Supplementary Materials for Degradation-aware Dynamic Schrödinger Bridge for Unpaired Image Restoration

Anonymous Author(s)
Affiliation
Address
email

Contents

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

A Theoretical Analysis

A.1 Tighter Generalization Error Bound

In this subsection, we provide a step-by-step deduction of the generalization error bounds for the proposed *Degradation-aware Dynamic Schrödinger Bridge (DDSB)*, and also demonstrate that it has a tighter generalization error bound than the Schrödinger Bridge (SB) baseline.

Definition 1. Generalization Error. *The generalization error $\mathcal{E}_{gen}$ quantifies the discrepancy between the predicted clean image $\hat{x}_c$ and the true clean image x_c, defined as:*

$$\mathcal{E}_{\text{gen}} = \mathbb{E}\left[\|\hat{x}_c - x_c\|^2\right], \tag{1}$$

where $\hat{x}_c$ denotes the clean image generated by the model, and x_c denotes the true clean image.

Lemma 1. Optimal Transport Formulation. *The problem of learning the optimal transport plan $\pi^\star$ between the degraded image distribution π_d and the clean image distribution π_c is formulated as:*

$$\pi^\star = \underset{\pi \in \Pi(\pi_d, \pi_c)}{\arg\min}\ \mathbb{E}_{(x_d, x_c)\sim\pi}\left[\|x_d - x_c\|^2\right] - 2\tau\,\mathcal{H}(\pi), \tag{2}$$

where $\mathcal{H}(\pi)$ denotes the joint entropy of the transport plan π, and τ denotes the regularization hyperparameter.

Proof: Please refer to [1] for the detailed proof. Briefly, the objective is to minimize the discrepancy between π_d and π_c while regularizing the transport plan π by its entropy. This ensures that the transport process remains smooth and avoids overfitting. The entropy term $\mathcal{H}(\pi)$ penalizes complex transport plans, promoting a smoother and more generalizable solution.

Theorem 1. Tighter Generalization Error Bound of the Proposed DDSB. *Compared to the Schrödinger Bridge baseline, the proposed Degradation-aware Dynamic Schrödinger Bridge (DDSB) has a tighter generalization error bound.*

Proof: The generalization error is linked to the entropy regularization term $\mathcal{H}(\pi)$. The term $\frac{1}{N}$ reflects the error due to the finite number of transport steps, while the term τ controls the smoothness of the transport plan. As the number of transport steps increases, the transport plan approximates the true coupling between π_d and π_c, thus reducing the generalization error. Moreover, increasing τ reduces the impact of entropy regularization, helping to prevent overfitting and further improving generalization. To summarize, the generalization error for the entropy-regularized optimal transport formulation, i.e. Schrödinger Bridge baseline, is bounded as:

$$\mathcal{E}_{\text{gen}}^{\text{SB}} \leq O\left(\frac{1}{N} + \tau\right), \tag{3}$$

where N denotes the number of transport steps, and τ denotes the temperature hyperparameter.

The DDSB method incorporates the *degradation-aware term* λ, which helps the model remain consistent with the degradation process during the restoration. This term reduces error accumulation during the iterative restoration process and improves the generalization performance. The total generalization error is composed of three components: 1. The transport error $O\left(\frac{1}{N}\right)$, which decreases as the number of transport steps increases. 2. The entropy regularization error $O(\tau)$, which ensures smooth transport. 3. The degradation fidelity term $O(\lambda)$, which helps the restoration process respect the underlying degradation. Thus, the generalization error for DDSB is bounded as:

$$\mathcal{E}_{\text{gen}}^{\text{DDSB}} \leq O\left(\frac{1}{N} + \tau + \lambda\right). \tag{4}$$

Combining Eq. 3 and Eq. 4, we conclude that:

$$\mathcal{E}_{\text{gen}}^{\text{DDSB}} \leq \mathcal{E}_{\text{gen}}^{\text{SB}} - O(\lambda), \tag{5}$$

demonstrating that the inclusion of the degradation-aware term λ results in a tighter error bound for DDSB.

