238	**+0.755**	+0.699	-0.077	-1.146
House (0.2)	**+0.220**	-0.078	-0.074	-0.679	**+0.467**	+0.057	-0.112	-1.446	**+0.252**	-0.026	-0.121	-0.943	**+0.581**	+0.319	-0.108	-1.274
Sailboat (0.05)	**+0.114**	-0.009	-0.036	-0.308	+0.333	**+0.523**	+0.054	-0.936	**+0.387**	+0.119	-0.096	-1.010	**+0.866**	+0.573	-0.049	-1.272
Sailboat (0.2)	**+0.221**	-0.094	-0.077	-0.704	**+0.445**	+0.203	-0.076	-1.276	**+0.289**	-0.082	-0.140	-0.992	**+0.657**	+0.274	-0.117	-1.206
News (0.05)	**+0.085**	-0.021	-0.017	-0.146	**+0.154**	-0.120	-0.103	-0.429	**+0.163**	-0.122	-0.114	-0.445	**+0.153**	+0.109	-0.050	-0.330
News (0.2)	**+0.162**	-0.005	-0.017	-0.261	**+0.196**	-0.032	-0.036	-0.365	**+0.212**	+0.023	-0.065	-0.475	+0.556	**+0.561**	-0.100	-1.036
Carphone (0.05)	**+0.508**	-0.060	-0.121	-1.124	**+0.200**	-0.139	-0.112	-0.579	**+0.362**	-0.046	-0.126	-0.843	**+0.392**	+0.242	-0.123	-0.908
Carphone (0.2)	**+0.872**	+0.003	-0.148	-1.690	**+0.439**	-0.031	-0.085	-0.883	**+0.478**	+0.041	-0.104	-1.025	**+1.459**	+1.082	-0.122	-2.180
Foreman (0.05)	**+0.523**	-0.112	-0.155	-1.245	**+0.334**	-0.127	-0.144	-0.762	**+0.283**	-0.001	-0.120	-0.603	**+0.549**	+0.196	-0.119	-1.158
Foreman (0.2)	**+0.879**	-0.052	-0.144	-1.777	**+0.690**	-0.067	-0.139	-1.473	**+0.347**	+0.007	-0.129	-0.915	**+1.485**	+0.984	-0.110	-2.342
Grandma (0.05)	**+0.107**	-0.059	-0.038	-0.226	**+0.163**	-0.117	-0.084	-0.460	**+0.268**	-0.145	-0.122	-0.715	**+0.274**	+0.064	-0.083	-0.715
Grandma (0.2)	**+0.172**	-0.098	-0.046	-0.357	**+0.219**	-0.098	-0.061	-0.470	**+0.588**	-0.076	-0.145	-1.264	**+0.601**	+0.242	-0.115	-1.285
Akiyo (0.05)	**+0.143**	-0.076	-0.042	-0.296	**+0.201**	-0.217	-0.115	-0.526	**+0.336**	-0.193	-0.134	-0.776	**+0.161**	-0.290	-0.171	-0.502
Akiyo (0.2)	**+0.218**	-0.066	-0.025	-0.382	**+0.426**	-0.140	-0.078	-0.876	**+0.800**	-0.125	-0.149	-1.649	**+1.308**	+0.138	-0.110	-1.987
Toy (0.05)	**+0.731**	-0.138	-0.147	-1.828	**+0.775**	+0.194	-0.079	-1.947	**+0.652**	+0.437	-0.089	-1.437	**+0.151**	+0.077	-0.138	-0.729
Toy (0.2)	**+0.435**	-0.081	-0.118	-1.167	+1.653	**+1.704**	+0.278	-2.924	**+0.926**	+0.597	-0.128	-2.102	**+0.698**	+0.616	-0.105	-1.443
Balloons (0.05)	**+0.419**	+0.075	-0.082	-0.925	**+0.939**	+0.790	+0.063	-2.014	+0.936	**+0.950**	-0.147	-1.793	+0.622	**+1.425**	-0.096	-1.225
Balloons (0.2)	**+0.365**	+0.001	-0.093	-0.866	**+2.424**	+2.203	+0.620	-3.660	**+1.850**	+1.248	-0.147	-2.935	+1.443	**+2.317**	-0.093	-2.020
Face (0.05)	**+0.541**	-0.041	-0.110	-1.267	**+0.629**	+0.156	-0.096	-1.547	**+0.724**	+0.674	-0.158	-1.698	+0.337	**+1.110**	-0.088	-0.852
