
ISUM: Inverse Problem Solver via Unbalanced Optimal Transport Map

Donggyu Lee*

Department of Mathematical Sciences
Seoul National University
dglee442@snu.ac.kr

Taekyung Lee*

Interdisciplinary Program in Artificial Intelligence
Seoul National University
dlxorud1231@snu.ac.kr

Jaewoong Choi†

Department of Statistics
Sungkyunkwan University
jaewoongchoi@skku.edu

Myungjoo Kang†

Department of Mathematical Sciences
Seoul National University
mkang@snu.ac.kr

Abstract

Solving unpaired inverse problems is critical in many applications where paired noisy measurements and clean targets are unavailable. We propose a novel approach for unpaired image inverse problems based on Unbalanced Optimal Transport (UOT). Our method formulates reconstruction as learning a transport map from noisy measurements to clean images, leveraging a likelihood-based cost function. By relaxing the strict marginal-matching constraint via UOT, our model enhances its robustness to multi-level observation noise. Moreover, under a quadratic cost function, our model effectively addresses linear inverse problems with unknown corruption operators. Our experiments demonstrate that our model achieves state-of-the-art performance on both linear (Gaussian deblurring and super-resolution) and nonlinear (high dynamic range reconstruction and nonlinear deblurring) image inverse problems.

1 Introduction

The inverse problem is defined as the task of reconstructing an unknown signal x from a possibly noisy measurement y [7, 26]. Many scientific and engineering tasks are formulated as inverse problems, including seismic imaging [20, 34], audio signal processing [22, 24], and medical imaging [6, 29]. In the computer vision domain, they are particularly relevant to various image reconstruction tasks, including image denoising [37], super-resolution [8, 30], and high dynamic range(HDR) reconstruction [36]. Formally, for a forward operator \mathcal{A} and a noisy measurement $\mathbf{y} \in \mathbb{R}^m$, the inverse problem can be formulated to find $\mathbf{x} \in \mathbb{R}^n$ from the following observation model:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \mathbf{n} \quad (1)$$

with measurement noise \mathbf{n} , which is typically assumed to be Gaussian $\mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}_m)$.

One of the main challenges in dealing with inverse problems is that the corresponding equations (Eq. 1) are often ill-posed and do not have a unique solution [9]. To address this issue, the Bayesian approach performs maximum a posteriori (MAP) estimation as follows:

$$\hat{\mathbf{x}}_{\text{MAP}}(\mathbf{y}) = \arg \max_{\mathbf{x}} p(\mathbf{x} \mid \mathbf{y}) = \arg \min_{\mathbf{x}} [-\log p(\mathbf{y} \mid \mathbf{x}) - \log p(\mathbf{x})]. \quad (2)$$

*Equal Contribution

†Correspondence, Equal Contribution

Intuitively, the MAP estimate addresses the ill-posedness of inverse problems by integrating prior knowledge of \mathbf{x} through $p(\mathbf{x})$ with the likelihood of the measurement. In this context, selecting a suitable prior $p(\mathbf{x})$ is essential to ensure accurate reconstruction. Recently, generative models such as GANs [2, 35], VAEs [14], Optimal Transport Map [19, 30], and Diffusion models [5, 29, 36] have demonstrated strong capabilities in modeling complex prior distributions.

We propose a novel approach, **ISUM**, which models an **Inverse problem Solver using Unbalanced Optimal Transport (UOT) Map**. We formulate the unpaired inverse problems as learning the transport map from the noisy measurement distribution to the target signal distribution using a cost function motivated by the MAP estimation principle. Furthermore, by leveraging the UOT framework, our model achieves robustness to multi-level observation noise through relaxation of the strict marginal constraint. The key contributions of our approach are as follows:

- We introduce the first model for unpaired inverse problems based on Unbalanced Optimal Transport, incorporating a likelihood-based cost function.
- Our model achieves state-of-the-art performance on various unpaired inverse problem benchmarks.
- By leveraging the UOT formulation, our model achieves robustness to multi-level observation noise, ensuring effectiveness in real-world settings.
- By adopting a quadratic cost, our model can address inverse problems with unknown corruption operators while retaining competitive performance.

2 Background

Notations and Assumptions Let \mathcal{X}, \mathcal{Y} be compact complete metric spaces with probability measures μ and ν , respectively, both absolutely continuous with respect to the Lebesgue measure. In the context of inverse problems, **μ and ν denote the distributions of noisy measurements and clean target signals**. For a measurable map T , $T_\# \mu$ denotes the pushforward of μ , and $\Pi(\mu, \nu)$ the set of couplings on $\mathcal{X} \times \mathcal{Y}$ with marginals μ and ν . For a function $f : \mathbb{R} \rightarrow [-\infty, \infty]$, we write its convex conjugate as $f^*(y) = \sup_{x \in \mathbb{R}} \{ \langle x, y \rangle - f(x) \}$.

2.1 Unbalanced Optimal Transport

Optimal Transport The optimal transport (OT) problem finds a mapping that transports a source distribution μ to a target distribution ν while minimizing a given cost function $c(\cdot, \cdot)$. Monge [25] first formulated this problem as finding a deterministic transport map T satisfying $T_\# \mu = \nu$.

$$C(\mu, \nu) := \inf_{T_\# \mu = \nu} \left[\int_{\mathcal{X}} c(\mathbf{x}, T(\mathbf{x})) d\mu(\mathbf{x}) \right]. \quad (3)$$

However, Monge’s OT problem (Eq. 3) is non-convex, and the existence of a solution T depends on the distributions μ and ν [33]. To overcome these limitations, Kantorovich proposes a formulation that relaxes the deterministic constraint on the transport map T [16]. Precisely, Kantorovich’s formulation finds a transport plan π for the OT problem as follows:

$$C_{ot}(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{X} \times \mathcal{Y}} c(\mathbf{x}, \mathbf{y}) d\pi(\mathbf{x}, \mathbf{y}) \right], \quad (4)$$

where $\Pi(\mu, \nu)$ denotes the set of joint distributions with marginals μ and ν . Unlike Monge’s OT problem, the Kantorovich formulation (Eq. 4) guarantees the existence of an optimal transport plan under mild assumptions. Moreover, when μ and ν are absolutely continuous with respect to the Lebesgue measure (our assumption), the solution of Monge’s problem T^* exists and the optimal transport plan π is characterized by this optimal transport map, i.e., $\pi^* = (Id \times T^*)_\# \mu$ [33].

