Beyond Unidirectional: Unsupervised Trajectory Learning for Omnidirectional Controllable Underwater Image Enhancement

Xiaofeng Cong
Hefei University of Technology
Hefei, China
cxf_svip@163.com

Xuhao Guo
Southeast University
Nanjing, China
220255561@seu.edu.cn

Kang Zheng
Southeast University
Nanjing, China
220255583@seu.edu.cn

Yu-Xin Zhang
Southeast University
Nanjing, China
yuxinzhang@seu.edu.cn

Jie Gui*
Southeast University
Nanjing, China
guijie@seu.edu.cn

Abstract

Existing underwater image enhancement (UIE) methods predominantly rely on unidirectional mapping, which often results in over-enhancement or under-enhancement, without user-friendly controllability. To address this, we propose a paradigm of unsupervised trajectory learning for Omnidirectional Controllable UIE that treats restoration as an omnidirectional dynamical system, namely OC-UIE. The main purpose of our work is the realization of a consistent omnidirectional mapping across the representation space, which allows the model to master complex underwater dynamics beyond traditional unidirectional constraints. To implement this, we propose an Omnidirectional Training that optimizes over arbitrary source-target positions sampled from the data distribution. We decompose this purpose into two synergistic stages: first, Representation Trajectory Flattening is employed to organize intricate degradations into a unified adaptation axis; second, Trajectory Omnidirectional Integration is introduced to model the enhancement as a path integral over the resulting manifold. By optimizing only the relative shifts, OC-UIE isolates structural preservation from fluid degradation dynamics. This approach provides a controllable space for intensity-parameterized enhancement with state-of-the-art performance. The code will be released at https://github.com/Xiaofeng-life/OC-UIE.

1 Introduction

Underwater sensing is vital for marine monitoring and exploration, yet its reliability is constrained by the optical properties of water. Absorption and scattering cause wavelength-dependent energy dissipation [64, 13]. The rapid attenuation of red and yellow light relative to blue and green bands distorts the spectral balance, resulting in a pervasive blue-green hue [59, 54, 30, 48, 46, 4, 61, 62, 11].

Existing underwater image enhancement methodologies are categorized into non-deep learning-based and deep learning-based paradigms [55, 56, 16, 33, 64, 53]. Non-deep learning-based approaches primarily rely on the formulation of physical imaging models and the application of empirical prior assumptions to restore image quality [6, 37]. While these methods demonstrate notable

*Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

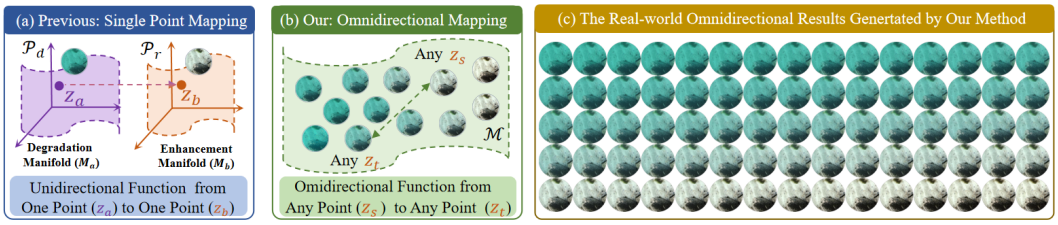

Figure 1: Contribution of our framework. (a) Unidirectional Training and Inference: Conventional paradigms are typically restricted to learning point-to-point mappings. (b) Omnidirectional Training: Our OC-UIE captures trajectory relationships between arbitrary states. (c) Omnidirectional Inference: Our OC-UIE enabling flexible transitions and can generate diverse controllable state.

effectiveness in controlled environments, their robustness is often limited by the significant variability and heterogeneity of optical conditions in complex marine scenes [25, 54, 3, 17, 44, 47].

Deep learning-based underwater restoration has evolved through two primary dimensions: structural backbones and training paradigms [60, 41, 45, 32, 51, 24, 21, 63]. Structurally, researchers have transitioned from convolutional neural networks (Shallow-UWNet [26], WaterNet [19]) to fine-grained local patterns, including HCLR-Net [60], MFEN [7], and SS-UIE [29]. Furthermore, the Transformer-based architectures [34], namely AutoEnhancer [39], UFormer [30] and LGT [34], are proposed for global context learning. Fourier-based operators for frequency-domain spectral attenuation are used by FLD [35] and CDF-UIE [52]. For the learning paradigms, the GAN-based UW-CycleGAN [49], WSDS-GAN [22] and INGCGAN [21], reconstruct fine-grained textures via adversarial training. While Flow-based model WaterFlow [57] utilize exact likelihood estimation for invertible mappings between distorted and clear scenes. Recently, diffusion-based methods DM-UIE [40], WF-Diff [58], SeaDiff [2] and WaterCDiff [42] achieve exceptional quality by iteratively denoising visual signals within the latent space of underwater scenes. However, recent efforts have prioritized controllable underwater image enhancement, recognizing that fixed-output models are often insufficient for the highly ill-posed nature of marine restoration [10, 18, 9]. Nevertheless, these methods still largely adhere to unidirectional trajectories, leaving a critical research gap for an omnidirectional controllable paradigm that empowers users to navigate freely across the entire degradation-restoration continuum.

Fundamentally, the bottleneck that prevents existing underwater image enhancement frameworks from attaining omnidirectional controllability lies in their rigid reliance on unidirectional, point-to-point mappings between predetermined boundary states. By formulating restoration through such a static and discrete lens, these paradigms conceptualize degradation and recovery as disjoint, isolated regimes rather than coupled states on a continuous manifold. Consequently, they fail to characterize the intricate velocity fields and intermediate spectral trajectories that govern underwater light transport, inherently neglecting the dynamic continuity required to traverse backward and forward across arbitrary optical depths. Specifically, this prevailing unidirectional formulation restricts the restoration process to an isolated, point-to-point deterministic mapping between predefined degradation and clean states, which can be formally expressed as:

$$\mathcal{F}_{\text{prev}} : z_a \in \mathcal{P}_d \longrightarrow z_b \in \mathcal{P}_r, \tag{1}$$

where z_a and z_b are two anchors. As shown in Fig. 1(a), existing unidirectional models only achieve uncontrollable one-way mapping. In contrast, our OC-UIE conceptualizes the state space as a unified manifold $\mathcal{M}$. By introducing a mapping function $\mathbf{V}(\sigma)$, our framework facilitates an omnidirectional traversal between arbitrary states z_s and z_t parameterized by τ_s and τ_t:

$$\mathcal{F}_{\text{omni}} : z_s \xrightarrow{\int_{\tau_s}^{\tau_t} \mathbf{V}(\sigma)d\sigma} z_t, \quad \forall z_s, z_t \in \mathcal{M}, \tag{2}$$

where $\int \mathbf{V}(\sigma)d\sigma$ encapsulates the trajectory integration discovered by our paradigm. As illustrated in Fig. 1(b), our method enables the connection between arbitrary points, achieving a consistent omnidirectional mapping. As demonstrated by the progressive state transitions in Fig. 1(c), the visualization illustrates a evolution of the underwater scene along the learned trajectory, performing smooth and logical state transitions with a fine-grained step. The main purpose of our work is the realization of a consistent omnidirectional mapping across the representation space, providing the fundamental capability to master complex underwater dynamics.

To realize this enhancement objective, we propose an Omnidirectional Training paradigm that systematically breaks free from static pairing constraints by formulating an optimization strategy over arbitrarily sampled source-target coordinate pairs across the underlying data distribution. The operational pipeline of this paradigm unfolds through two mutually reinforced phases. First, Representation Trajectory Flattening is introduced to untangle the convoluted, high-dimensional space of aquatic degradations. By establishing two diametrically opposed operator bases along an antipodal trajectory, it aligns heterogeneous optical distortions along a unified adaptation axis, successfully revealing the intrinsic geometric anchors of the latent representation manifold. Building upon this regularized latent landscape, Trajectory Omnidirectional Integration then formulates the continuous transition between arbitrary points as a dynamic path integral over the state manifold. By distilling the learned trajectory velocity field into directional steering operators, this stage achieves state-dependent, fine-grained visual modulation that faithfully maintains structural consistency across arbitrary degradation intensities. The main contributions of this work are summarized as follows:

- We propose the OC-UIE framework, which redefines underwater enhancement as an omnidirectional dynamical system. By shifting from static mappings to an **Omnidirectional Training** strategy, our approach enables the model to master consistent state transitions, facilitating seamless navigation across arbitrary states with state-of-the-art fidelity.
- We introduce **Representation Trajectory Flattening** as a mechanism to organize the complex, high-dimensional underwater degradation space into a structured adaptation axis. This method effectively decomposes intricate degradation dynamics into polarity-aligned bases, ensuring that opposing enhancement dimensions are decoupled and mapped onto a geometrically interpretable manifold.
- We develop **Trajectory Omnidirectional Integration** to formulate the enhancement process as a path integral via a geometry-aware controller. By distilling the learned trajectory dynamics, this approach enables intensity-parameterized transformations along the manifold path, providing omnidirectional results across diverse aquatic environments.

2 Related Work

Non-deep Learning-based Methods. Despite extensive research, underwater image restoration remains formidable due to non-uniform light propagation [24, 21, 41, 45]. Early efforts focused on stochastic modeling, statistical heuristics, and physical priors [6, 37, 31, 36]. IBLA [31] restores images by analyzing the relationship between blurriness and absorption to estimate scene depth. ULAP [36] utilizes an attenuation prior to constrain the physical imaging model's inversion. To reduce computational complexity, HLRP [64] employs a Hyper-Laplacian reflectance prior, decomposing optimization into tractable sub-problems. Similarly, SMBL [37] combines background light estimation with multi-prior fusion to compensate for wavelength-dependent attenuation. However, these prior-driven frameworks often struggle with chromatic consistency and structural fidelity in stochastic real-world environments [10, 60]. Their rigid reliance on idealized assumptions frequently produces artifacts or fails to satisfy diverse user preferences [39, 30, 34].

