SEA-PACE: Semi-Supervised Underwater Image Enhancement via Gaussian Process–Assisted Self-Paced Learning

Jingyang Wang[1*], Hengyue Bi[1*], Jingchao Cao[1†], Feng Gao[1], Junyu Dong[1]

[1]Faculty of Information Science and Engineering, Ocean University of China
wangjingyang6966@stu.ouc.edu.cn, bihengyue@stu.ouc.edu.cn, caojingchao@ouc.edu.cn, gaofeng@ouc.edu.cn, dongjunyu@ouc.edu.cn

Abstract

The scarcity of paired data severely limits the performance and generalization of learning-based underwater image enhancement (UIE) methods. This challenge is particularly prominent in scenes with complex degradations. Semi-supervised learning has emerged as a promising solution by enabling the utilization of large-scale unlabeled data. However, its effectiveness is limited by the use of static, model-agnostic metrics for pseudo-label reliability assessment. To address this, we propose SEA-PACE, a novel semi-supervised framework that integrates model-aware uncertainty modeling and self-paced consistency learning to fully exploit unlabeled data for UIE. Specifically, we design a Model-Aware Reliability Estimator (MARE) that quantifies the uncertainty of the teacher model's predictions through Gaussian Process Regression in latent feature space. The resulting uncertainty is then transformed into reliability weights via a rank-based mapping. Additionally, we apply the Self-Paced Consistency Learning (SPCL) strategy that employs a loss-aware schedule to dynamically prioritize high-confidence pseudo-labels, gradually incorporating more challenging samples during training. Extensive experiments on several public UIE benchmarks demonstrate that SEA-PACE consistently surpasses state-of-the-art methods in both visual quality and generalization capability.

Code — https://github.com/haigouboss/SEA-PACE

Figure 1: An illustration of pseudo-labels generated by SemiUIR (Huang et al. 2023) and our proposed SEA-PACE. Compared with the input, pseudo-labels from SemiUIR suffer from over-enhancement and semantic inconsistency, exhibiting low perceptual quality and compromising consistency learning. In contrast, SEA-PACE enables model-aware selection and adaptive updating of pseudo-labels, ensuring better structural fidelity and more reliable supervision.

Introduction

Underwater image enhancement (UIE) is essential for underwater vision systems, aiming to recover high-quality images from inputs degraded by wavelength-dependent absorption and scattering (Chiang and Chen 2011). Over the past few decades, numerous methods have been developed to tackle the challenges posed by complex degradations, such as color distortion, low contrast, and blurred details, in real-world underwater scenarios. Notably, recent learning-based approaches (Peng, Zhu, and Bian 2023; Zhou et al. 2024) have achieved remarkable progress in UIE.

However, the impressive gains of these methods rest on an assumption that is rarely satisfied in practice: the availability of abundant paired data. In reality, such data are virtually unattainable because the underwater imaging process is irreversible, which precludes the capture of pristine reference scenes. To mitigate this bottleneck, semi-supervised frameworks, particularly Mean-Teacher variants (Huang et al. 2023; Chen et al. 2024; Xu et al. 2024), leverage large repositories of unlabeled images through teacher-student consistency learning. Nevertheless, their effectiveness remains constrained because pseudo-label selection depends exclusively on model-agnostic, static image quality assessment (IQA) (Cao et al. 2021) metrics such as MUSIQ (Ke et al. 2021), which introduces inherent bias into the learning process, and often leads to the generation of visual artifacts. As illustrated in Figure 1, these metrics often produce pseudo-labels of low perceptual quality with content inconsistencies, especially in scenes with complex degradations. Such pseudo-labels not only serve as unreliable proxies for

[*]These authors contributed equally.
[†]Corresponding author.

model-aware reliability quantification but also mislead the teacher-student consistency learning process. Furthermore, current methods apply consistency regularization uniformly to all unlabeled samples, disregarding their individual reliability and thereby compromising both training efficiency and convergence stability.

To address these limitations, we propose SEA-PACE, a semi-supervised UIE framework incorporating Gaussian Process-assisted Self-Paced Learning to enable more adaptive, model-aware use of unlabeled data. Specifically, our framework introduces a novel Model-Aware Reliability Estimator (MARE) that estimates sample-wise reliability in the latent feature space. Concretely, MARE first leverages Gaussian Process Regression (GPR) (Rasmussen 2003) to quantify the epistemic uncertainty of each unlabeled sample by modeling its latent distribution conditioned on features from the labeled data. This uncertainty is then transformed into a relative reliability weight via a rank-based transformation. Such estimation captures the model's confidence in transferring supervised UIE knowledge to unlabeled data, effectively bridging the learned representations and pseudo-label selection.