Unbalanced Optimal Transport Note that in the standard OT problem (Eq. 4), the transport plan π is constrained to have marginal distributions μ and ν . The *Unbalanced Optimal Transport (UOT)* problem [3, 23] is a generalized formulation that relaxes this marginal constraint:

$$C_{uot}(\mu, \nu) = \inf_{\pi \in \mathcal{M}_+(\mathcal{Y} \times \mathcal{X})} \left[\int_{\mathcal{Y} \times \mathcal{X}} c(\mathbf{y}, \mathbf{x}) d\pi(\mathbf{y}, \mathbf{x}) + D_{\Psi_1}(\pi_0 | \mu) + D_{\Psi_2}(\pi_1 | \nu) \right] \quad (5)$$

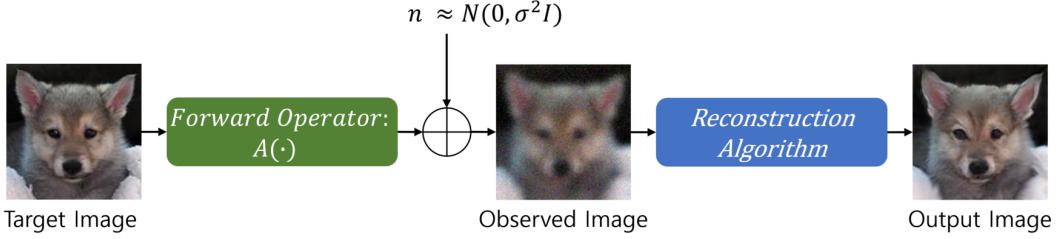


Figure 1: **Visualization of the inverse problem** (Eq. 1). Our model learns the reconstruction mapping from noisy observations to target images (*Output Image*).

where $\mathcal{M}_+(\mathcal{Y} \times \mathcal{X})$ denotes the set of positive Radon measures on $\mathcal{Y} \times \mathcal{X}$. The terms D_{Ψ_1} and D_{Ψ_2} are penalty terms on the marginal distributions defined as an f -divergences as follows:

$$D_{\Psi_i}(\pi_j | \eta) = \int \Psi_i \left(\frac{d\pi_j(x)}{d\eta(x)} \right) d\eta(x) \quad (6)$$

with convex functions Ψ_i . By introducing these penalty terms, the two marginal distributions of the optimal UOT plan π^* are softly matched to μ and ν , i.e., $\pi_0^* \approx \mu, \pi_1^* \approx \nu$. This relaxation enables UOT to achieve robustness against outliers [1, 27] and effectively address class imbalance between marginal distributions [10, 21].

3 Method

In this section, we present our UOT-based model for inverse problems, called **ISUM**. Our model is based on formulating the inverse problem solver as an Unbalanced Optimal Transport Map from the noisy measurements \mathbf{y} to the target signal \mathbf{x} (Eq. 1). In Sec 3.1, we introduce our optimal transport formulation of inverse problems. In Sec 3.2, we describe the learning objective, which is derived as a generalization of the UOTM framework [4] under the generalized cost function.

3.1 Optimal Transport Formulation of Unpaired Inverse Problems

Task Formulation Our main task is the **unpaired image inverse problems with a known corruption operator \mathcal{A} and unknown noise level $\sigma_y > 0$ (Eq. 1)**. (The case of an unknown corruption operator will be discussed in Sec 4.3.) Formally, let

$$Y = \{\mathbf{y}_i : \mathbf{y}_i \in \mathcal{Y}, i = 1, \dots, M\} \quad \text{and} \quad X = \{\mathbf{x}_j : \mathbf{x}_j \in \mathcal{X}, j = 1, \dots, N\}. \quad (7)$$

denote the sets of *noisy measurements* and *target signals (clean image)*, respectively (Fig. 1). We assume the unpaired setting, i.e., Y and X are independently sampled from the source distribution μ and the target distribution ν . Then our goal is to learn an inverse problem solver T :

$$T : \mathcal{Y} \rightarrow \mathcal{X}, \quad \mathbf{y} \mapsto T(\mathbf{y}) \quad (8)$$

from these unpaired training data. Following the MAP approach (Eq. 2), T should satisfy the following conditions:

- (i) The estimate $T(\mathbf{y})$ follows the target signal distribution, i.e., $T_{\#}\mu = \nu$.
- (ii) For each given \mathbf{y}_0 , $T(\mathbf{y}_0)$ maximizes the log-posterior $\log p(\cdot | \mathbf{y}_0)$.

Our main observation is that these two conditions can be naturally interpreted through the (Unbalanced) Optimal Transport (Eq. 3 and 5). Formally, in the standard OT problem, the Optimal Transport Map (OT Map) T^* is defined as the transport cost minimizer over valid transport maps:

- (a) Valid transport maps satisfying $T_{\#}^*\mu = \nu$.
- (b) Each \mathbf{y} is mapped to minimize the transport cost $\int_{\mathcal{Y}} c(\mathbf{y}, T^*(\mathbf{y})) d\mu(\mathbf{y})$.

Therefore, the OT Map T^* naturally satisfies the first condition (i) of the inverse problem solver. **Our approach is to design an appropriate cost function $c(\cdot, \cdot)$ so that the second condition (ii) is also satisfied.**

Likelihood Cost and MAP Estimate Our goal is to solve inverse problems using the optimal transport (OT) framework. From a MAP estimation perspective, one can consider finding the transport plan that minimizes the negative log-posterior $-\log p(\mathbf{x} \mid \mathbf{y})$. This can be expressed as follows in the standard OT framework:

$$C_{ot}(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{Y} \times \mathcal{X}} -\log p(\mathbf{x} \mid \mathbf{y}) d\pi(\mathbf{y}, \mathbf{x}) \right]. \quad (9)$$

$$= \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{Y} \times \mathcal{X}} -\log p(\mathbf{y} \mid \mathbf{x}) - \log p(\mathbf{x}) + \log p(\mathbf{y}) d\pi(\mathbf{y}, \mathbf{x}) \right]. \quad (10)$$

$$= \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{Y} \times \mathcal{X}} -\log p(\mathbf{y} \mid \mathbf{x}) d\pi(\mathbf{y}, \mathbf{x}) \right]. \quad (11)$$