Face (0.2)	**+0.727**	-0.005	-0.129	-1.830	**+1.076**	+0.484	-0.077	-2.399	**+1.310**	+0.791	-0.132	-2.625	+0.939	**+1.159**	-0.106	-1.766
Flowers (0.05)	**+0.421**	-0.100	-0.089	-1.036	**+0.476**	-0.019	-0.092	-1.234	**+0.281**	+0.220	-0.082	-0.834	+0.084	**+0.150**	-0.098	-0.671
Flowers (0.2)	**+0.227**	-0.045	-0.078	-0.652	**+0.293**	+0.146	-0.044	-0.797	+0.570	**+0.622**	-0.101	-1.448	+0.624	**+0.851**	-0.091	-1.317
Washington DC (0.05)	**+0.060**	-0.056	-0.067	-0.248	**+0.677**	+0.266	+0.052	-1.576	**+0.251**	+0.095	-0.150	-0.893	**+0.121**	+0.054	-0.095	-0.485
Washington DC (0.2)	**+0.121**	+0.033	-0.059	-0.320	**+1.251**	+1.245	+0.382	-2.356	**+0.336**	+0.321	-0.134	-1.039	+0.369	**+0.430**	-0.111	-0.783
Botswana (0.05)	**+0.070**	-0.067	-0.074	-0.316	**+0.741**	+0.469	+0.358	-1.791	**+0.507**	+0.107	-0.128	-1.020	**+0.365**	+0.354	-0.110	-0.845
Botswana (0.2)	**+0.040**	+0.027	-0.017	-0.117	**+1.064**	+1.030	+0.801	-2.298	+0.120	**+0.209**	-0.074	-0.399	**+0.484**	+0.481	-0.110	-1.056
Salinas (0.05)	**+0.094**	-0.039	-0.044	-0.242	**+3.000**	+2.653	+1.772	-4.028	**+1.200**	+0.845	-0.098	-2.079	+0.432	**+1.729**	-0.125	-0.938
Salinas (0.2)	**+0.162**	-0.006	-0.023	-0.369	**+3.910**	+3.768	+2.644	-4.562	**+1.006**	+0.568	-0.135	-2.321	**+1.805**	+1.070	-0.097	-2.378

Table C.4: Complete ablation results at sampling ratios 0.1 and 0.2. Base denotes the original reconstruction; Opt. denotes continued optimization; Gated Opt. applies the O-O-F-based optimization gate; Resid. applies residual transfer without continued optimization; w/o Gate combines continued optimization and residual transfer without the gate; and Full denotes the complete SDR framework. All entries report PSNR (dB). NA indicates that the optimization-related component is not applicable to TRLRF.

Sampling Ratio = 0.1								Sampling Ratio = 0.2							
Dataset	Method	Base	Opt.	Gated Opt.	Resid.	w/o Gate	Full	Dataset	Method	Base	Opt.	Gated Opt.	Resid.	w/o Gate	Full
Airplane	Rep-TRFD	26.29	26.68	26.68	26.79	27.08	27.08	Airplane	Rep-TRFD	30.48	30.88	30.88	31.10	31.42	31.42
	DRO-TFF	23.09	23.04	23.04	24.56	25.01	25.01		DRO-TFF	28.10	28.22	28.22	28.96	29.18	29.18
	HLRTF	25.54	25.49	25.54	25.98	25.94	25.98		HLRTF	28.53	28.48	28.53	29.11	29.05	29.11
	TRLRF	17.42	NA	NA	18.55	NA	18.55		TRLRF	23.66	NA	NA	24.36	NA	24.36
Peppers	Rep-TRFD	26.03	26.48	26.48	26.60	26.98	26.98	Peppers	Rep-TRFD	29.57	30.17	30.17	30.25	30.86	30.86
	DRO-TFF	23.52	23.78	23.78	24.34	24.54	24.54		DRO-TFF	28.48	28.52	28.52	29.25	29.37	29.37