We conclude this subsection by the following remark. This theoretical analysis demonstrates that the proposed DDSB method provides superior generalization performance compared to the SB baseline.

A.2 Connection to Physical Model based Degradation

In this subsection, we provide a theoretical analysis of the *Degradation-aware Dynamic Schrödinger Bridge (DDSB)*, focusing on its connection to *degradation-aware* techniques and the integration of a *physical degradation model*. This analysis formalizes how the proposed method ensures realistic image restoration by incorporating the degradation process into the transport framework and learning to reverse it.

Degradation-aware Transport (DOT) and its Role. The core innovation in DDSB lies in the *Degradation-aware Optimal Transport (DOT)* term. Traditional Schrödinger Bridge (SB) methods focus on learning a transport plan between the degraded and clean image distributions, without considering the physical degradation process that generated the degraded image. DDSB, on the other hand, introduces a *degradation-aware component*, which ensures that the transport process respects the underlying degradation dynamics at every step.

Let π_d and π_c denote the degraded and clean image distributions, respectively. The objective in DDSB is to learn a transport plan $\pi^\star$ that minimizes the discrepancy between these distributions while respecting the degradation model. The optimal transport cost is modified to include the degradation-aware term:

$$c_{\mathrm{DOT}}(\boldsymbol{x}_d, \boldsymbol{x}_c) = \|\boldsymbol{x}_d - \boldsymbol{x}_c\|^2 + \lambda \cdot \|D_\phi(\boldsymbol{x}_c) - \boldsymbol{x}_d\|^2, \tag{6}$$

where $\boldsymbol{x}_d$ is a sample from the degraded image distribution π_d, $\boldsymbol{x}_c$ is a clean image sample, $D_\phi(\boldsymbol{x}_c)$ denotes the degraded version of $\boldsymbol{x}_c$ predicted by the learned degradation model D_ϕ, and λ is a hyperparameter controlling the strength of the degradation fidelity term.

The term $\lambda \cdot \|D_\phi(\boldsymbol{x}_c) - \boldsymbol{x}_d\|^2$ enforces that the transport process not only minimizes the discrepancy between $\boldsymbol{x}_d$ and $\boldsymbol{x}_c$, but also ensures that the restoration process is consistent with the learned degradation model D_ϕ. This degradation-aware term helps maintain physical realism during the restoration process and prevents the generation of unrealistic artifacts.

Physical Degradation Model D_ϕ and its Role. In real-world image restoration, the degradation process is typically complex, involving factors such as noise, blur, and distortion. In many cases, it is impractical to assume a simple degradation model (such as Gaussian noise). Therefore, DDSB uses a *learnable degradation model* D_ϕ, which is trained to simulate various degradation processes, such as motion blur, fog, and noise. The model D_ϕ learns to predict the degradation process for a given clean image, providing a more accurate representation of real-world degradation dynamics.

The degradation model D_ϕ is integrated into the transport process by penalizing discrepancies between the degraded version of the predicted clean image and the degraded input image. Specifically, at each transport step, we introduce a degradation consistency term that ensures the restored image $\boldsymbol{x}_c$ remains aligned with the degradation process at each intermediate step. This term is added to the transport cost as follows:

$$\mathcal{L}_{\mathrm{DOT}} = \mathbb{E}_{q_\theta(\boldsymbol{x}_d)} \left[\|\boldsymbol{x}_d - \boldsymbol{x}_c(\boldsymbol{x}_d)\|^2 + \lambda \|D_\phi(\boldsymbol{x}_c(\boldsymbol{x}_d)) - \boldsymbol{x}_d\|^2 \right], \tag{7}$$

where $\boldsymbol{x}_d$ is a degraded sample, $\boldsymbol{x}_c(\boldsymbol{x}_d)$ denotes the predicted clean image at a given transport step, and $D_\phi(\boldsymbol{x}_c(\boldsymbol{x}_d))$ denotes the degraded version of $\boldsymbol{x}_c(\boldsymbol{x}_d)$ predicted by the degradation model D_ϕ.