Note that the equality (11) holds because of the marginal constraints on the transport plan $\pi \in \Pi(\mu, \nu)$. Therefore, under the assumption of Gaussian noise, the problem can be formulated as follows:

$$C_{ot}(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{Y} \times \mathcal{X}} \|\mathbf{y} - \mathcal{A}(\mathbf{x})\|_2^2 d\pi(\mathbf{y}, \mathbf{x}) \right]. \quad (12)$$

and based on this formulation, we introduce a *likelihood-based cost function* $c_l(\mathbf{y}, \mathbf{x}) = \|\mathbf{y} - \mathcal{A}(\mathbf{x})\|_2^2$. Furthermore, we incorporate the standard quadratic cost function, and consequently propose the following overall cost:

$$c(\mathbf{y}, \mathbf{x}) = \tau (c_l(\mathbf{y}, \mathbf{x}) + c_q(\mathbf{y}, \mathbf{x})) \quad \text{where } c_l(\mathbf{y}, \mathbf{x}) = \|\mathcal{A}(\mathbf{x}) - \mathbf{y}\|_2^2 \text{ and } c_q(\mathbf{y}, \mathbf{x}) = \|\mathbf{y} - \mathbf{x}\|_2^2 \quad (13)$$

Under a weak assumption, this additional quadratic term can allow us to guarantee the existence of the OT Map T^* (For more details, see Appendix B). Furthermore, when the corruption operator is unknown, our model can be extended by relying solely on the quadratic cost (Section 4.3).

3.2 Proposed Model

The Unbalanced Optimal Transport (Eq. 5) is a generalization of the standard Optimal Transport by relaxing the marginal distribution constraints. Such flexibility is particularly advantageous in our unpaired inverse problem setting, where exact marginal matching constraint $T_{\#}\mu = \nu$ as in (a) can be too restrictive: UOT mitigates this issue by allowing marginal discrepancies that lead to greater adaptability in matching target distributions (Eq. 5). Moreover, the flexibility provides the capability to address class imbalance between the source and target distributions [10, 21], which is beneficial in cases involving multiple measurement noise levels (Sec 4.2). In such cases, several noisy measurements \mathbf{y} may correspond to the same target signal \mathbf{x} . In this regard, this scenario can be interpreted as one instance of class imbalance. Leveraging these benefits, we design our model based on the UOT Map.

We estimate the UOT Map T_u^* via the UOTM framework [4], which utilizes the semi-dual formulation of Unbalanced Optimal Transport [32] (Eq. 14).

$$C_{uot}(\mu, \nu) = \sup_{v \in \mathcal{C}} \left[\int_{\mathcal{X}} -\Psi_1^*(-v^c(\mathbf{y})) d\mu(\mathbf{y}) + \int_{\mathcal{Y}} -\Psi_2^*(-v(\mathbf{x})) d\nu(\mathbf{x}) \right], \quad (14)$$

where v is the potential function and $v^c(\mathbf{y}) = \inf_{\mathbf{x} \in \mathcal{X}} (c(\mathbf{y}, \mathbf{x}) - v(\mathbf{x}))$ with the cost function defined in Eq. 13. Then, following the approach in [11] and [18], we approximate the UOT Map T_u^* with a parametrized T_θ as follows:

$$T_\theta(\mathbf{y}) \in \operatorname{arginf}_{\mathbf{x} \in \mathcal{X}} [c(\mathbf{y}, \mathbf{x}) - v(\mathbf{x})] \Leftrightarrow v^c(\mathbf{y}) = c(\mathbf{y}, T_\theta(\mathbf{y})) - v(T_\theta(\mathbf{y})). \quad (15)$$

Consequently, replacing v^c in Eq. 14 with Eq. 15 and parametrizing the potential function as a neural network v_ϕ , we define the learning objective as follows:

$$\mathcal{L}_{v_\phi, T_\theta} = \inf_{v_\phi} \int_{\mathcal{X}} \Psi_1^* \left(-\inf_{T_\theta} [c(\mathbf{y}, T_\theta(\mathbf{y})) - v_\phi(T_\theta(\mathbf{y}))] \right) d\mu(\mathbf{y}) + \int_{\mathcal{Y}} \Psi_2^* (-v_\phi(\mathbf{x})) d\nu(\mathbf{x}). \quad (16)$$

where Ψ_1^* and Ψ_2^* are monotone increasing convex functions determined by the marginal distribution penalties in the UOT problem (Eq. 5). Note that when Ψ_1 and Ψ_2 are the convex indicator function of the set $\{1\}$, it follows that $\Psi_1^*(\cdot) = \Psi_2^*(\cdot) = \text{Identity}(\cdot)$, which indicates that the UOT map can be regarded as a generalization of the OT map.

Table 1: **Quantitative results on various inverse problems.** We conduct experiments on the following four inverse problems: Gaussian deblurring, super-resolution, high dynamic range reconstruction, and nonlinear deblurring. The **boldface** and underlined values indicate the best and second-best performance. Our model outperforms existing approaches, with minor exceptions on AFHQ super-resolution (second-best PSNR) and HDR reconstruction (second-best FID).

(a) Gaussian deblurring								
Method	FFHQ				AFHQ			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	20.11	0.6035	0.209	52.901	19.99	0.5472	0.273	58.927
OTUR [35]	<u>23.82</u>	<u>0.7106</u>	<u>0.136</u>	<u>24.337</u>	<u>23.91</u>	<u>0.6777</u>	<u>0.165</u>	<u>30.773</u>
RCOT [30]	22.07	0.5492	0.279	123.692	22.34	0.5365	0.268	132.465
Ours	24.06	0.7139	0.124	21.210	24.22	0.6804	0.139	12.566
(b) Super-resolution 4×								
Method	FFHQ				AFHQ			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	20.13	0.6257	0.209	50.066	20.14	0.5833	0.261	44.252
OTUR [35]	<u>24.09</u>	<u>0.7243</u>	<u>0.129</u>	<u>22.751</u>	24.71	0.7079	<u>0.148</u>	<u>19.575</u>
RCOT [30]	24.05	0.6820	0.260	118.776	25.04	0.7137	0.208	69.072
Ours	24.35	0.7371	0.118	19.475	<u>24.97</u>	0.7142	0.133	15.939
(c) High Dynamic Range reconstruction								
Method	FFHQ				AFHQ			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	21.24	0.7978	0.138	25.842	23.36	0.8179	0.127	10.528
OTUR [35]	<u>25.32</u>	<u>0.8545</u>	<u>0.076</u>	16.458	<u>26.25</u>	<u>0.8542</u>	<u>0.091</u>	7.227
RCOT [30]	19.26	0.6755	0.133	33.422	18.99	0.7060	0.166	27.767
Ours	26.02	0.8642	0.064	<u>20.840</u>	26.40	0.8653	0.074	<u>7.897</u>
(d) Nonlinear Deblurring								
Method	FFHQ				AFHQ			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	21.37	0.7373	0.157	43.661	23.03	0.7271	0.158	17.377
OTUR [35]	<u>26.94</u>	<u>0.8594</u>	<u>0.068</u>	<u>12.538</u>	<u>26.09</u>	<u>0.8253</u>	<u>0.092</u>	<u>7.651</u>
RCOT [30]	25.14	0.7221	0.151	52.268	24.48	0.7172	0.153	29.902
Ours	28.52	0.8841	0.051	11.370	27.74	0.8589	0.069	5.113