	HLRTF	23.86	24.11	24.11	24.48	24.64	24.64		HLRTF	27.98	27.94	27.98	28.73	28.65	28.73
	TRLRF	15.07	NA	NA	16.17	NA	16.17		TRLRF	21.81	NA	NA	22.58	NA	22.58
House	Rep-TRFD	25.37	25.51	25.51	25.66	25.75	25.75	House	Rep-TRFD	28.76	28.82	28.82	28.97	29.04	29.04
	DRO-TFF	23.01	23.07	23.07	23.70	23.77	23.77		DRO-TFF	26.89	26.99	26.99	27.37	27.46	27.46
	HLRTF	24.06	24.11	24.11	24.39	24.45	24.45		HLRTF	27.37	27.34	27.37	27.63	27.57	27.63
	TRLRF	17.73	NA	NA	18.81	NA	18.81		TRLRF	22.90	NA	NA	23.48	NA	23.48
Sailboat	Rep-TRFD	25.03	25.22	25.22	25.23	25.36	25.36	Sailboat	Rep-TRFD	28.64	28.80	28.80	28.88	29.03	29.03
	DRO-TFF	23.37	23.52	23.52	23.80	23.94	23.94		DRO-TFF	27.15	27.23	27.23	27.55	27.67	27.67
	HLRTF	23.46	23.42	23.46	23.76	23.74	23.76		HLRTF	27.10	27.06	27.06	27.41	27.35	27.35
	TRLRF	16.66	NA	NA	17.62	NA	17.62		TRLRF	22.20	NA	NA	22.85	NA	22.85
News	Rep-TRFD	32.48	32.67	32.67	32.62	32.78	32.78	News	Rep-TRFD	34.76	35.10	35.10	34.95	35.26	35.26
	DRO-TFF	29.21	29.16	29.21	29.33	29.28	29.33		DRO-TFF	31.74	32.01	32.01	31.93	32.20	32.20
	HLRTF	30.40	30.45	30.45	30.57	30.61	30.61		HLRTF	33.19	33.28	33.28	33.42	33.49	33.49
	TRLRF	24.58	NA	NA	24.89	NA	24.89		TRLRF	25.54	NA	NA	26.09	NA	26.09
Carphone	Rep-TRFD	30.81	30.97	30.97	31.43	31.67	31.67	Carphone	Rep-TRFD	32.50	32.75	32.75	33.37	33.62	33.62
	DRO-TFF	29.02	29.08	29.08	29.23	29.28	29.28		DRO-TFF	30.88	31.16	31.16	31.34	31.60	31.60
	HLRTF	29.06	29.08	29.08	29.44	29.44	29.44		HLRTF	32.11	32.14	32.14	32.62	32.62	32.62
	TRLRF	24.98	NA	NA	25.66	NA	25.66		TRLRF	25.72	NA	NA	27.18	NA	27.18
Foreman	Rep-TRFD	27.69	27.87	27.87	28.31	28.55	28.55	Foreman	Rep-TRFD	29.42	29.68	29.68	30.25	30.56	30.56
	DRO-TFF	25.81	25.86	25.86	26.18	26.23	26.23		DRO-TFF	27.42	27.53	27.53	28.09	28.22	28.22
	HLRTF	25.06	25.08	25.08	25.35	25.36	25.36		HLRTF	28.24	28.35	28.35	28.63	28.69	28.69
	TRLRF	21.94	NA	NA	22.87	NA	22.87		TRLRF	22.57	NA	NA	24.06	NA	24.06
Grandma	Rep-TRFD	36.94	37.26	37.26	37.07	37.39	37.39	Grandma	Rep-TRFD	39.34	39.88	39.88	39.52	40.05	40.05
	DRO-TFF	33.84	33.94	33.94	33.96	34.06	34.06		DRO-TFF	36.06	36.43	36.43	36.29	36.65	36.65
	HLRTF	34.77	34.83	34.83	35.16	35.22	35.22		HLRTF	36.91	37.11	37.11	37.56	37.70	37.70
	TRLRF	28.59	NA	NA	28.93	NA	28.93		TRLRF	30.41	NA	NA	31.01	NA	31.01
Akiyo	Rep-TRFD	36.48	36.87	36.87	36.71	37.02	37.02	Akiyo	Rep-TRFD	38.97	39.91	39.91	39.14	40.13	40.13
	DRO-TFF	32.77	32.86	32.77	32.98	33.06	32.98		DRO-TFF	35.09	35.61	35.61	35.54	36.04	36.04
	HLRTF	33.93	34.14	34.14	34.37	34.55	34.55		HLRTF	36.39	36.60	36.60	37.23	37.40	37.40