Deep Learning-based Methods. In contrast, deep-learning-based methods have revolutionized UIE by designing diverse network architectures to capture high-dimensional feature representations. These architectures range from foundational Convolutional Neural Networks [26, 19] and Attention-based mechanisms [60, 7, 29] to more advanced paradigms, including Transformer-based frameworks [39, 30, 34], Fourier-inspired operators [35, 52], Mamba-based state space models [38], and Diffusion-based generative models [58, 2, 42]. Shallow-UWNet [26] adopts a lightweight, streamlined convolutional architecture featuring cross-layer connectivity to facilitate an efficient end-to-end mapping for underwater image restoration. HCLR-Net [60] proposes a hybrid attention mechanism specifically engineered to concurrently capture and refine complex underwater distortion features, effectively mitigating spectral and spatial degradation through multi-scale feature dependencies. CDF-UIE [52] introduces a spatial-frequency decoupling block that decomposes latent features into spatial and Fourier components, leveraging the global receptive field of Fourier transforms to achieve high-fidelity spectral reconstruction. PGMamba [38] introduces a physical model-guided Mamba framework that integrates underwater optical principles with a spatial-aware global scanning strategy to achieve efficient long-range dependency modeling and enhanced restoration fidelity. SeaDiff [2] introduces a degradation-aware diffusion model that embeds wavelength- and distance-

dependent physical priors into the generative process, effectively overcoming data scarcity to achieve high-fidelity restoration.

Despite their impressive restoration performance, these deep learning-based architectures remain inherently constrained by their deterministic formulation, mapping a single degraded observation to a fixed point in the solution space [10, 18, 9]. Such a rigid setting proves inadequate for underwater image restoration, which is fundamentally a severely ill-posed inverse problem characterized by massive solution ambiguity, where multiple plausible clean scenes often correspond to the same degraded measurement under varying illumination and optical water types.

Recent advancements have prioritized controllable underwater image enhancement, acknowledging that deterministic models are often inadequate for this highly ill-posed restoration task. Existing strategies, such as CECF [10] and PWAE [18], achieve condition-aware enhancement via auxiliary guidance, while UIESS [9] employs latent disentanglement for test-time intensity adjustment. Despite their efficacy, these frameworks treat enhancement as heuristic style transfer rather than inverse imaging. By manipulating latent style codes instead of physical parameters, they decouple from the spectral dynamics essential for consistent and controllable restoration. In contrast, OC-UIE treats enhancement as a continuous, omnidirectional dynamical system over an operator manifold. Instead of heuristic style shifts, it optimizes path integrals along a geometrically aligned trajectory, enabling physically grounded, bidirectional transitions between arbitrary degradation states.

Algorithm 1: Pipeline of the OC-UIE

Input: Model weights $\{\mathcal{E}, \Phi, \mathcal{D}\}$
Output: Model weights $\pi_{\mathcal{E}}$ and $\pi_{\mathcal{D}}$, π_{Φ}^{α}, π_{Φ}^{β}
// Trajectory Flattening
Initialize $\mathcal{B}_{\mathcal{E}} = \{\phi_{\mathcal{E}}, \psi_{\mathcal{E}}\}$, $\mathcal{B}_{\mathcal{D}} = \{\phi_{\mathcal{D}}, \psi_{\mathcal{D}}\}$, and shared U-Net operator ω_{Φ};
while *not converged* **do**
  // Independent sampling
  Sample $x_\alpha \sim \mathcal{X}_\alpha$ and $x_\beta \sim \mathcal{X}_\beta$;
  Compute $z_\alpha = \mathcal{E}(x_\alpha)$ and $z_\beta = \mathcal{E}(x_\beta)$;
  // Update basis operators
  Update $\Theta \leftarrow \Theta - \eta \nabla_{\Theta}(\mathcal{L}_m + \lambda_p \mathcal{L}_d)$;
end
// Save weights for flattening
$\mathcal{W}_{\mathcal{E}} = \{\mathbf{W}_{\phi}^{\mathcal{E}}, \mathbf{W}_{\psi}^{\mathcal{E}}\}$, $\mathcal{W}_{\mathcal{D}} = \{\mathbf{W}_{\phi}^{\mathcal{D}}, \mathbf{W}_{\psi}^{\mathcal{D}}\}$;
// Trajectory Integration
Initialize trajectory bases π_{Φ}^{α} and π_{Φ}^{β};
while *not converged* **do**
  Sample states τ_s, τ_t from $\mathcal{U}(0, 1)$;
  // Generate state samples
  $\mathbf{W}_{\tau}^{\mathcal{E}} = (1-\tau)\mathbf{W}_{\psi}^{\mathcal{E}} + \tau \mathbf{W}_{\phi}^{\mathcal{E}}$;
  $\mathbf{W}_{\tau}^{\mathcal{D}} = (1-\tau)\mathbf{W}_{\psi}^{\mathcal{D}} + \tau \mathbf{W}_{\phi}^{\mathcal{D}}$;
  $x_{\tau}^{s} \leftarrow$ Inference with $(\mathbf{W}_{\tau}^{\mathcal{E}}[s], \mathbf{W}_{\tau}^{\mathcal{D}}[s])$;
  $x_{\tau}^{t} \leftarrow$ Inference with $(\mathbf{W}_{\tau}^{\mathcal{E}}[t], \mathbf{W}_{\tau}^{\mathcal{D}}[t])$;
  // Omnidirectional optimization
  $\hat{x}_{\tau}^{t} = \mathcal{D}\left(\Phi_{pre}(\mathcal{E}(x_{\tau}^{s})) + \Delta\pi_{\Phi}\right)$;
  // Optimize trajectory bases
  $\{\pi_{\Phi}^{\alpha}, \pi_{\Phi}^{\beta}\} \leftarrow \{\pi_{\Phi}^{\alpha}, \pi_{\Phi}^{\beta}\} - \eta \nabla_{\{\pi_{\Phi}^{\alpha}, \pi_{\Phi}^{\beta}\}} \mathcal{L}_o$;
end

3 Methods

As illustrated in Fig. 2, our approach comprises two distinct yet synergistic stages: Representation Trajectory Flattening and Trajectory Omnidirectional Integration. We use a frozen encoder $\mathcal{E}$, U-Net Φ, and decoder $\mathcal{D}$, parameterized via low-rank based operators [50]. We define Initial Network as shown in Fig. 2(a)①: $\mathcal{E}$ and $\mathcal{D}$ as pairs $\mathcal{B}_{\mathcal{E}} = \{\phi_{\mathcal{E}}, \psi_{\mathcal{E}}\}$ and $\mathcal{B}_{\mathcal{D}} = \{\phi_{\mathcal{D}}, \psi_{\mathcal{D}}\}$, representing enhancement and degradation dynamics, and a shared operator ω_{Φ} captures representation features. The Initial Network utilizes $\mathcal{G}_{\phi} = \{\phi_{\mathcal{E}}, \omega_{\Phi}, \phi_{\mathcal{D}}\}$ and $\mathcal{G}_{\psi} = \{\psi_{\mathcal{E}}, \omega_{\Phi}, \psi_{\mathcal{D}}\}$ to organize the degradation space. During integration, as shown in Fig. 2(b)②, we refine the framework using unified pairs $\mathcal{B}_{\mathcal{E}} = \{\pi_{\mathcal{E}}\}$ and $\mathcal{B}_{\mathcal{D}} = \{\pi_{\mathcal{D}}\}$ to maintain structural consistency. Crucially, the U-Net is expanded into two trajectory bases, π_{Φ}^{α} and π_{Φ}^{β}, serving as antipodal polarities for omnidirectional navigation. Section 3.1 elaborates on the Representation Trajectory Flattening, where we describe the organization of polarity-aligned adaptation bases along a unified latent axis. Section 3.2 details the Trajectory Omnidirectional Integration in OC-UIE, explaining how we formulate the enhancement as a path integral over the learned manifold to achieve omnidirectional control via self-guided optimization.

3.1 Representation Trajectory Flattening

In the trajectory flattening stage, we identify two diametrically opposed operators, Ω_{α} (degradation-oriented) and Ω_{β} (restoration-oriented), to define the extreme polarities of the latent manifold $\mathcal{M}$.

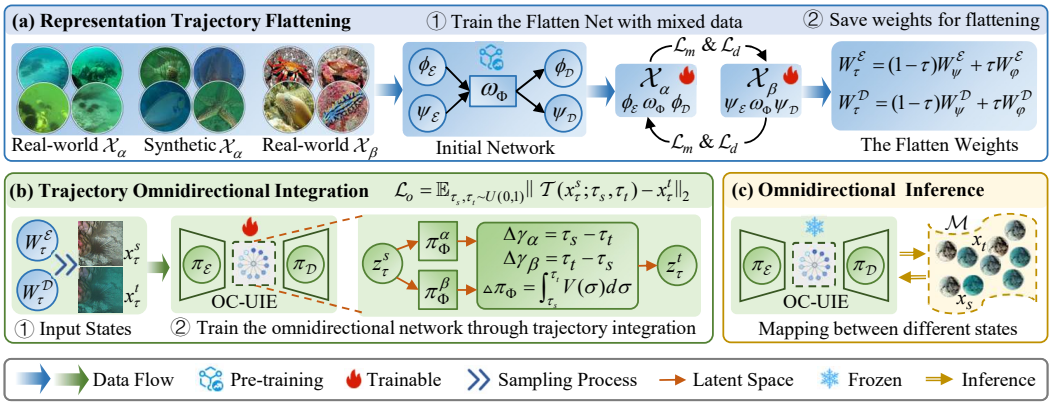

Figure 2: Framework of OC-UIE. (a) Representation Trajectory Flattening by learning the basis for dynamic representations. (b) Trajectory Omnidirectional Integration training by the proposed omnidirectional loss. (c) Inference stage with controllable behavior by user-defined state x_s to x_t.