4 Experiments

Datasets and Metrics For our experiments, we use the AFHQ and FFHQ datasets as target distributions. We generate the source distributions by applying a corruption operator to each image dataset. For the unpaired setting, we randomly shuffle target \mathbf{x} and source \mathbf{y} to remove pairing information. For quantitative evaluation, we measure image restoration quality in terms of both pixel-level fidelity and perceptual image quality, using metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Perceptual Image Patch Similarity (LPIPS) [38] and Fréchet Inception Distance (FID) [15]. Note that we compute FID using both the training and test sets. We adopt the experimental setting of [17, 28] for super-resolution, while for the other tasks we use the setting from [36]. For super-resolution, due to the dimensional mismatch between \mathbf{y} and \mathbf{x} , we apply bicubic interpolation Q to \mathbf{y} and define the modified quadratic cost as $c_q(\mathbf{y}, \mathbf{x}) = \|Q(\mathbf{y}) - \mathbf{x}\|_2^2$. In all experiments, we assume additive white Gaussian noise with $\sigma_y = 0.05$ as the measurement noise. Further implementation details are provided in Appendix A.

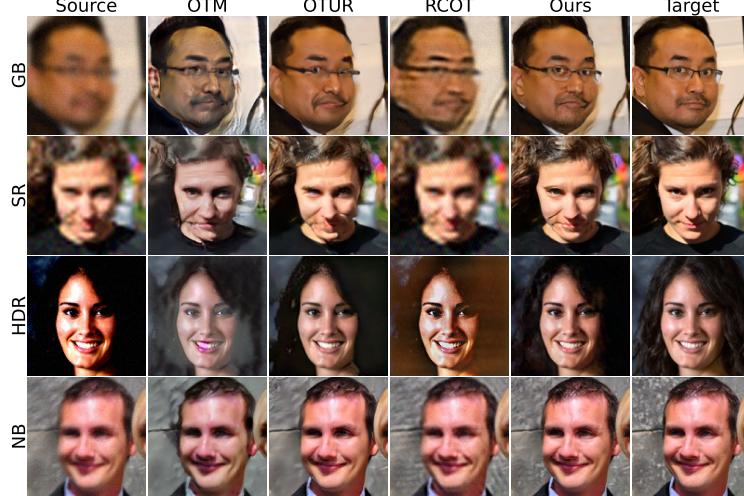


Figure 2: **Comparison of inverse problem solvers on FFHQ** for four tasks: Gaussian deblurring (GB), Super-resolution (SR), High dynamic range reconstruction (HDR), and Nonlinear deblurring (NB). Our model produces higher fidelity images with well-preserved textures.

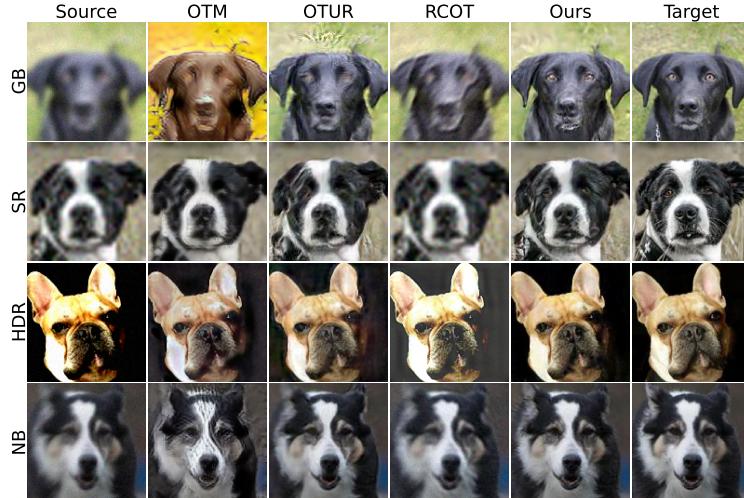


Figure 3: **Comparison of inverse problem solvers on AFHQ** for four tasks: Gaussian deblurring (GB), Super-resolution (SR), High dynamic range reconstruction (HDR), and Nonlinear deblurring (NB). Our model produces higher fidelity images with well-preserved textures.

4.1 Solving inverse problem

We evaluate our model on four inverse problem tasks: Gaussian deblurring, super-resolution, high dynamic range reconstruction, and nonlinear deblurring. We benchmark our model against prior approaches for unpaired inverse problems, based on GANs (*OTUR* [35]) and OT Maps (*NOT* [19] and *RCOT* [30]). Table 1 presents the results for each task. Overall, our model consistently surpasses existing methods, with the only exceptions being a comparable second-best PSNR score on AFHQ in the super-resolution task and the second-best FID score in HDR reconstruction. These results suggest that our method attains higher fidelity in recovering target signals relative to other approaches.

Figures 2 and 3 qualitatively illustrate the superior image fidelity of our model. NOT does not tend to overemphasize features such as contours and textures, while RCOT often fails to reliably remove degradations. Although RCOT achieves a higher PSNR on AFHQ, its reconstructions appear noticeably inferior under qualitative inspection. OTUR mitigates degradations effectively but often leads to distortion or excessive smoothing of fine details. Also, in HDR restoration task, OTUR

Table 2: **Quantitative results under multi-level observation noise** for four inverse problems on FFHQ. Our model exhibits superior robustness across noise levels.

(a) Linear inverse problems: Gaussian deblurring and Super-resolution (4×).

Method	Gaussian Deblurring				Super Resolution 4×			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	19.07	0.5491	0.250	98.558	18.91	0.5653	0.264	99.053
OTUR [35]	22.55	0.6417	0.174	67.323	23.20	0.6681	0.178	70.223
Ours (OT-based)	22.87	0.6562	0.160	91.309	23.21	0.6732	0.170	85.674
Ours	23.04	0.6649	0.154	65.664	23.30	0.6864	0.157	58.406

(b) Nonlinear inverse problems: High Dynamic Range and Nonlinear Deblurring.