Building upon MARE, we further propose a Self-Paced Consistency Learning (SPCL) strategy that fosters a flexible and adaptive learning process within the Mean Teacher framework. Specifically, SPCL dynamically adjusts the training contribution of each unlabeled sample according to its estimated reliability, enabling the model to initially focus on the most reliable samples. Gradually, more challenging samples are incorporated and persistently unreliable ones are progressively down-weighted or excluded, thereby reducing confirmation bias (Arazo et al. 2020) and enhancing training stability.

Overall, our contributions are summarized as follows:

- We propose **SEA-PACE**, a semi-supervised framework for underwater image enhancement that integrates unlabeled data through reliability-guided consistency learning, effectively addressing the issue of data scarcity.
- We design a **Model-Aware Reliability Estimator** that employs Gaussian Process Regression to estimate epistemic uncertainty and produce sample-wise reliability weights in the latent feature space.
- We develop a **Self-Paced Consistency Learning** strategy that dynamically adjusts the learning pace for each unlabeled sample based on its reliability, thereby promoting stable and progressive training.

Related Work

Supervised UIE

Supervised UIE methods typically learn a direct mapping from degraded inputs to their high-quality counterparts using paired datasets (Cong et al. 2024). A primary focus of these methods is designing network architectures capable of learning robust feature representations, often incorporating building blocks like CNNs (Zhang et al. 2024; Bi et al. 2025), attention mechanisms (Cao et al. 2025), Transformers (Peng, Zhu, and Bian 2023), wavelet decomposition (Jamadandi and Mudenagudi 2019), and Fourier operations (Peng and Bian 2025). Beyond architectural innovations, another line of research concentrates on specialized learning strategies to better combat the complex degradations in the underwater domain. These include techniques such as adversarial learning (Cong et al. 2023), transfer learning (Wang et al. 2023), rank learning (Guo et al. 2023), and contrastive learning (Zhou et al. 2024). Despite their success with well-designed network architectures and learning strategy, supervised UIE methods are fundamentally limited by the availability and diversity of high-quality paired training data, which constrains their generalization in real-world underwater scenarios.

Unsupervised UIE

To alleviate the data-hunger issue, unsupervised techniques of UIE have been explored. These methods enable model training on unpaired data by either introducing physical priors (Fu et al. 2022; Yan et al. 2023) or employing cycle-consistency constraint (Yin et al. 2024). For example, HybrUR (Yan et al. 2023) integrates the Jaffe-McGlamery model into a GAN to simulate underwater degradation and enforce physics-based constraints. Yin et al. (Yin et al. 2024) adopt a cycle-consistent framework with bidirectional disentanglement and a double-order contrastive loss for better domain alignment. However, their performance is often hindered by unstable training dynamics and highly scene-dependent enhancement outcomes, primarily due to the complete absence of ground-truth supervision and the challenge of formulating sufficiently robust constraints to regularize the inherently ill-posed nature of the UIE task.

Semi-Supervised UIE

Conceptually acting as a bridge between supervised and unsupervised learning, semi-supervised learning focus on how to leverage both paired and unpaired data to improve the model's generalization performance. Existing semi-supervised UIE methods mainly adopt the teacher-student scheme within the Mean Teacher (Tarvainen and Valpola 2017) framework. In this framework, the role of the teacher model is to generate pseudo-labels for the unlabeled data. The student model is trained on both paired data and pseudo-labeled unpaired data. For instance, SemiUIR (Huang et al. 2023) employs an IQA metric MUSIQ to score pseudo-labels and maintains a reliable bank to store the best-ever pseudo-label for each sample, while SemiUIE (Xu et al. 2024) and UWFormer (Chen et al. 2024) rely on URanker or handcrafted perceptual losses. However, these approaches share two key limitations: (1) relying solely on external metrics without explicitly modeling pseudo-label reliability; (2) treating all unlabeled samples equally, ignoring the varying reliability and influence of pseudo-labels. In contrast, our method introduces a model-aware reliability estimator that models and quantifies the reliability of each pseudo-label, and further incorporate a self-paced consistency learning strategy that adaptively selects samples of varying reliability across different training stages, better aligning with the model's evolving learning capacity.