We define the flattening mapping $\mathcal{F}$ as a traversal along the axis $\mathcal{S}$:

$$\mathcal{F}(z_\alpha;\mathcal{S}) : z_\alpha \in \mathcal{M}_\alpha \xrightarrow{\mathcal{S}=\text{span}(\Omega_\alpha - \Omega_\beta)} z_\beta \in \mathcal{M}_\beta, \quad \text{s.t.} \quad \langle \frac{\Omega_\alpha}{\|\Omega_\alpha\|}, \frac{\Omega_\beta}{\|\Omega_\beta\|} \rangle \to -1, \tag{3}$$

where $z = \mathcal{E}(x)$. This alignment flattens complex dynamics into a unified trajectory, providing a structured geometry for omnidirectional navigation. During training as shown in Fig. 2(a)①, the pretrained weights Φ_{pre} remain frozen, while the low-rank trajectory operators $\Theta = \{\phi_\alpha, \phi_\beta\}$ are optimized. To enforce the desired geometric relationship, we impose a diametric alignment loss $\mathcal{L}_d$ on the low-rank decomposition matrices $\mathbf{A}_m$ and $\mathbf{B}_m$ across all layers:

$$\mathcal{L}_d = \frac{1}{M} \sum_{m=1}^{M} \sum_{\mathbf{W}\in\mathbf{A},\mathbf{B}} \left| \langle \bar{\mathbf{W}}_m^{(\alpha)}, \bar{\mathbf{W}}_m^{(\beta)} \rangle + 1 \right|_F^2, \tag{4}$$

where $\bar{\mathbf{W}}$ denotes the normalized weight vector. This constraint forces the low-rank updates to anchor at antipodal locations on the manifold, ensuring that the discovered α and β dynamics represent inverted structural transitions. This stage aims to discover the low-rank basis operators through an unsupervised cycle-consistency training framework [49, 28]. We construct synthetic data using [1] and incorporate real-world images [15] to facilitate auxiliary training. We learn the mapping between manifolds by minimizing a joint objective that combines reconstruction and adversarial alignment:

$$\mathcal{L}_m = \sum_{k\in\{\phi,\psi\}} \left(\|\mathcal{G}_{\bar{k}}(\mathcal{G}_k(x_k)) - x_k\|_1 + \lambda_a \mathbb{E}_x \left[-\log \sigma(D_k(\mathcal{G}_k(x))) \right] \right), \tag{5}$$

where $\mathcal{D}_k$ denotes the discriminator, forcing the low-rank bases to synthesize target-consistent details through a sigmoid loss. $\mathcal{G}_k$ and $\mathcal{G}_{\bar{k}}$ denote the forward and reverse translation operators, respectively. The total objective is formulated as $\mathcal{L}_{total} = \mathcal{L}_m + \lambda_p \mathcal{L}_d$, ensuring both manifold alignment and the diametric geometric constraints of the discovered operators. To generate continuous enhancement trajectories for the subsequent integration stage, we perform interpolation directly in the low-rank weight space between the discovered endpoints. Given the low-rank basis weights for the encoder $\mathcal{W}_\mathcal{E} = \{\mathbf{W}_\phi^\mathcal{E}, \mathbf{W}_\psi^\mathcal{E}\}$ and the decoder $\mathcal{W}_\mathcal{D} = \{\mathbf{W}_\phi^\mathcal{D}, \mathbf{W}_\psi^\mathcal{D}\}$, the Symmetric Interpolation at any intermediate state $\tau \in [0, 1]$ is formulated as:

$$\mathbf{W}_\tau^\mathcal{E} = (1-\tau)\mathbf{W}_\psi^\mathcal{E} + \tau \mathbf{W}_\phi^\mathcal{E}, \mathbf{W}_\tau^\mathcal{D} = (1-\tau)\mathbf{W}_\psi^\mathcal{D} + \tau \mathbf{W}_\phi^\mathcal{D}, \tag{6}$$

where $\mathbf{W}_\tau^\mathcal{E}$ and $\mathbf{W}_\tau^\mathcal{D}$ represent the synchronized low-rank adaptations applied to the frozen encoder and decoder, respectively. By modulating these weights while keeping the U-Net basis ω_Φ fixed, we synthesize a sequence of training samples that traverse the geodesic path of the underwater manifold. This weight-space interpolation ensures that the generated samples capture the fluid transition of spectral and scattering variations, providing a dense supervision signal for the subsequent trajectory integration stage.

3.2 Trajectory Omnidirectional Integration

To provide a physical grounding for the latent evolution, we relate our trajectory polarities to the revised underwater image formation model [1]. According to [1], the intensity of an image channel $c \in \{R, G, B\}$ formed at the sensor at a distance z is obtained by integrating the spectral signals:

$$I_c = \frac{1}{\kappa} \int_{\lambda_1}^{\lambda_2} S_c(\lambda)\rho(\lambda)E(d,\lambda)e^{-\beta(\lambda)z}d\lambda + \frac{1}{\kappa} \int_{\lambda_1}^{\lambda_2} S_c(\lambda)B^{\infty}(\lambda)(1-e^{-\beta(\lambda)z})d\lambda, \tag{7}$$

where $S_c(\lambda)$ represents the sensor response for wavelength λ, $\rho(\lambda)$ is the object reflectance, $E(d,\lambda)$ denotes the irradiance at depth d, B^{∞} is the wideband backscatter, κ is a scalar governing image exposure and camera pixel geometry. This integral formulation encapsulates the non-linear spectral attenuation $\beta(\lambda)$ and the range-dependent backscatter into a physical framework.

To bridge the gap between physical degradation and latent evolution, we relate the basis to the spectral attenuation laws. However, as the spectral response is implicitly coupled within the latent space, we reformulate the enhancement as an integral manifold mapping over the discrete degradation intensity τ. We define the basis π_{Φ}^{α} and π_{Φ}^{β} as the boundary states corresponding to the minimum and maximum optical thicknesses $\{\gamma_{min}, \gamma_{max}\}$. Any intermediate state γ_t is formulated as the accumulation of state γ_s along the manifold:

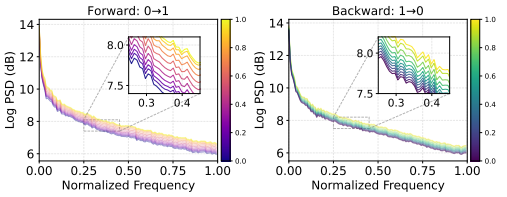

Figure 3: Spectral analysis of image restoration and degradation trajectories.

$$\Omega(\gamma_t) = \Omega(\gamma_s) + \int_{\tau_s}^{\tau_t} \mathcal{V}(\tau)\, d\tau, \tag{8}$$

where $\mathcal{V}(\tau)$ represents the mapping function of the degradation-restoration continuum. Since the integration is numerically intractable, our trajectory integration stage distills this manifold into a deterministic controller $\mathcal{C}$ that learns to compute the weight increments between arbitrary points. Given a source level τ_s and a target level τ_t sampled during training, the steering coefficients $\Delta\gamma$ are defined by the relative displacement between these states:

$$\Delta\gamma_{\alpha} = (1-\tau_t) - (1-\tau_s), \Delta\gamma_{\beta} = \tau_t - \tau_s, \tag{9}$$

where $\Delta\gamma_{\alpha}$ and $\Delta\gamma_{\beta}$ capture the residual dynamics required for weight migration as shown in Fig. 2(b)②. To ensure stable and controllability, we formulate the total latent shift $\Delta\omega_{\Phi}$ as the integral projection of the basis along the transition path:

$$\Delta\omega_{\Phi}(\tau_s, \tau_t) = \int_{\tau_s}^{\tau_t} \mathbf{V}(\sigma)d\sigma = \Delta\gamma_{\alpha}\pi_{\Phi}^{\alpha} + \Delta\gamma_{\beta}\pi_{\Phi}^{\beta}, \tag{10}$$

where $\tau_s, \tau_t \in [0,1]$ denote the source and target degradation levels, $\mathbf{V}(\sigma)$ represents the weight evolution field that characterizes the instantaneous transition between anchors. This incremental formulation enables the model to navigate the degradation-restoration continuum by accumulating infinitesimal shifts along the learned manifold. The enhancement process for an input latent representation $z = \mathcal{E}(x)$ is formulated as a precise weight migration within the operator space:

$$\hat{z}_t = \Phi_{pre}(z) + \Delta\gamma_{\alpha}\pi_{\Phi}^{\alpha}(z) + \Delta\gamma_{\beta}\pi_{\Phi}^{\beta}(z), \tag{11}$$

where Φ_{pre} denotes the frozen pre-trained backbone, and the dual trajectory bases $\pi_{\Phi}^{\alpha}, \pi_{\Phi}^{\beta}$ provide the directional refinements. Consequently, the complete transition operator $\mathcal{T}$ at an arbitrary state (τ_s, τ_t) is defined by the feature-wise modulation governed by the cumulative latent shift:

$$\mathcal{T}(x; \tau_s, \tau_t) = \mathcal{D}\left(\Phi_{pre}(z) \oplus \Delta\pi_{\Phi}(\tau_s, \tau_t)\right), \tag{12}$$

where $\oplus$ signifies the additive modulation of the trajectory-guided dynamics. Fig. 3 illustrates the spectral evolution along the $0 \to 1$ (left) and $1 \to 0$ (right) trajectories via Radially Averaged Power Spectral Density. The profiles exhibit a monotonic shift that preserves the natural power-law distribution. This laminar behavior demonstrates that the operator manifold facilitates stable energy compensation across the frequency spectrum without introducing artifacts. The smooth decay

Table 1: Quantitative evaluation results, where **bold** denotes the best performance. The CIQA and URANK denote CLIPIQA and URANKER, respectively.