This ensures that each intermediate image in the restoration trajectory follows the degradation process as closely as possible, improving the physical realism of the restoration.

Degradation Amplification for Realism. In this approach, the degradation model D_ϕ is not only used to predict the degradation but is also employed to *amplify the degradation* at each step. This means that the model simulates the inverse of the restoration process by artificially degrading the predicted clean image, making the difference between the degraded image $\boldsymbol{x}_d$ and the predicted clean image more pronounced. This helps in preventing the model from generating unrealistic images that would not correspond to any physical degradation.

Formally, the degradation model D_ϕ is learned to *amplify* the degradation, i.e., simulate the inverse process of restoration. The objective is to penalize large deviations between the degraded input $\boldsymbol{x}_d$ and the predicted clean image $\boldsymbol{x}_c$ by introducing a degradation consistency term, as described in the DOT loss:

$$\|D_\phi(\boldsymbol{x}_c) - \boldsymbol{x}_d\|^2, \tag{8}$$

which ensures that the model's restoration trajectory remains consistent with the degradation process, leading to more realistic results.

Integrating the Degradation Model into the Transport Process. The key idea behind DDSB is to integrate the *degradation model* D_ϕ into the *optimal transport* framework. The method learns a transport plan $\pi^\star$ that minimizes the distance between the degraded and clean images while ensuring that intermediate steps in the restoration process respect the physical degradation dynamics. This is achieved by adding the degradation-aware term to the transport cost.

At each step of the restoration, we aim to *align the transport plan with the degradation process*. The total loss for DDSB can be expressed as the sum of the standard SB loss and the degradation-aware term:

$$\mathcal{L}_{\text{DDSB}} = \mathcal{L}_{\text{SB}} + \mathcal{L}_{\text{DOT}}, \tag{9}$$

where $\mathcal{L}_{\text{SB}}$ denotes the standard Schrödinger Bridge loss (transport loss). In addition, $\mathcal{L}_{\text{DOT}}$ denotes the degradation-aware optimal transport loss, which ensures that the restoration process respects the degradation model D_ϕ.

This dual loss function encourages the model to minimize the transport cost while ensuring that the restored images are consistent with the degradation process, leading to more realistic and physically plausible restoration results.

We conclude this subsection by the following remark. The introduction of the degradation-aware term λ in the DOT loss improves the generalization of the restoration model. By enforcing consistency with the degradation process, DDSB ensures that the model can generalize well to unseen degradation types. This also reduces the risk of generating unrealistic artifacts, which is a common problem in unpaired image restoration tasks.

The *Degradation-aware Dynamic Schrödinger Bridge (DDSB)* method introduces a novel integration of *degradation-aware optimal transport* and a *physical degradation model*. By incorporating the degradation model D_ϕ into the transport process, DDSB ensures that the restoration trajectory remains consistent with the physical degradation process, improving the realism of the restored images. This is achieved by penalizing discrepancies between the degraded input and the predicted clean images, leading to a more physically plausible restoration process that reduces artifacts and improves generalization to unseen degradation types.

This theoretical analysis, along with the introduced degradation-aware and dynamic transport components, lays the foundation for DDSB's superior performance in unpaired image restoration, addressing the key challenges of realism and generalization.

B More Visual Results

B.1 On Multi-Task Image Restoration

More visual results of multi-task image restoration are provided in Fig. 1 and Fig. 2.

B.2 On Generalized Haze Removal

More visual results of generalized haze removal are provided in Fig. 3 and Fig. 4.

References

[1] Aude Genevay. *Entropy-regularized optimal transport for machine learning*. PhD thesis, Université Paris sciences et lettres, 2019.