Method	High Dynamic Range				Nonlinear Deblurring			
	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
NOT [19]	21.01	0.7728	0.153	62.641	21.50	0.7201	0.177	76.795
OTUR [35]	24.29	0.8098	0.103	52.149	25.23	0.7763	0.114	61.244
Ours (OT-based)	25.92	0.8488	0.086	49.750	26.90	0.8307	0.087	51.805
Ours	25.44	0.8330	0.084	51.300	27.25	0.8427	0.077	43.450

produces artifacts around the dog’s head boundary and insufficient recovery of hair texture. In contrast, our model generates clean and sharp reconstructions with textures preserved in high fidelity. Additional qualitative results for each task are provided in Appendix C.2.

4.2 Multi noise level observation

Our method employs the UOT Map, which, unlike the standard OT Map, provides additional flexibility for each sample $T(y)$ (Sec. 3.2). While this flexibility is often associated with handling class imbalance [10, 13, 21], in our formulation, it also enables robustness to **multi-level observation noise**. Most unpaired inverse problem solvers [30, 35] assume a single fixed noise level; however, in practice, the noise level can vary significantly across images. To assess robustness under such multi-level noise conditions, we designed an experiment where noise is drawn from a mixture of four levels: for each level, we construct the source distribution by randomly and independently sampling degraded images in a 4:3:2:1 ratio. For the evaluation phase, we apply all four types of noise to 500 images, generating 2,000 degraded samples in total. We compare our method against NOT, OTUR, and the OT-variant of ours (Ours-OT), in which the two convex conjugate functions in our framework are replaced with the identity function. These constitute the OT Map-based baselines, whereas ours relies on the UOT Map.

As shown in Table 2, our model outperforms the other OT-based methods on both linear and nonlinear inverse problems. The Figure 4 illustrates the impact of increasing noise levels. For clarity, note that each column presents reconstructed samples from a single model for each method, evaluated across multiple input measurements y with varying noise levels σ_y . In the multi-noise setting, NOT fails to reliably restore images, and OTUR struggles to reconstruct under high noise. Similarly, Ours-OT also shows image quality degradation as noise intensity increases. In contrast, our model successfully restores images and maintains relatively low variance even under strong noise.

4.3 Likelihood cost ablation

We conduct an ablation study on our proposed cost function, which consists of two terms: the problem-agnostic quadratic cost c_q and the problem-adaptive likelihood-based cost c_l . To assess the contribution of each term, we analyze the performance when either component is removed. The results are summarized in Fig. 5 (See Appendix C.1 for metric table). We conduct experiments on two inverse problems, Gaussian deblurring and super-resolution. In most cases, the full cost (with both terms) achieves the best performance. Notably, when relying solely on c_q , our method still outperforms state-of-the-art baselines such as OTUR. This clearly demonstrates that our model remains effective for inverse problems even under unknown corruption operators.



Figure 4: **Comparison of inverse problem solvers under multi-level observation noise** on FFHQ for the nonlinear deblurring (see Appendix C.2 for additional tasks). Our method demonstrates stable reconstruction quality with substantially smaller variations across noise levels.

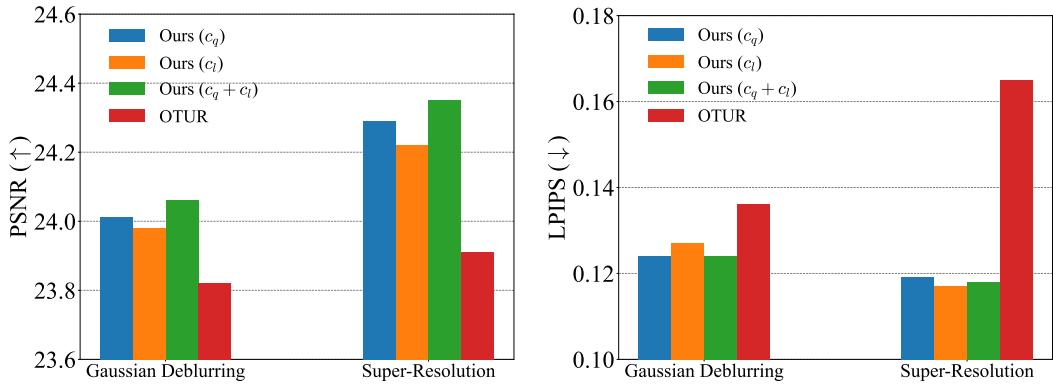


Figure 5: **Ablation study on the cost function $c(y, x)$** . The combination of the likelihood cost c_l and the quadratic cost c_q yields the best overall performance. Notably, even the blind-corruption version of our method (using only c_q) surpasses the previous state-of-the-art OTUR model.

5 Conclusion

We propose ISUM, a UOT-based method for solving inverse problems. Our approach formulates inverse problems within the (Unbalanced) Optimal Transport framework and introduces a cost function with two terms: a quadratic cost and a likelihood-based cost. Experiments show that ISUM achieves outperforming performance compared to existing methods. Also, our model demonstrate both the superior robustness to multi-level observation noise and the effectiveness of the proposed likelihood-based cost function. Moreover, even with unknown corruption operators, our framework achieves competitive performance using only quadratic cost, demonstrating its broad applicability to inverse problems with unknown corruption operators. A limitation of our work is that the method requires training from scratch. Expanding our approach to training-free or fine-tuning settings will be investigated in our future work.

Acknowledgements

This work was supported by the NRF grant [RS-2024-00421203, RS-2024-00406127] and the IITP grant funded by the Korea government (MSIT) [RS2021-II211343-GSAI, Artificial Intelligence Graduate School Program (Seoul National University)]. Jaewoong was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) [RS-2024-00349646].