Method	Challenge-60			U45			UCCS			EUVP		
	UIQM	CIQA	URANK	UIQM	CIQA	URANK	UIQM	CIQA	URANK	UIQM	CIQA	URANK
MIP	3.627	0.516	0.099	3.875	0.488	0.110	2.465	0.339	-0.597	2.904	0.414	0.495
SMBL	3.863	0.527	0.870	3.544	0.492	0.771	4.048	0.355	1.853	3.470	0.420	1.298
IBLA	3.474	0.527	0.441	2.770	0.531	0.522	2.480	0.369	-0.216	2.898	0.437	0.929
ULAP	4.152	0.538	-0.081	3.550	0.532	-0.126	2.810	0.380	-0.086	3.705	0.426	0.801
Semi-UIR	4.098	0.541	1.417	4.431	0.547	2.032	4.049	0.439	1.335	3.660	0.455	1.818
UIE-DM	3.874	0.485	1.531	4.097	0.515	1.983	4.166	0.281	1.468	3.652	0.355	1.897
CECF	4.127	0.461	1.501	4.354	0.500	2.013	4.286	0.380	1.474	3.681	0.444	1.792
HCLRNet	4.100	**0.554**	1.253	4.293	0.541	1.970	3.935	0.390	1.088	3.680	0.451	1.706
AquaCl	4.359	0.510	1.571	4.586	**0.550**	**2.093**	4.272	0.334	1.348	3.730	0.426	1.623
ERD	4.298	0.430	1.524	4.405	0.530	1.939	4.232	0.404	1.552	3.858	0.366	1.908
CDF	3.757	0.439	1.467	3.851	0.398	1.873	3.971	0.383	1.549	3.490	0.402	1.880
MFEN	4.263	0.521	1.400	4.427	0.545	1.994	4.280	0.415	1.408	3.940	0.451	1.790
Ours	**4.769**	0.479	**1.958**	**4.854**	0.503	2.007	**4.815**	**0.526**	**2.224**	**4.587**	**0.475**	**2.235**

confirms the framework's ability to maintain structural consistency during bidirectional navigation, ensuring that both restoration and degradation remain statistically regular. To enforce the model's capability for seamless navigation across the entire degradation-restoration continuum, we define the omni-directional objective $\mathcal{L}_o$. This loss function is formulated as an expectation over arbitrary transitions within the learned manifold:

$$\mathcal{L}_o = \mathbb{E}_{\tau_s,\tau_t\sim\mathcal{U}(0,1)} \left\| \mathcal{T}(x_\tau^s; \tau_s, \tau_t) - x_\tau^t \right\|_2, \tag{13}$$

where x_τ^s and x_τ^t represent underwater observations sampled at random intensities τ_s and τ_t from a uniform distribution $\mathcal{U}(0,1)$. Unlike traditional unidirectional restoration losses, $\mathcal{L}_o$ supervises the controller to learn a consistent field for any arbitrary starting point and destination. Fig. 2(c) shows the inference stage. The training details are in Algorithm 1.

4 Experiments

Baseline Comparisons and Datasets. To demonstrate the efficacy of our method, we benchmark it against a comprehensive suite of state-of-the-art (SOTA) underwater image restoration methods. These baselines are categorized into two groups: traditional prior-based approaches, including MIP [6], SMBL [37], IBLA [31], and ULAP [36]; and recent learning-based frameworks, namely Semi-UIR [14], UIE-DM [40], CECF [10], HCLRNet [60], AquaCl [8], ERD [5], CDF [52], and MFEN [7]. Learning-based baselines are retrained on the [19]. We focus on evaluating all models on a consistent set of challenging real-world benchmarks, such as UCCS [23], the real-world test set of EUVP [15], U45 [20], and the Challenging-60 [19].

Evaluation Metrics. To quantitatively evaluate the image enhancement performance, we adopt three widely established non-reference quality metrics: UIQM [27], CLIPIQA [43], and URANKER [12].

Implementation Details. The rank r for the low-rank weight matrices in both the VAE and UNet is set to 8. The model is optimized using the Adam optimizer with a learning rate of 1×10^{-4} and a

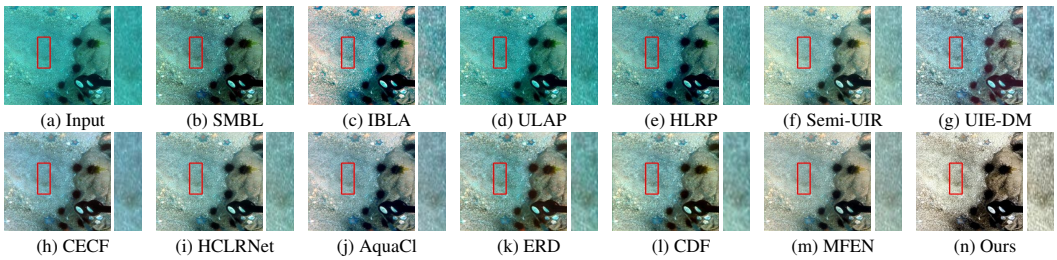

(a) Input (b) SMBL (c) IBLA (d) ULAP (e) HLRP (f) Semi-UIR (g) UIE-DM
(h) CECF (i) HCLRNet (j) AquaCl (k) ERD (l) CDF (m) MFEN (n) Ours

Figure 4: Visual comparisons with underwater image enhancement methods. The red box indicates the zoomed-in region, which is displayed on the right side of the image.

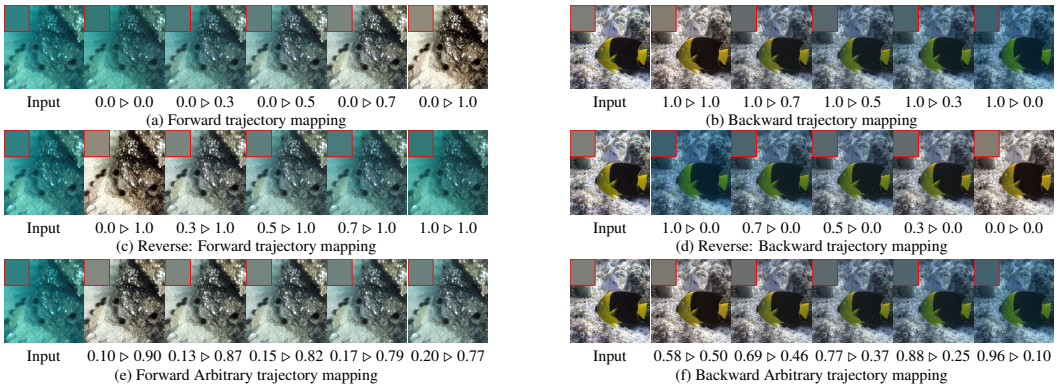

Figure 5: Visual results of omnidirectional trajectory mapping results.

Table 2: Evaluation results under different trajectory paths. Panels (c) and (d) reorganize the evaluation results from (a) and (b), respectively, without additional model inference.

(a) Forward trajectory mapping, from the distorted images to clean images.

$0:T$	0.0 : 0.0	0.0 : 0.1	0.0 : 0.2	0.0 : 0.3	0.0 : 0.4	0.0 : 0.5	0.0 : 0.6	0.0 : 0.7	0.0 : 0.8	0.0 : 0.9	0.0 : 1.0
UIQM	3.185	3.445	3.652	3.880	4.142	4.419	4.706	4.952	4.996	**5.061**	4.998
URANK	0.376	0.645	0.851	1.102	1.479	1.877	2.233	2.603	2.856	3.200	**3.371**

(b) Reverse: Forward trajectory mapping, from the clean images to distorted images.

$S:0$	1.0 : 1.0	1.0 : 0.9	1.0 : 0.8	1.0 : 0.7	1.0 : 0.6	1.0 : 0.5	1.0 : 0.4	1.0 : 0.3	1.0 : 0.2	1.0 : 0.1	1.0 : 0.0
UIQM	**5.634**	5.545	5.371	5.211	5.058	4.919	4.785	4.644	4.508	4.362	4.200
URANK	**3.101**	2.974	2.590	2.261	2.069	1.949	1.834	1.692	1.462	1.301	1.214

(c) Backward trajectory mapping from (a)

$S:T$	0.0 : 1.0	0.3 : 1.0	0.5 : 1.0	0.7 : 1.0	1.0 : 1.0
UIQM	**4.998**	4.952	4.419	3.880	3.185
URANK	**3.371**	2.603	1.877	1.102	0.376

(d) Reverse: Backward trajectory mapping from (b)

$S:T$	1.0 : 0.0	0.7 : 0.0	0.5 : 0.0	0.3 : 0.0	0.0 : 0.0
UIQM	4.200	4.644	4.919	5.211	**5.634**
URANK	1.214	1.692	1.949	2.261	**3.101**

(e) Arbitrary trajectory mapping

$S:T$	0.10:0.90	0.13:0.87	0.15:0.82	0.17:0.79	0.20:0.77
UIQM	**5.085**	5.026	4.884	4.765	4.629
URANK	**2.856**	2.715	2.486	2.301	2.138

(f) Reverse: Arbitrary trajectory mapping

$S:T$	0.58:0.50	0.69:0.46	0.77:0.37	0.88:0.25	0.96:0.10
UIQM	**5.572**	5.335	5.058	4.737	4.424
URANK	**3.027**	2.480	2.069	1.794	1.366

batch size of 1. The balancing hyper-parameters for the loss functions, λ_a and λ_p, are empirically set to 0.5 and 1.0, respectively. All experiments are implemented in PyTorch and conducted on a single NVIDIA GeForce RTX 4090 GPU.

4.1 Real-world Enhancement and Omnidirectional Trajectory Mapping Results

Visual and Quantitative Comparison with Enhancement Methods. Table 1 presents the quantitative evaluation results, which confirm that our method consistently attains robust and competitive performance across all metrics. Fig. 4 illustrates the visual comparisons between our OC-UIE and existing state-of-the-art methods. The results demonstrate that our proposed approach achieves superior performance in both color correction and structural recovery.