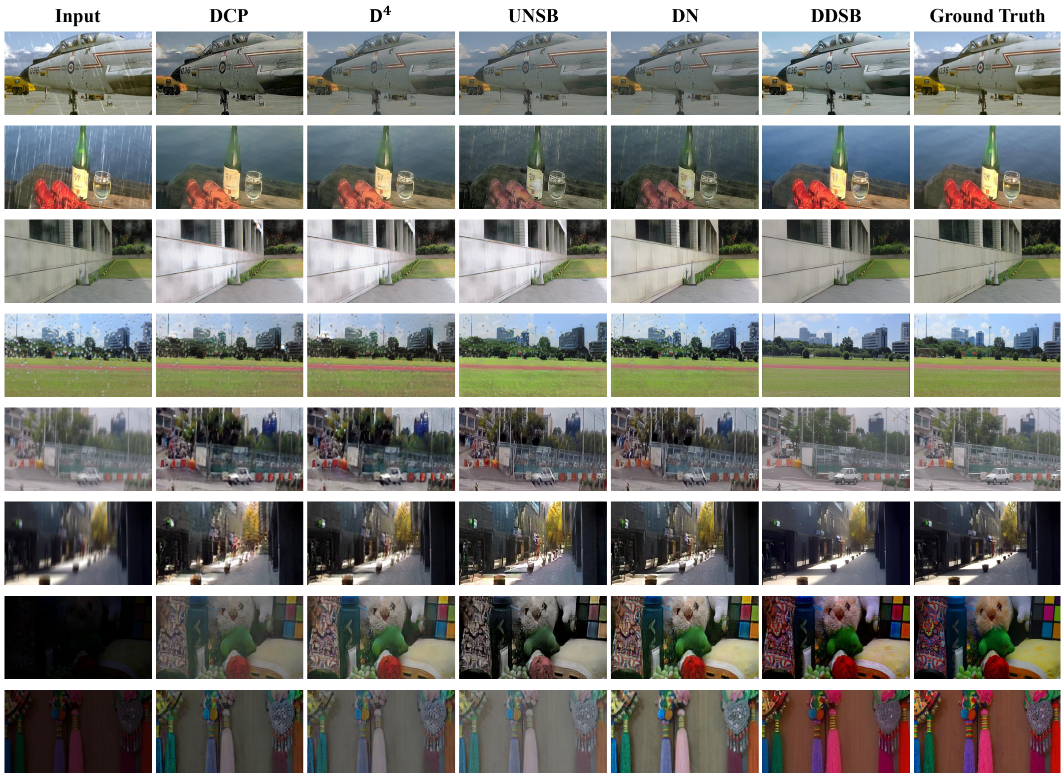

Figure 1: More results of multi-task image restoration

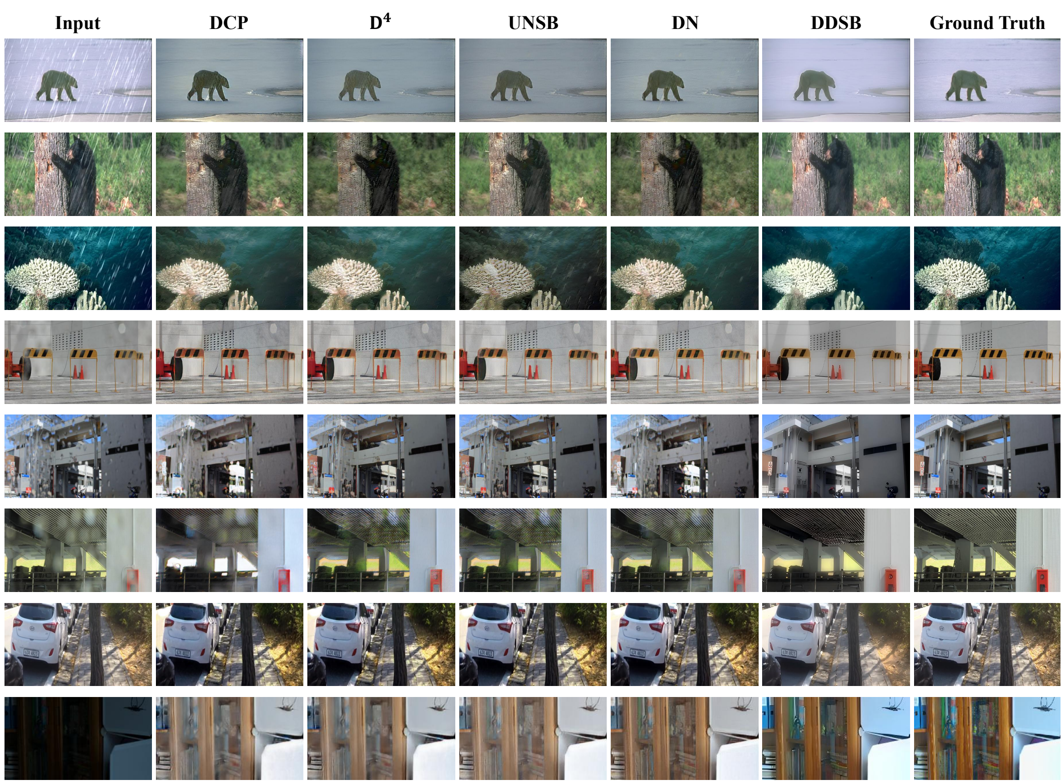

Figure 2: More results of