References

- [1] Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. *Advances in Neural Information Processing Systems*, 33:12934–12944, 2020.
- [2] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. In *International conference on machine learning*, pages 537–546. PMLR, 2017.
- [3] Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11):3090–3123, 2018.
- [4] Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Generative modeling through the semi-dual formulation of unbalanced optimal transport. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [5] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.
- [6] Hyungjin Chung, Dohoon Ryu, Michael T McCann, Marc L Klasky, and Jong Chul Ye. Solving 3d inverse problems using pre-trained 2d diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22542–22551, 2023.
- [7] Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024.
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014.
- [9] Heinz W Engl and Ronny Ramlau. Regularization of inverse problems. In *Encyclopedia of applied and computational mathematics*, pages 1233–1241. Springer, 2015.
- [10] Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and Fabian J Theis. Unbalancedness in neural monge maps improves unpaired domain translation. In *The Twelfth International Conference on Learning Representations*, 2024.
- [11] Jiaojiao Fan, Shu Liu, Shaojun Ma, Yongxin Chen, and Hao-Min Zhou. Scalable computation of monge maps with general costs. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.
- [12] Albert Fathi and Alessio Figalli. Optimal transportation on non-compact manifolds. *Israel Journal of Mathematics*, 175(1):1–59, 2010.
- [13] Milena Gazdieva, Jaemoo Choi, Alexander Kolesov, Jaewoong Choi, Petr Mokrov, and Alexander Korotin. Robust barycenter estimation using semi-unbalanced neural optimal transport. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [14] Hwan Goh, Sheroze Sheriffdeen, Jonathan Wittmer, and Tan Bui-Thanh. Solving bayesian inverse problems via variational autoencoders. *arXiv preprint arXiv:1912.04212*, 2019.

[15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

[16] Leonid Vitalevich Kantorovich. On a problem of monge. *Uspekhi Mat. Nauk*, pages 225–226, 1948.

[17] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in neural information processing systems*, 35:23593–23606, 2022.

[18] Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, Alexander Filippov, and Evgeny Burnaev. Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark. *Advances in neural information processing systems*, 34:14593–14605, 2021.

[19] Alexander Korotin, Daniil Selikhanovych, and Evgeny Burnaev. Neural optimal transport. In *The Eleventh International Conference on Learning Representations*, 2023.

[20] Patrick Lailly. As a sequence of before stack migrations. In *Conference on inverse scattering–Theory and application*, volume 11, page 206, 1983.

[21] Taekyung Lee, Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Unpaired point cloud completion via unbalanced optimal transport. In *Forty-second International Conference on Machine Learning*, 2025.

[22] Jean-Marie Lemercier, Julius Richter, Simon Welker, Eloi Moliner, Vesa Välimäki, and Timo Gerkmann. Diffusion models for audio restoration. *arXiv preprint arXiv:2402.09821*, 2024.

[23] Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures. *Inventiones mathematicae*, 211(3):969–1117, 2018.

[24] Eloi Moliner, Jaakko Lehtinen, and Vesa Välimäki. Solving audio inverse problems with a diffusion model. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023.

[25] Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704, 1781.

[26] Adnan Qayyum, Inaam Ilahi, Fahad Shamshad, Farid Boussaid, Mohammed Bennamoun, and Junaid Qadir. Untrained neural network priors for inverse imaging problems: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):6511–6536, 2022.

[27] Thibault Séjourné, Gabriel Peyré, and François-Xavier Vialard. Unbalanced optimal transport, from theory to numerics. *arXiv preprint arXiv:2211.08775*, 2022.

[28] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *ICLR*, 2023.

[29] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021.

[30] Xiaole Tang, Xin Hu, Xiang Gu, and Jian Sun. Residual-conditioned optimal transport: Towards structure-preserving unpaired and paired image restoration. volume 235 of *Proceedings of Machine Learning Research*, pages 47757–47777. PMLR, 21–27 Jul 2024.

[31] Phong Tran, Anh Tuan Tran, Quynh Phung, and Minh Hoai. Explore image deblurring via encoded blur kernel space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11956–11965, 2021.

[32] Adrien Vacher and François-Xavier Vialard. Semi-dual unbalanced quadratic optimal transport: fast statistical rates and convergent algorithm. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 34734–34758. PMLR, 23–29 Jul 2023.

[33] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

- [34] Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6):WCC1–WCC26, 2009.
- [35] Wei Wang, Fei Wen, Zeyu Yan, and Peilin Liu. Optimal transport for unsupervised denoising learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):2104–2118, 2022.
- [36] Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 20895–20905, 2025.
- [37] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.

A Implementation details

All models introduced in this paper are trained for 60,000 iterations, and we report the best results with respect to LPIPS and FID, even if they occur at intermediate iterations.

Ours In our model, unless otherwise specified, the settings follow those of UOTM [4] on CelebA-256. Our framework jointly learns the potential v_ϕ and the Optimal Transport Map T_θ . For the potential v_ϕ , we adopt the potential architecture from UOTM [4], while for OT Map T_θ , we employ the generator architecture from OTUR [35]. The learning rate for the potential is $lr_{v_\phi} = 5.0 \times 10^{-5}$, and the learning rate for the OT Map is $lr_{T_\theta} = 1.0 \times 10^{-4}$. The cost intensity hyperparameter τ is fixed to 0.001. The batch size during training is fixed at 32. The convex conjugate Ψ_i^* is derived from the generator function of the KL divergence, $\Psi(x) = \begin{cases} x \log x - x + 1, & \text{if } x > 0 \\ \infty, & \text{if } x \leq 0 \end{cases}$, which defines the associated f -divergence and yields the explicit form $\Psi_i^*(t) = e^t - 1$. In the case of Ours-OT, we instead set $\Psi_i^*(\cdot) = \text{Identity}(\cdot)$, while keeping all other configurations identical to those of Ours.

Baselines For NOT [19], we employ the generator and discriminator of UOTM and adopt its hyperparameter settings. For OTUR [35] and RCOT [30], we strictly follow all configurations as proposed in their original models.

Corruption operators For super-resolution, we adopt the setting of [17, 28] for super-resolution. For other tasks, we use the setting from [36]. Precisely,

- **Gaussian deblurring** (kernel size 61×61 and kernel standard deviation 3.0):

$$\mathbf{y} = \mathbf{k} * \mathbf{x} + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}_m) \quad (17)$$

where k is the Gaussian kernel and $*$ denotes the convolution operator.

- **Super-resolution** (4×4 patch downsampling):

$$\mathbf{y} = \mathbf{x} \downarrow_4 + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}_m) \quad (18)$$

- **High Dynamic Range (HDR) reconstruction** (scale factor 2.0, clipping to $[-1, 1]$):

$$\mathbf{y} = \text{clip}(2\mathbf{x}, -1, 1) + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}_m) \quad (19)$$

- **Nonlinear deblurring** (\mathcal{A} : neural operator from [31]):

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}_m) \quad \text{for pretrained operator } \mathcal{A} \quad (20)$$

B Twist condition

In this section, we discuss conditions ensuring the existence of the OT map. First, we introduce the condition referred to as the *twist condition*.