Visual Demonstration of Omnidirectional Control. Fig. 5 provides a qualitative demonstration of the omnidirectional control and geometric continuity. Subfigures (a) and (b) illustrate the model's capacity for unrestricted forward and backward evolution, maintaining rigorous spectral consistency at every interval. The global stability of the trajectory field is further validated in subfigures (c) and (d); regardless of the starting anchor point, the system consistently converges toward the designated boundary states (0.0 or 1.0), confirming the consistency of the learned mapping field. Most significantly, subfigures (e) and (f) showcase arbitrary point-to-point navigation between any two coordinates in the representation space (e.g., $0.10 \rightarrow 0.90$ or $0.96 \rightarrow 0.10$). This confirms that OC-UIE fundamentally transcends conventional rigid, point-to-point mappings by reformulating the underwater restoration task as an intrinsically traversable dynamical system defined over a continuous operator manifold.

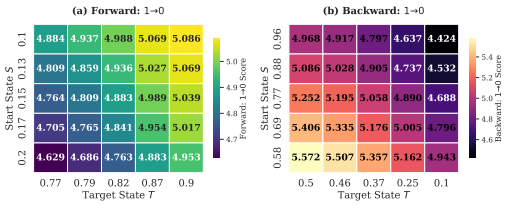

Figure 6: The UIQM heatmap of the any point to any point mapping $S \to T$.

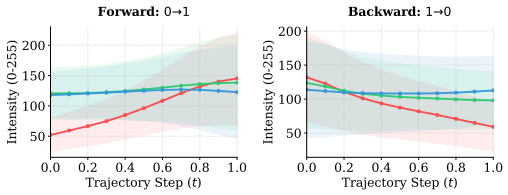

Figure 7: Variation of RGB color channel intensities during the different processes.

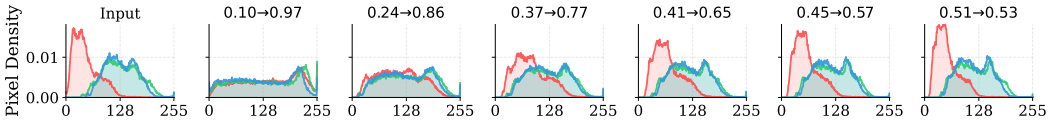

Figure 8: Histogram evolution of the arbitrary trajectory mapping $S \to T$.

Table 3: Ablation results of SI of URANKER.

Method	0→0.0	0→0.2	0→0.4	0→0.6	0→0.8	0→1.0
w/o SI	0.091	0.668	0.748	0.805	0.835	**0.867**
Ours	-1.647	-0.865	-0.492	-0.345	0.102	**0.867**

Table 4: Ablation results of MT.

Settings	UCCS			EUVP-330		
	UIQM	CIQA	URANK	UIQM	CIQA	URANK
w/o MT	4.380	0.460	1.908	4.176	**0.478**	1.957
Ours	**4.815**	**0.526**	**2.224**	**4.587**	0.475	**2.235**

Quantitative Validation of Omnidirectional Control. Table 2 provides a quantitative validation of the framework's geometric continuity and omnidirectional flexibility corresponding to Fig. 5. Across all experimental configurations, the scores demonstrate an overall positive correlation with the target parameter T: as T increases from 0.0 to 1.0, the image quality improves consistently, regardless of whether the trajectory is sampled at coarse or fine-grained intervals. Furthermore, the successful execution of the Reverse ($1.0 \to T$) and Arbitrary ($S \to T$) paths confirms that the OC-UIE paradigm does not merely learn a static enhancement mapping, but rather masters the underlying dynamics of the operator manifold, allowing for stable and predictable navigation between any two states in the degradation-restoration space. As illustrated in Fig. 6 that corresponding to the Fig. 5 (e) and (f), our model exhibits exceptional stability across arbitrary navigation paths. The consistent performance gain from various starting points S to target states T demonstrates that the learned operator manifold maintains a coherent mapping field, ensuring predictable restoration regardless of the initial degradation level.

Channel Evolution and Reversible Symmetry. As shown in Fig. 7 corresponding to the path in Fig. 5(a) and (b), the color channel evolution confirms the robustness of our omnidirectional learning. In the forward path ($0 \to 1$), the red channel exhibits a smooth, monotonic recovery from its initial attenuation, while green and blue channels remain stable. Conversely, the backward path ($1 \to 0$) accurately simulates degradation by gradually diminishing the red component of the clear image.

Pixel-level Distribution Dynamics. Fig. 8 illustrates the pixel density evolution, highlighting the framework's ability to precisely modulate color distributions. From Input to target states, the red channel's histogram progressively expands from a compressed, low-intensity range toward a normalized distribution, while blue and green channels maintain statistical stability. This progressive dispersion verifies that the framework selectively compensates for severe red-wavelength attenuation while avoiding over-saturation in well-preserved bands.

4.2 Ablation Study and Discussions

Ablation on Symmetric Interpolation: Symmetric interpolation (SI) in Eq. 6 is central to our trajectory decomposition. Its contribution is validated via an ablation study, denoted as w/o SI. As shown in Table 3, removing SI leads to a marked reduction in the range of image quality variation, indicating that SI extends the feature-space distance between degraded and restored states to provide superior controllability. These findings are consistent with the visual evidence in Fig. 9(a), where the contrast between different enhancement levels becomes less pronounced in w/o SI.

Ablation on Mixed-Training: Incorporating both synthetic and real-world data through hybrid training is essential for model robustness. We evaluate this strategy by denoting its omission as w/o MT. The results in Table 4 reveal a substantial drop in quantitative scores without MT. This is further supported by Fig. 9(b), where models trained exclusively on real-world data exhibit inferior visual restoration compared to the full training setup.

Sensitivity Analysis on Loss Weighting: Fig. 9(c) evaluates the model's sensitivity to loss weights ($\lambda_a : \lambda_p$). Across a range of scales (0.2 to 5.0), both UIQM and URanker scores fluctuate within a narrow margin, with no catastrophic performance degradation observed. The peak performance occurs near balanced configurations, confirming that our OC-UIE framework is robust to hyperparameter selection, effectively maintaining high restoration quality without requiring exhaustive tuning.

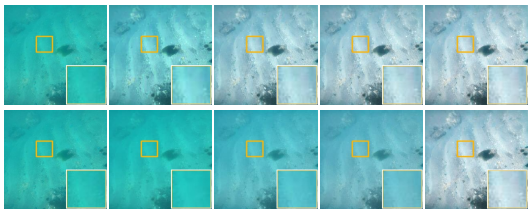

(a) Ablation of SI. The first and second row are w/o SI and Ours.

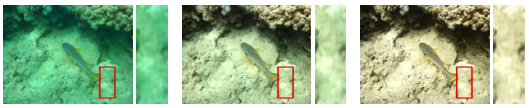

(b) Visual ablation study of MT.

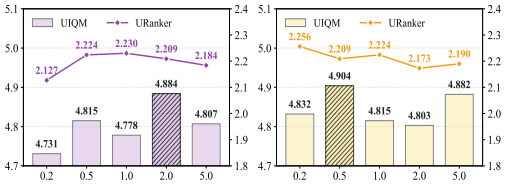

(c) Evaluation of the sensitivity of performance to the loss weights.

Figure 9: Visual ablation results and hyperparameters sensitivity evaluations.

5 Conclusions

In conclusion, we present OC-UIE, a paradigm shift from rigid unidirectional mapping to an unsupervised Omnidirectional Training method for underwater image enhancement. The main purpose of our work is the realization of a consistent omnidirectional mapping across the representation space, which allows the model to be trained beyond unidirectional constraints. By optimizing an omnidirectional objective and employing representation trajectory flattening and trajectory omnidirectional integration, our framework masters consistent state transitions that decouple structural preservation from degradation dynamics. By optimizing only the relative shifts, OC-UIE isolates structural preservation from fluid degradation dynamics. This approach transforms underwater restoration into a geometrically interpretable dynamical system, enabling intensity-parameterized control and state-of-the-art performance.

Acknowledgments and Disclosure of Funding

This work was supported by the National Natural Science Foundation of China (Grant No. 62676086). This work was also supported in part by the CIPS-SMP-Zhipu Large Model Fund. The authors declare no competing interests.

References

[1] Derya Akkaynak and Tali Treibitz. A revised underwater image formation model. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6723–6732, 2018.

[2] Hengyue Bi, Long Chen, Jingchao Cao, Jingyang Wang, Jinghao Sun, Yuan Rao, and Junyu Dong. Seadiff: Underwater image enhancement with degradation-aware diffusion model. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.

[3] Xuewen Bing, Wenqi Ren, Yang Tang, Gary G Yen, and Qiyu Sun. Domain adaptation for in-air to underwater image enhancement via deep learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2023.

[4] Jiangzhong Cao, Zekai Zeng, Xu Zhang, Huan Zhang, Chunling Fan, Gangyi Jiang, and Weisi Lin. Unveiling the underwater world: Clip perception model-guided underwater image enhancement. *Pattern Recognition*, 162:111395, 2025.

[5] Jingchao Cao, Wangzhen Peng, Yutao Liu, Junyu Dong, Patrick Le Callet, and Sam Kwong. Erd: Encoder-residual-decoder neural network for underwater image enhancement. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.

[6] Nicholas Carlevaris-Bianco, Anush Mohan, and Ryan M Eustice. Initial results in underwater single image dehazing. In *Oceans 2010 Mts/IEEE Seattle*, pages 1–8, 2010.

[7] Ashutosh Chauhan, Dakshi Goel, Meghna Kapoor, Badri Narayan Subudhi, and Vinit Jakhetiya. Illumination aware multi-scale attention fusion model for underwater image enhancement. *IEEE Transactions on Artificial Intelligence*, 2025.