multi-task image restoration

Figure 3: More results of generalized haze removal

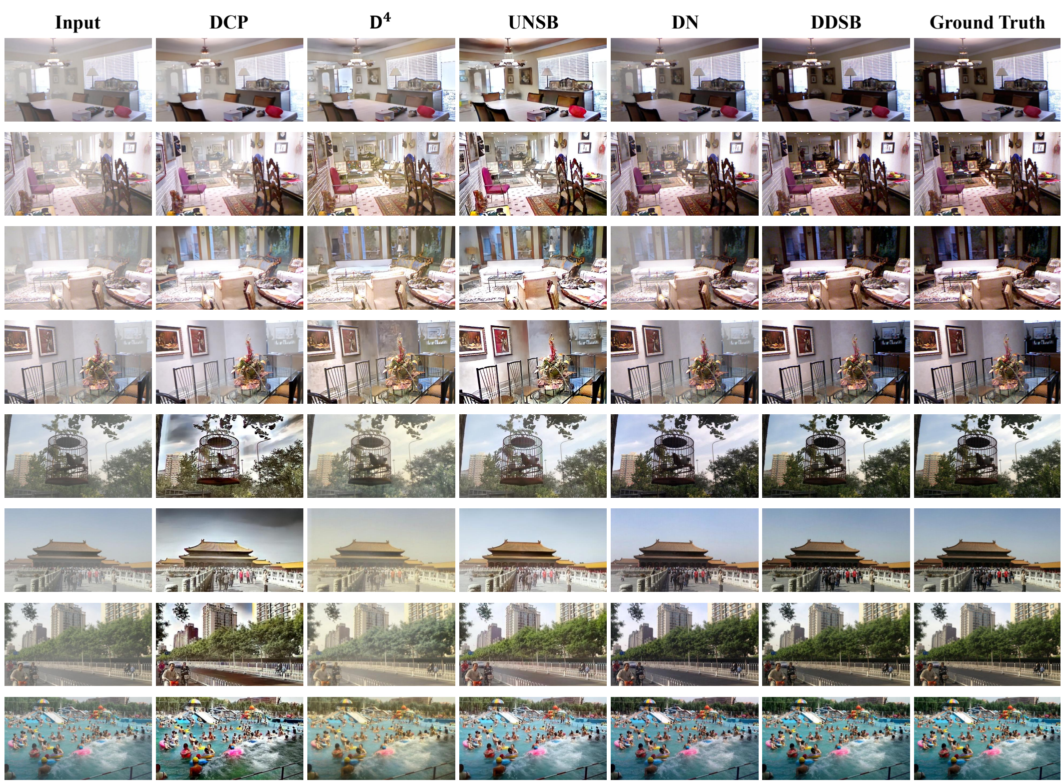

Figure 4: More results of generalized haze removal