Definition B.1 (Left twist condition). *Let M and N be a n -dimensional manifold and N be a Polish space. Let $c : M \times N \rightarrow \mathbb{R}$ be a cost function and μ and ν be two probability measures on M and N , respectively. For a given cost function $c(\mathbf{y}, \mathbf{x})$, we define the skew left Legendre transform as the partial map*

$$\begin{aligned} \Lambda_c^l : M \times N &\rightarrow T^* M \\ \Lambda_c^l(\mathbf{y}, \mathbf{x}) &= \left(\mathbf{y}, \frac{\partial c}{\partial \mathbf{y}}(\mathbf{y}, \mathbf{x}) \right) \end{aligned}$$

whose domain of definition is

$$\mathcal{D}(\Lambda_c^l) = \left\{ (\mathbf{y}, \mathbf{x}) \in M \times N : \frac{\partial c}{\partial \mathbf{y}}(\mathbf{y}, \mathbf{x}) \text{ exists} \right\}.$$

We say that c satisfies the left twist condition if Λ_c^l is injective on $\mathcal{D}(\Lambda_c^l)$.

Now, we state a theorem that guarantees the existence and uniqueness of the optimal transport map.

Theorem B.1 ([12]). *Let M be a smooth (second countable) manifold, and N be a Polish space, and consider μ and ν (Borel) probability measures on M and N respectively. Assume that the cost $c : M \times N \rightarrow \mathbb{R}$ is lower semicontinuous and bounded from below. Moreover, assume that the following conditions hold:*

1. *the family of maps $\mathbf{y} \mapsto c(\mathbf{y}, \mathbf{x}) = c_{\mathbf{x}}(\mathbf{y})$ is locally semi-concave in \mathbf{y} locally uniformly in \mathbf{x} ,*
2. *the cost c satisfies the left twist condition,*
3. *the measure μ gives zero mass to countably $(n - 1)$ -Lipschitz sets,*
4. *the infimum in the Kantorovich problem $C(\mu, \nu) = \arg \min_{\gamma \in \Pi} \{ \int c(\mathbf{y}, \mathbf{x}) d\gamma \}$ is finite.*

Then there exists a borel map $T : M \rightarrow N$, which is an optimal transport map from μ to ν for the cost c . Moreover, the map T is unique μ -a.e., and any plan $\gamma_c \in \Pi(\mu, \nu)$ optimal for the cost c is concentrated on the graph of T .

Proof. See [12]. □

Remark. Denote $c(\mathbf{y}, \mathbf{x}; \lambda) = c_l(\mathbf{y}, \mathbf{x}) + \lambda c_q(\mathbf{y}, \mathbf{x})$ for $\lambda \geq 0$. Note that under our mild assumption in Section 2, all conditions of the above theorem except condition 2 (left twist condition) are satisfied with cost function $c(\mathbf{y}, \mathbf{x}; \lambda)$. To check the left twist condition, it is enough to show the injectivity of $\frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; \lambda)$ with respect to \mathbf{x} . However, for $\lambda = 0$ (i.e., when only the likelihood-based cost term is used), the map $\mathbf{x} \mapsto \frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; 0) = 2(\mathbf{y} - \mathcal{A}(\mathbf{x}))$ is not injective due to the ill-posedness of \mathcal{A} .

To handle this issue, we incorporate the standard quadratic cost c_q . Note that in many tasks such as Gaussian deblurring or HDR reconstruction, one can readily verify that $\frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; \lambda) = 2(\mathbf{y} - \mathcal{A}(\mathbf{x})) + 2\lambda(\mathbf{y} - \mathbf{x})$ is injective with respect to \mathbf{x} with $\lambda = 1$. In general, when \mathcal{A} is Lipschitz continuous, we can choose λ such that the function $\mathbf{x} \mapsto \frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; \lambda)$ is injective.

Claim 1. Assume that \mathcal{A} is Lipschitz continuous. Then for some $c > 0$, the function $\mathbf{x} \mapsto \frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; \lambda)$ is injective for all $\lambda > c$.

Proof. Let c be the Lipschitz constant of the operator \mathcal{A} . Then the following equation holds:

$$\frac{\partial}{\partial \mathbf{y}} c(\mathbf{y}, \mathbf{x}; \lambda) = 2(\mathbf{y} - \mathcal{A}(\mathbf{x})) + 2\lambda(\mathbf{y} - \mathbf{x}) = (2 + 2\lambda)\mathbf{y} - 2(\lambda\mathbf{x} + \mathcal{A}(\mathbf{x})) \quad (21)$$

Thus it is enough to show that $\lambda\mathbf{x} + \mathcal{A}(\mathbf{x})$ is injective. Also, for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$, $\lambda\mathbf{x}_1 + \mathcal{A}(\mathbf{x}_1) = \lambda\mathbf{x}_2 + \mathcal{A}(\mathbf{x}_2)$ implies that $\lambda\|\mathbf{x}_1 - \mathbf{x}_2\| = \|\mathcal{A}(\mathbf{x}_1) - \mathcal{A}(\mathbf{x}_2)\| \leq c\|\mathbf{x}_1 - \mathbf{x}_2\|$. Thus letting $\lambda > c$, the above result implies that $\mathbf{x}_1 = \mathbf{x}_2$ and $\lambda\mathbf{x} + \mathcal{A}(\mathbf{x})$ is injective. □

C Addition results

C.1 Additional Quantitative Results

Table 3: Ablation study on the cost function $c(\mathbf{y}, \mathbf{x})$, investigated on the FFHQ dataset

	cost term		Gaussian Deblurring				Super Resolution 4×			
	Quad term	IP term	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
Ours	✓		24.01	0.7130	0.124	21.516	24.29	0.7332	0.119	20.131
		✓	23.98	0.7140	0.1270	25.608	24.22	0.7354	0.117	19.085
	✓	✓	24.06	0.7139	0.124	21.210	24.35	0.7371	0.118	19.475
OTUR [35]			23.82	0.7106	0.136	24.337	23.91	0.6777	0.165	30.773

C.2 Additional Qualitative Results



Figure 6: Additional qualitative results of our model for the Gaussian deblurring task on FFHQ (degraded (Left) → clean (Right)).



Figure 7: Additional qualitative results of our model for the Gaussian deblurring task on AFHQ (degraded (Left) → clean (Right)).