[8] Xuhang Chen, Zimeng Li, Xiujun Zhang, Yingtie Lei, Yihang Dong, Xiaofeng Zhang, and Chi-Man Pun. Aquaclarity: Vision transformer enhances underwater image via nonlinear feature fusion and color fidelity improvement. *IEEE Transactions on Consumer Electronics*, 2025.

[9] Yu-Wei Chen and Soo-Chang Pei. Domain adaptation for underwater image enhancement via content and style separation. *IEEE Access*, 10:90523–90534, 2022.

[10] Xiaofeng Cong, Jie Gui, and Junming Hou. Underwater organism color fine-tuning via decomposition and guidance. In *AAAI Conference on Artificial Intelligence*, volume 38, pages 1389–1398, 2024.

[11] Guodong Fan, Yu Zhou, Jingchun Zhou, Yakun Ju, Guang-Yong Chen, Jinjiang Li, and Alex C Kot. Dcd-uie: Decoupled chromatic diffusion model for underwater image enhancement. *IEEE Transactions on Image Processing*, 2026.

[12] Chunle Guo, Ruiqi Wu, Xin Jin, Linghao Han, Weidong Zhang, Zhi Chai, and Chongyi Li. Underwater ranker: Learn which is better and how to be better. In *AAAI Conference on Artificial Intelligence*, pages 702–709, 2023.

[13] Guojia Hou, Nan Li, Peixian Zhuang, Kunqian Li, Haihan Sun, and Chongyi Li. Non-uniform illumination underwater image restoration via illumination channel sparsity prior. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(2):799–814, 2024.

[14] Shirui Huang, Keyan Wang, Huan Liu, Jun Chen, and Yunsong Li. Contrastive semi-supervised learning for underwater image restoration via reliable bank. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 18145–18155, 2023.

[15] Md Jahidul Islam, Youya Xia, and Junaed Sattar. Fast underwater image enhancement for improved visual perception. *IEEE Robotics and Automation Letters*, 5(2):3227–3234, 2020.

[16] Qiuping Jiang, Yaozu Kang, Zhihua Wang, Wenqi Ren, and Chongyi Li. Perception-driven deep underwater image enhancement without paired supervision. *IEEE Transactions on Multimedia*, 2023.

[17] Qun Jiang, Yunfeng Zhang, Fangxun Bao, Xiuyang Zhao, Caiming Zhang, and Peide Liu. Two-step domain adaptation for underwater image enhancement. *Pattern Recognition*, 122: 108324, 2022.

[18] Guisik Kim, Sung Woo Park, and Junseok Kwon. Pixel-wise wasserstein autoencoder for highly generative dehazing. *IEEE Transactions on Image Processing*, 30:5452–5462, 2021.

[19] Chongyi Li, Chunle Guo, Wenqi Ren, Runmin Cong, Junhui Hou, Sam Kwong, and Dacheng Tao. An underwater image enhancement benchmark dataset and beyond. *IEEE Transactions on Image Processing*, 29:4376–4389, 2019.

[20] Hanyu Li, Jingjing Li, and Wei Wang. A fusion adversarial underwater image enhancement network with a public test dataset. *arXiv preprint arXiv:1906.06819*, 2019.

[21] Weiming Li, Xuelong Wu, Shuaishuai Fan, Songjie Wei, and Glyn Gowing. Ingc-gan: An implicit neural-guided cycle generative approach for perceptual-friendly underwater image enhancement. *IEEE Transactions on Neural Networks and Learning Systems*, 2025.

[22] Qiong Liu, Qi Zhang, Wei Liu, Wenbai Chen, Xinwang Liu, and Xiangke Wang. Wsds-gan: A weak-strong dual supervised learning method for underwater image enhancement. *Pattern Recognition*, page 109774, 2023.

[23] Risheng Liu, Xin Fan, Ming Zhu, Minjun Hou, and Zhongxuan Luo. Real-world underwater enhancement: Challenges, benchmarks, and solutions under natural light. *IEEE transactions on circuits and systems for video technology*, 30(12):4861–4875, 2020.

[24] Xiaokai Liu, Yutong Jiang, Yangyang Wang, Taifei Liu, and Jie Wang. Mda-net: A multidistribution aware network for underwater image enhancement. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–13, 2025.

[25] Pan Mu, Hanning Xu, Zheyuan Liu, Zheng Wang, Sixian Chan, and Cong Bai. A generalized physical-knowledge-guided dynamic model for underwater image enhancement. In *ACM International Conference on Multimedia*, pages 7111–7120, 2023.

[26] Ankita Naik, Apurva Swarnakar, and Kartik Mittal. Shallow-uwnet: Compressed model for underwater image enhancement (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 15853–15854, 2021.

[27] Karen Panetta, Chen Gao, and Sos Agaian. Human-visual-system-inspired underwater image quality measures. *IEEE Journal of Oceanic Engineering*, 41(3):541–551, 2015.

[28] Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. One-step image translation with text-to-image models. *arXiv preprint arXiv:2403.12036*, 2024.

[29] Lintao Peng and Liheng Bian. Adaptive dual-domain learning for underwater image enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 6461–6469, 2025.

[30] Lintao Peng, Chunli Zhu, and Liheng Bian. U-shape transformer for underwater image enhancement. *IEEE Transactions on Image Processing*, 2023.

[31] Yan-Tsung Peng and Pamela C Cosman. Underwater image restoration based on image blurriness and light absorption. *IEEE transactions on image processing*, 26(4):1579–1594, 2017.

[32] Hao Qi, Huiyu Zhou, Junyu Dong, and Xinghui Dong. Deep color-corrected multi-scale retinex network for underwater image enhancement. *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

[33] Yuan Rao, Wenjie Liu, Kunqian Li, Hao Fan, Sen Wang, and Junyu Dong. Deep color compensation for generalized underwater image enhancement. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.

[34] Jiashuo Shang, Ying Li, Hu Xing, and Jingyi Yuan. Lgt: Luminance-guided transformer-based multi-feature fusion network for underwater image enhancement. *Information Fusion*, 118: 102977, 2025.

[35] Jingyu Song, Haiyong Xu, Gangyi Jiang, Mei Yu, Yeyao Chen, Ting Luo, and Yang Song. Frequency domain-based latent diffusion model for underwater image enhancement. *Pattern Recognition*, 160:111198, 2025.

[36] Wei Song, Yan Wang, Dongmei Huang, and Dian Tjondronegoro. A rapid scene depth estimation model based on underwater light attenuation prior for underwater image restoration. In *Advances in Multimedia Information Processing Pacific-Rim Conference on Multimedia*, pages 678–688, 2018.

[37] Wei Song, Yan Wang, Dongmei Huang, Antonio Liotta, and Cristian Perra. Enhancement of underwater images with statistical model of background light and optimization of transmission map. *IEEE Transactions on Broadcasting*, 66(1):153–169, 2020.

[38] Zijun Tan, Chuan Fu, Tan Guo, Zhixiong Nan, Pengzhan Zhou, Xinggan Peng, and Fulin Luo. Pgmamba: A physical model-guided global mamba for underwater image enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 9359–9367, 2026.

[39] Yi Tang, Takafumi Iwaguchi, Hiroshi Kawasaki, Ryusuke Sagawa, and Ryo Furukawa. Autoenhancer: Transformer on u-net architecture search for underwater image enhancement. In *Proceedings of the Asian Conference on Computer Vision*, pages 1403–1420, 2022.

[40] Yi Tang, Hiroshi Kawasaki, and Takafumi Iwaguchi. Underwater image enhancement by transformer-based diffusion model with non-uniform sampling for skip strategy. In *ACM International Conference on Multimedia*, pages 5419–5427, 2023.

[41] Hao Wang, Weibo Zhang, Lu Bai, and Peng Ren. Metalantis: A comprehensive underwater image enhancement framework. *IEEE Transactions on Geoscience and Remote Sensing*, 2024.

[42] Hao Wang, Weibo Zhang, Yinghao Xu, Huanyu Li, and Peng Ren. Watercyclediffusion: Visual-textual fusion empowered underwater image enhancement. *Information Fusion*, page 103693, 2025.

[43] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 2555–2563, 2023.

[44] Zhengyong Wang, Liquan Shen, Mai Xu, Mei Yu, Kun Wang, and Yufei Lin. Domain adaptation for underwater image enhancement. *IEEE Transactions on Image Processing*, 32:1442–1457, 2023.

[45] Zhiheng Wu, Zhengxing Wu, Xingyu Chen, Yue Lu, and Junzhi Yu. Self-supervised underwater image generation for underwater domain pre-training. *IEEE Transactions on Instrumentation and Measurement*, 2024.

[46] Yaofeng Xie, Lingwei Kong, Kai Chen, Ziqiang Zheng, Xiao Yu, Zhibin Yu, and Bing Zheng. Uveb: A large-scale benchmark and baseline towards real-world underwater video enhancement. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2024.

[47] Chengpei Xu, Wenhao Zhou, Zhixiong Huang, Yuanfang Zhang, Yan Zhang, Weimin Wang, and Feng Xia. Fusion-based graph neural networks for synergistic underwater image enhancement. *Information Fusion*, 117:102857, 2025.

[48] Xinwei Xue, Zexuan Li, Long Ma, Qi Jia, Risheng Liu, and Xin Fan. Investigating intrinsic degradation factors by multi-branch aggregation for real-world underwater image enhancement. *Pattern Recognition*, 133:109041, 2023.

[49] Haorui Yan, Zhenwei Zhang, Jing Xu, Tingting Wang, Ping An, Aobo Wang, and Yuping Duan. Uw-cyclegan: Model-driven cyclegan for underwater image restoration. *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

[50] Yaming Yang, Dilxat Muhtar, Yelong Shen, Yuefeng Zhan, Jianfeng Liu, Yujing Wang, Hao Sun, Weiwei Deng, Feng Sun, Qi Zhang, et al. Mtl-lora: Low-rank adaptation for multi-task learning. In *AAAI Conference on Artificial Intelligence*, volume 39, pages 22010–22018, 2025.