	TRLRF	28.52	NA	NA	28.90	NA	28.90		TRLRF	29.19	NA	NA	30.50	NA	30.50
Toy	Rep-TRFD	43.68	44.88	44.88	43.98	45.53	45.53	Toy	Rep-TRFD	48.29	51.59	51.59	48.76	52.03	52.03
	DRO-TFF	41.09	42.28	42.28	41.95	43.28	43.28		DRO-TFF	43.35	46.07	46.07	44.96	47.72	47.72
	HLRTF	40.89	41.12	41.12	41.69	41.85	41.85		HLRTF	45.32	46.34	46.34	46.45	47.27	47.27
	TRLRF	31.22	NA	NA	31.69	NA	31.69		TRLRF	32.14	NA	NA	32.84	NA	32.84
Balloons	Rep-TRFD	49.55	51.87	51.87	49.82	52.24	52.24	Balloons	Rep-TRFD	51.60	55.86	55.86	52.18	56.22	56.22
	DRO-TFF	46.60	49.16	49.16	47.38	50.67	50.67		DRO-TFF	48.53	52.36	52.36	49.14	54.78	54.78
	HLRTF	44.45	44.49	44.49	45.79	45.80	45.80		HLRTF	48.69	48.96	48.96	50.49	50.81	50.81
	TRLRF	32.51	NA	NA	33.48	NA	33.48		TRLRF	34.82	NA	NA	36.26	NA	36.26
Face	Rep-TRFD	47.94	48.54	48.54	48.64	49.17	49.17	Face	Rep-TRFD	51.94	53.45	53.45	52.65	54.18	54.18
	DRO-TFF	46.57	47.33	47.33	47.07	48.14	48.14		DRO-TFF	49.20	51.61	51.61	49.58	52.68	52.68
	HLRTF	43.90	43.85	43.90	45.11	45.04	45.11		HLRTF	47.93	48.14	48.14	49.29	49.45	49.45
	TRLRF	33.51	NA	NA	33.98	NA	33.98		TRLRF	38.48	NA	NA	39.42	NA	39.42
Flowers	Rep-TRFD	45.34	45.87	45.87	45.50	46.26	46.26	Flowers	Rep-TRFD	50.07	51.75	51.75	50.10	51.98	51.98
	DRO-TFF	44.47	45.81	45.81	44.65	46.17	46.17		DRO-TFF	47.74	51.48	51.48	47.91	51.77	51.77
	HLRTF	42.17	42.24	42.24	42.55	42.60	42.60		HLRTF	46.47	46.68	46.68	47.06	47.25	47.25
	TRLRF	28.34	NA	NA	28.61	NA	28.61		TRLRF	33.22	NA	NA	33.84	NA	33.84
Washington DC	Rep-TRFD	45.84	46.80	46.80	45.97	46.86	46.86	Washington DC	Rep-TRFD	48.37	49.51	49.51	48.39	49.64	49.64
	DRO-TFF	39.56	41.62	41.62	40.07	42.41	42.41		DRO-TFF	41.02	43.18	43.18	42.02	44.43	44.43
	HLRTF	41.67	41.69	41.69	42.02	42.04	42.04		HLRTF	46.53	46.67	46.67	46.88	47.00	47.00
	TRLRF	28.31	NA	NA	28.55	NA	28.55		TRLRF	30.08	NA	NA	30.45	NA	30.45
Botswana	Rep-TRFD	42.88	43.19	43.19	42.90	43.22	43.22	Botswana	Rep-TRFD	45.37	46.06	46.06	45.39	46.10	46.10
	DRO-TFF	37.75	38.71	38.71	38.46	39.62	39.62		DRO-TFF	39.39	40.27	40.27	40.23	41.34	41.34
	HLRTF	40.32	40.66	40.66	40.50	40.82	40.82		HLRTF	44.04	44.11	44.11	44.16	44.23	44.23
	TRLRF	28.37	NA	NA	28.69	NA	28.69		TRLRF	30.17	NA	NA	30.65	NA	30.65
Salinas	Rep-TRFD	51.56	53.29	53.29	51.73	53.44	53.44	Salinas	Rep-TRFD	52.85	55.61	55.61	52.88	55.77	55.77
	DRO-TFF	40.56	41.10	41.10	43.52	44.68	44.68		DRO-TFF	41.46	42.32	42.32	44.36	46.23	46.23
	HLRTF	45.81	45.81	45.81	46.50	46.54	46.54		HLRTF	50.04	50.01	50.01	50.97	51.02	51.02
	TRLRF	35.35	NA	NA	36.36	NA	36.36		TRLRF	35.02	NA	NA	36.82	NA	36.82