Figure 8: Additional qualitative results of our model for the super-resolution task on FFHQ (degraded (Left) → clean (Right)).



Figure 9: Additional qualitative results of our model for the super-resolution task on AFHQ (degraded (Left) → clean (Right)).



Figure 10: Additional qualitative results of our model for the HDR reconstruction task on FFHQ (degraded (Left) → clean (Right)).



Figure 11: Additional qualitative results of our model for the HDR reconstruction task on AFHQ (degraded (Left) → clean (Right)).



Figure 12: Additional qualitative results of our model for the nonlinear deblurring task on FFHQ (degraded (Left) → clean (Right)).



Figure 13: Additional qualitative results of our model for the nonlinear deblurring task on AFHQ (degraded (Left) \rightarrow clean (Right)).

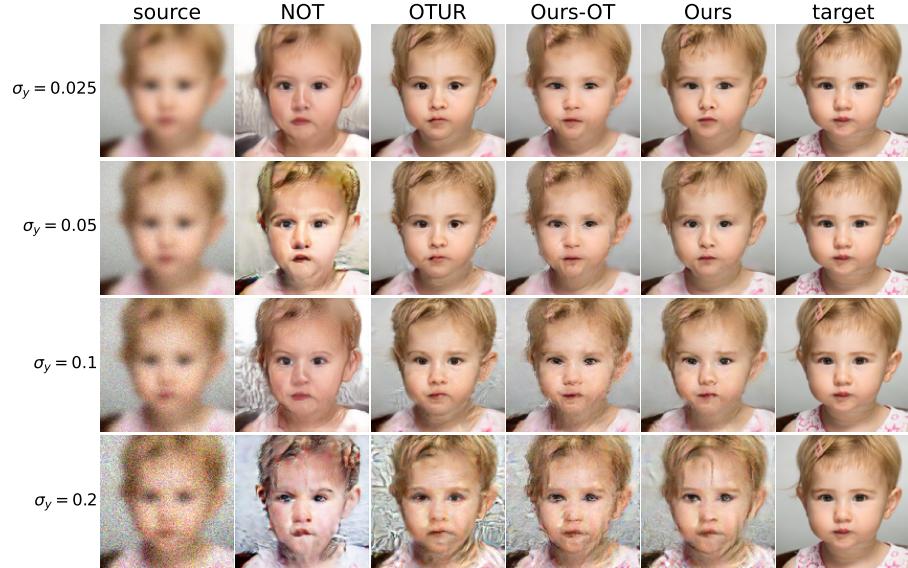


Figure 14: **Comparison of inverse problem solvers under multi-level observation noise on FFHQ for the Gaussian deblurring).**

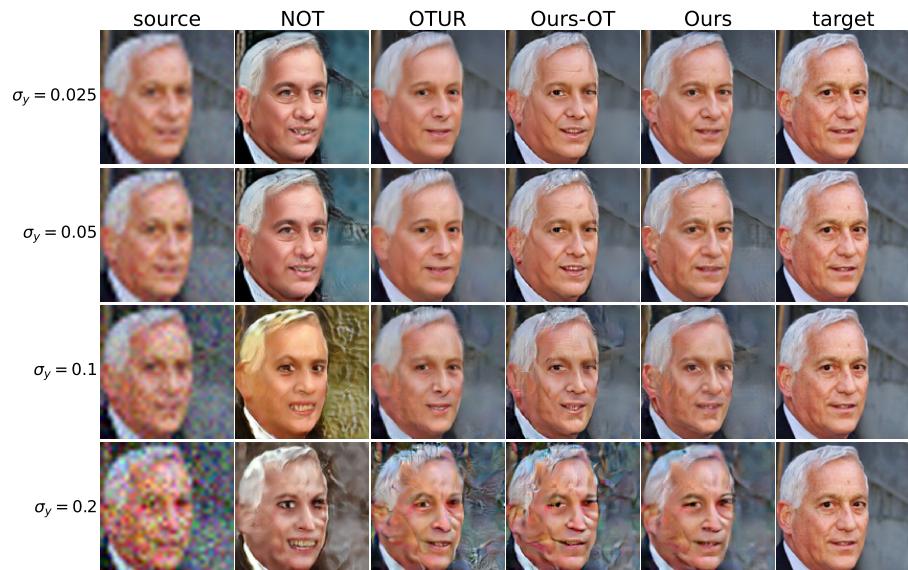


Figure 15: **Comparison of inverse problem solvers under multi-level observation noise on FFHQ for the super-resolution).**

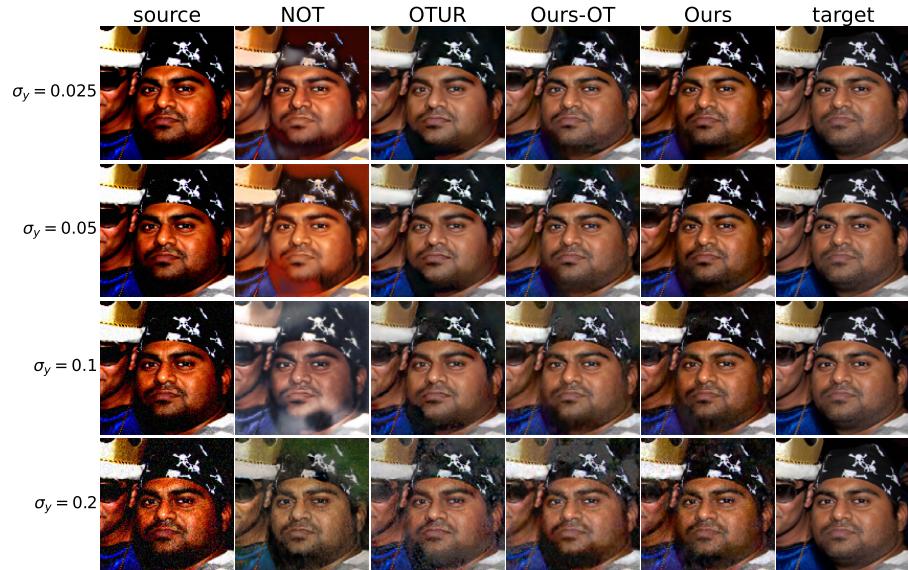


Figure 16: **Comparison of inverse problem solvers under multi-level observation noise on FFHQ for the HDR reconstruction).**



Figure 17: **Comparison of inverse problem solvers under multi-level observation noise on FFHQ for the nonlinear-deblurring).**