[51] Meng Yu, Liquan Shen, Zhengyong Wang, and Xia Hua. Task-friendly underwater image enhancement for machine vision applications. *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

[52] Haopeng Zhang, Hongli Xu, Xiaosheng Yu, Xiangyue Zhang, Xiujing Gao, and Chengdong Wu. Cdf-uie: Leveraging cross-domain fusion for underwater image enhancement. *IEEE Transactions on Geoscience and Remote Sensing*, 2025.

[53] Weidong Zhang, Peixian Zhuang, Hai-Han Sun, Guohou Li, Sam Kwong, and Chongyi Li. Underwater image enhancement via minimal color loss and locally adaptive contrast enhancement. *IEEE Transactions on Image Processing*, 31:3997–4010, 2022.

[54] Weidong Zhang, Qingmin Liu, Huimin Lu, Jianping Wang, and Jing Liang. Underwater image enhancement via wavelet decomposition fusion of advantage contrast. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.

[55] Weidong Zhang, Muzi Wang, Peixian Zhuang, and Dahai Liu. Underwater image enhancement via advantage feature weighted fusion. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.

[56] Yibo Zhang, Jingchun Zhou, Ruijie Liu, Dehuan Zhang, and Qiuping Jiang. Perception-consistent underwater image enhancement for sustainable agricultural visual sensing. *IEEE Transactions on Consumer Electronics*, 2025.

[57] Zengxi Zhang, Zhiying Jiang, Jinyuan Liu, Xin Fan, and Risheng Liu. Waterflow: heuristic normalizing flow for underwater image enhancement and beyond. In *ACM International Conference on Multimedia*, pages 7314–7323, 2023.

[58] Chen Zhao, Weiling Cai, Chenyu Dong, and Chengwei Hu. Wavelet-based fourier information interaction with frequency diffusion adjustment for underwater image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8281–8291, 2024.

[59] Jingchun Zhou, Qilin Gai, Dehuan Zhang, Kin-Man Lam, Weishi Zhang, and Xianping Fu. Iacc: Cross-illumination awareness and color correction for underwater images under mixed natural and artificial lighting. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–15, 2024.

[60] Jingchun Zhou, Jiaming Sun, Chongyi Li, Qiuping Jiang, Man Zhou, Kin-Man Lam, Weishi Zhang, and Xianping Fu. Hclr-net: Hybrid contrastive learning regularization with locally randomized perturbation for underwater image enhancement. *International Journal of Computer Vision*, pages 1–25, 2024.

[61] Jingchun Zhou, Wenbo Guo, Dehuan Zhang, Guangyong Chen, Qiuping Jiang, and Ghulam Muhammad. Sp-uiqa: Fusion segmentation priors and multi-scale features for underwater image quality assessment. *Information Fusion*, page 104257, 2026.

[62] Jingchun Zhou, Chunjiang Liu, Qiuping Jiang, Xianping Fu, Junhui Hou, and Xuelong Li. Semantic contrast for domain-robust underwater image quality assessment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

[63] Jingchun Zhou, Jiongjiong Wang, Bing Long, Dehuan Zhang, and Qiuping Jiang. Underwater image stitching via optimal seam estimation and multi-band fusion. *Pattern Recognition*, page 113459, 2026.

[64] Peixian Zhuang, Jiamin Wu, Fatih Porikli, and Chongyi Li. Underwater image enhancement with hyper-laplacian reflectance priors. *IEEE Transactions on Image Processing*, 31:5442–5455, 2022.

A Supplementary Materials

Analysis of Prompts. Quantitative evaluation across ten radically different semantic contexts in Fig. 10 demonstrates the exceptional robustness of our method, where the left is for the input distorted scene $0.1 \rightarrow 0.9$ in Fig. 5(a), and the right is for the input clean scene $0.9 \rightarrow 0.1$ in Fig. 5(b). When comparing the neutral baseline "a picture with content" against extreme environmental interference such as "a vast Saharan desert" or "futuristic cyberpunk city", the non-reference metrics including UIQM and URanker exhibit remarkable stability. This numerical convergence confirms that the restoration process is governed by the underlying physical priors of the underwater scene and remains effectively immune to the stochastic interference of textual prompts.

More Real-world Results. The analysis of RGB channel evolution in Fig. 11 and spectral dynamics in Fig. 12, which all corresponding to Fig. 13 (a) and (b), further confirms the framework's ability to maintain reversible symmetry and structural regularity across the entire operator manifold as concluded in the main paper. We provide extended results of omnidirectional control in Fig. 14, and Table 5 which corresponding to Fig. 13. It yield conclusions consistent with findings in main paper. Fig. 15 presents additional comparisons on real-world scenarios, which further demonstrates that the proposed model achieves superior overall performance in color correction and image restoration.

Inference Speed. Table 6 demonstrates that our proposed model achieves a highly competitive balance between restoration quality and computational efficiency, enabling rapid real-time inference.

Limitations. This work focuses on continuous degradation modeling for underwater imagery. While our framework may be applicable to tasks like dehazing or deraining, its effectiveness across these distinct physical manifolds remains to be extensively validated.

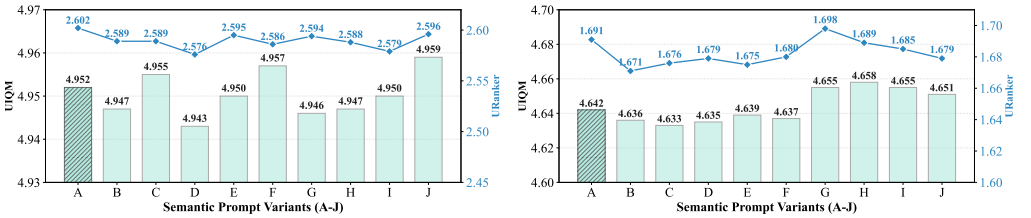

Figure 10: Sensitivity of performance to the prompts of forward (left) and backward (right).

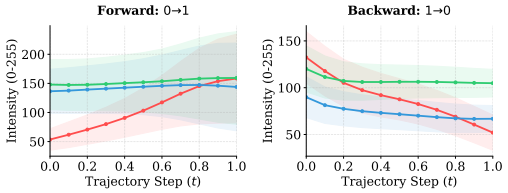

Figure 11: Results of RGB color channel intensities during the different processes.

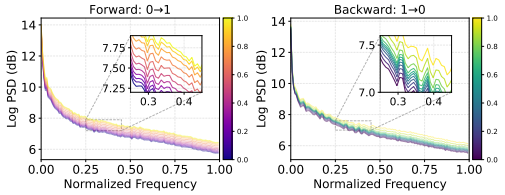

Figure 12: Analysis of spectral image restoration and degradation trajectories.

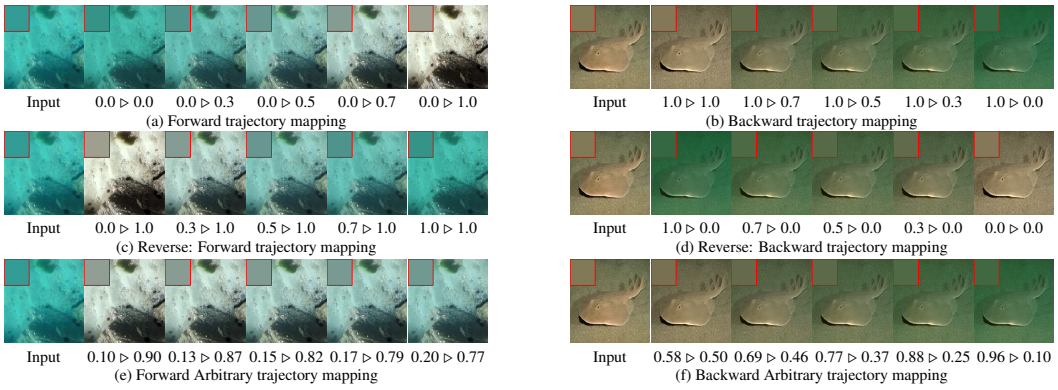

Figure 13: Visual results of omnidirectional trajectory mapping results.

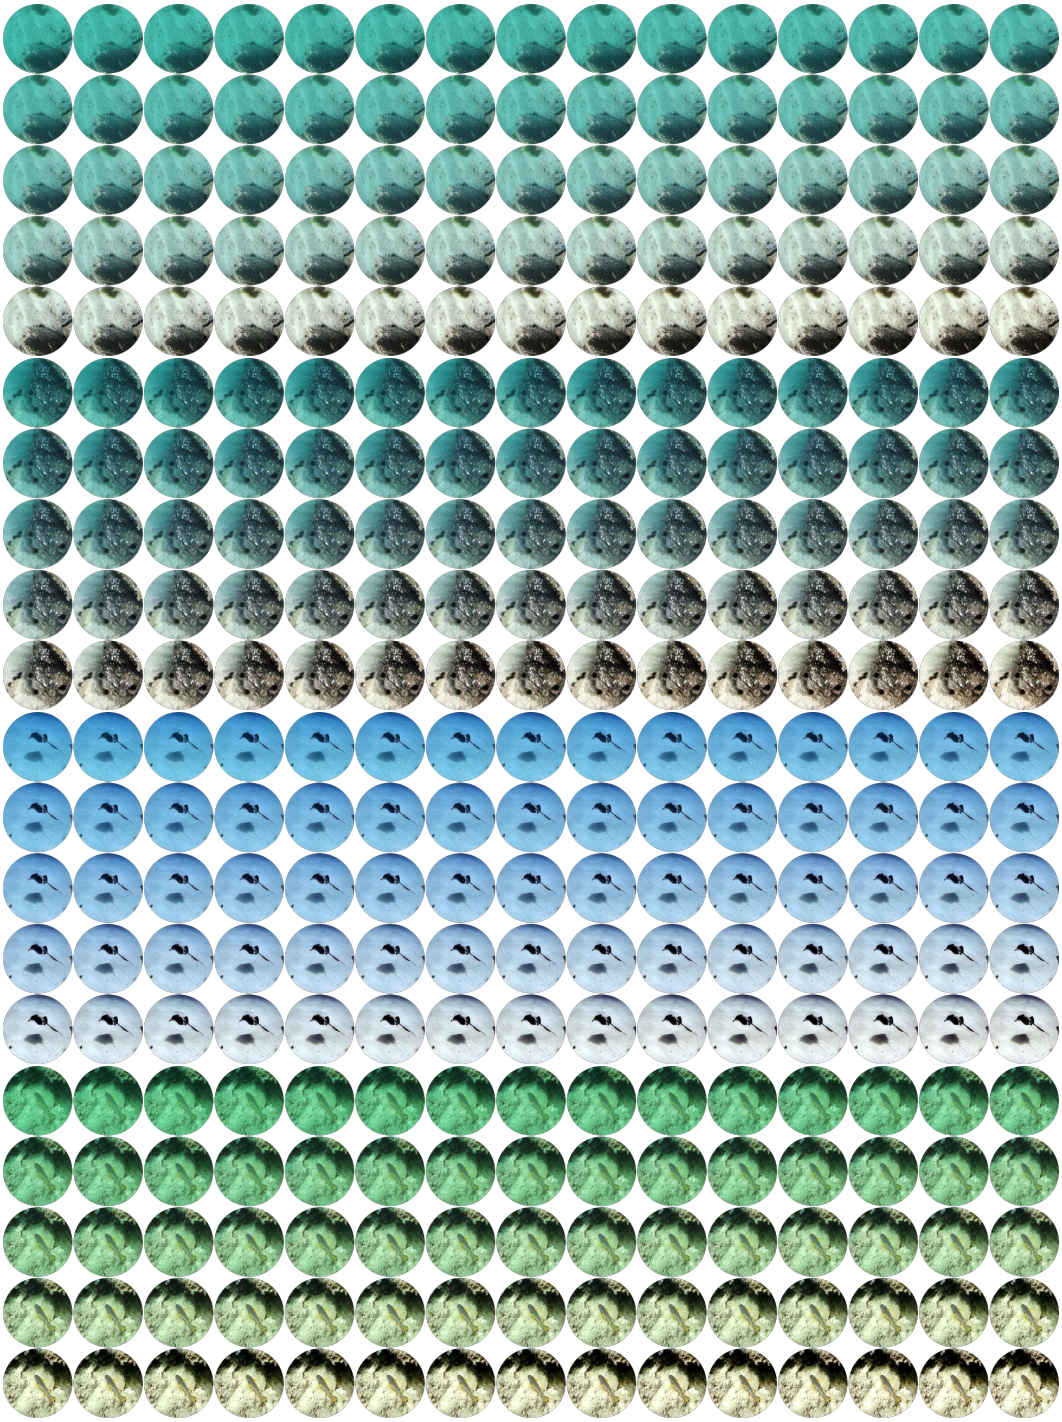

Figure 14: Visualizing the continuous degradation manifold. This figure displays 75 states uniformly sampled from the learned representation trajectory (15×5 grid).

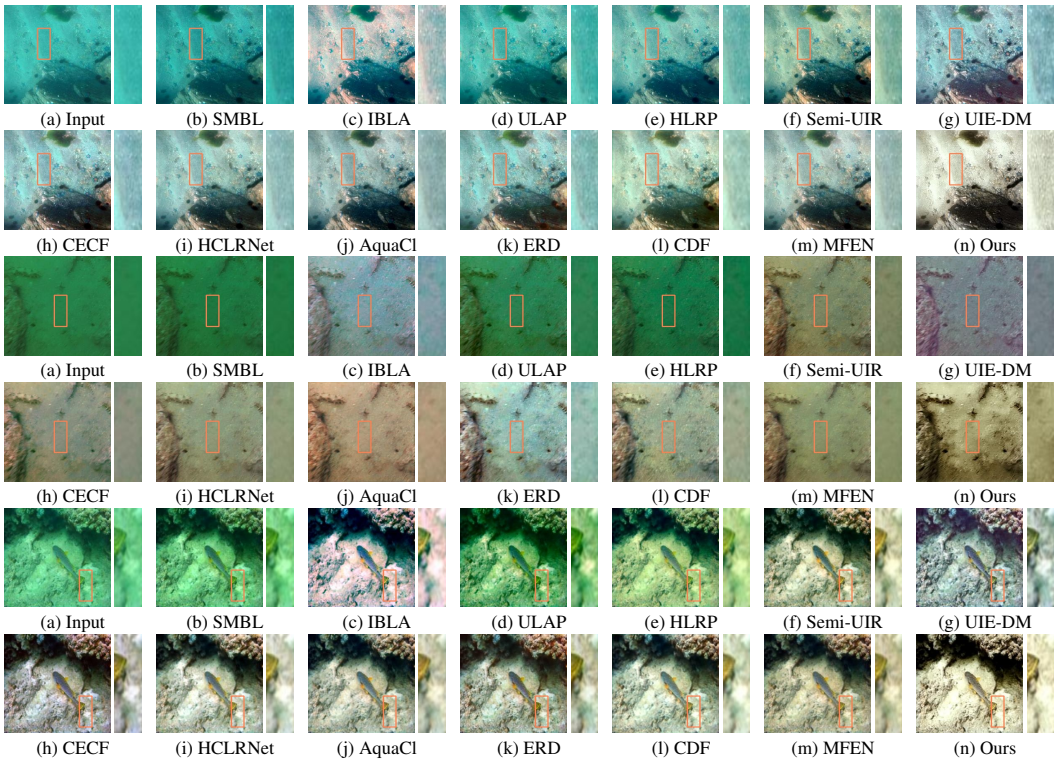

Figure 15: Visual comparisons with underwater image enhancement methods.

Table 5: Evaluation results under different trajectory paths.

(a) Forward trajectory mapping, from the distorted images to clean images.

$0:T$	0.0 : 0.0	0.0 : 0.1	0.0 : 0.2	0.0 : 0.3	0.0 : 0.4	0.0 : 0.5	0.0 : 0.6	0.0 : 0.7	0.0 : 0.8	0.0 : 0.9	0.0 : 1.0
UIQM	1.803	2.069	2.323	2.582	2.869	3.257	3.685	4.136	4.364	4.735	**4.820**
URANK	-0.522	-0.317	-0.078	0.195	0.400	0.610	1.029	1.420	1.622	1.927	**2.139**

(b) Reverse: Forward trajectory mapping, from the clean images to distorted images.

$S:0$	1.0 : 1.0	1.0 : 0.9	1.0 : 0.8	1.0 : 0.7	1.0 : 0.6	1.0 : 0.5	1.0 : 0.4	1.0 : 0.3	1.0 : 0.2	1.0 : 0.1	1.0 : 0.0
UIQM	**5.114**	4.853	4.515	4.287	4.090	3.938	3.777	3.614	3.408	3.200	2.978
URANK	**0.193**	-0.075	-0.336	-0.547	-0.717	-0.862	-1.012	-1.174	-1.329	-1.464	-1.587

(c) Backward trajectory mapping from (a)

$S:T$	0.0 : 1.0	0.3 : 1.0	0.5 : 1.0	0.7 : 1.0	1.0 : 1.0
UIQM	**4.820**	4.136	3.257	2.582	1.803
URANK	**2.139**	1.420	0.610	0.195	-0.522

(d) Reverse: Backward trajectory mapping from (b)

1.0 : 0.0	0.7 : 0.0	0.5 : 0.0	0.3 : 0.0	0.0 : 0.0
2.978	3.614	3.938	4.287	**5.114**
-1.587	-1.174	-0.862	-0.547	**0.193**

(e) Arbitrary trajectory mapping

$S:T$	0.10:0.90	0.13:0.87	0.15:0.82	0.17:0.79	0.20:0.77
UIQM	**4.363**	4.289	4.016	3.775	3.549
URANK	**1.622**	1.487	1.349	1.142	0.864

(f) Reverse: Arbitrary trajectory mapping

0.58:0.50	0.69:0.46	0.77:0.37	0.88:0.25	0.96:0.10
4.892	4.440	4.091	3.739	3.278
-0.013	-0.405	-0.716	-1.059	-1.410

Table 6: Inference speeds of different methods.

Method	MIP	SMBL	IBLA	ULAP	HCLR	UIE-DM	AquaCl	CDF	MFEN	Ours
Time (ms)	1542	306.3	391.8	549.2	83.5	205.4	98.62	42.3	34.6	140.2

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Sufficient experimental results demonstrate the effectiveness and superiority of the proposed method.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The limitations of this approach are further elaborated in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: We validate our hypothesis through both empirical analysis and theoretical derivation.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide comprehensive experimental details in this manuscript and will release the source code upon acceptance.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All datasets used in this study are publicly available, and the experimental code will be released to the public

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Comprehensive implementation details are provided in this manuscript.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: All experimental results are statistically significant and documented to ensure full reproducibility.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The details regarding computational resources used for the experiments are provided in the manuscript.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: All code is implemented based on open-source frameworks and complies with the standard requirements.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: The proposed algorithm is purely an underwater image enhancement model and poses no negative ethical implications.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.

- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: There are no ethical risks associated with the model and its related resources.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All resources used in this paper, including datasets and pre-trained models, are publicly open-sourced.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: No additional dedicated resources are involved in this work, and thus this concern is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: We only propose an image processing algorithm and involve no human experiments in this work.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: We only propose an underwater image processing algorithm and involve no human experiments in this work.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: We adopt an LLM to polish the paper's grammar, check spelling, and refine academic expressions.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.