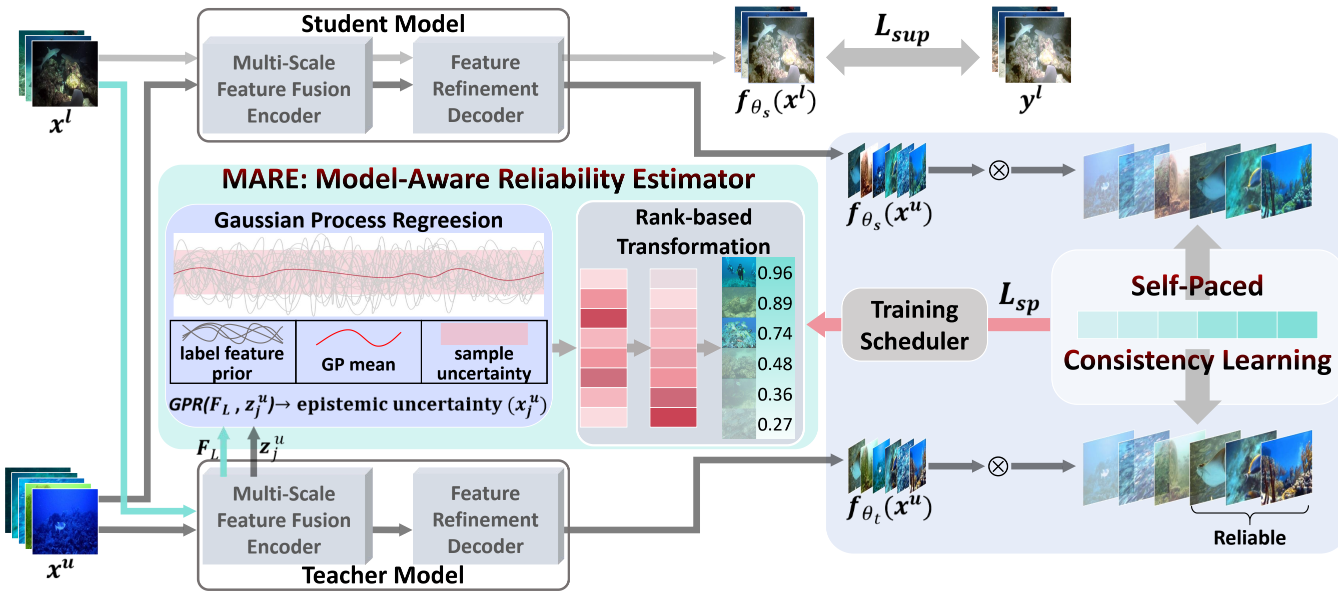

Figure 2: SEA-PACE is a semi-supervised UIE framework featuring a Model-Aware Reliability Estimator (MARE) and a Self-Paced Consistency Learning (SPCL) strategy. Our framework adopts a student-teacher architecture with identical networks built from a Multi-Scale Feature Fusion Encoder and a Feature Refinement Decoder. MARE leverages GPR to estimate epistemic uncertainty and assign rank-based reliability weights to pseudo-labels. SPCL leverages the resulting weights to modulate the consistency loss, while a training scheduler adjusts the learning pace accordingly.

Proposed Method

Task Definition and Notation

Data Setup. Let $\mathcal{D}_L = \{(x_i^l, y_i^l)\}_{i=1}^N$ be a small set of **labeled** underwater image pairs, where x_i^l is a degraded input and y_i^l its high-quality reference. Let $\mathcal{D}_U = \{x_j^u\}_{j=1}^M$ ($M \gg N$) be a large **unlabeled** collection drawn from the same domain, with $\mathcal{D}_L \cap \mathcal{D}_U = \emptyset$. The complete training pool is $\mathcal{D} = \mathcal{D}_L \cup \mathcal{D}_U$.

Model Components:

- a *student* model $f_{\theta_s} : \mathbb{R}^{3\times H\times W} \to \mathbb{R}^{3\times H\times W}$,
- a *teacher* model f_{θ_t} obtained as an exponential moving average (EMA) of the student (Eq. (4)),
- a *Model-Aware Reliability Estimator* that yields a reliability weight $\omega_j \in (0,1)$ for every unlabeled sample x_j^u.

Learning Objective. The student model's parameters θ_s are learned by minimizing a supervised–unsupervised composite loss:

$$\mathcal{L}(\theta_s) = \underbrace{\mathcal{L}_{\text{sup}}(\theta_s; \mathcal{D}_L)}_{\text{paired}} + \gamma \underbrace{\mathcal{L}_{\text{sp}}(\theta_s, \theta_t; \mathcal{D}_U)}_{\text{unpaired}}, \tag{1}$$

where $\gamma > 0$ balances the two terms. The self-paced consistency learning weight γ is adjusted dynamically with training epoch t following an exponential warm-up schedule (Liu et al. 2021), defined as $\gamma(t) = 0.2 \times e^{-5(1-t/200)^2}$.

Supervised Loss. Given a labeled pair (x_i^l, y_i^l) we combine an ℓ_1 pixel loss with a perceptual loss (Huo, Li, and Zhu 2021):

$$\mathcal{L}_{\text{sup}} = \left\| f_{\theta_s}(x_i^l) - y_i^l \right\|_1 + \lambda_{\text{per}} \mathcal{L}_{\text{per}}\left(f_{\theta_s}(x_i^l),\, y_i^l\right), \tag{2}$$

where ℓ_1 *term* enforces strict pixel fidelity, and *perceptual term* $\mathcal{L}_{\text{per}}$ computes an ℓ_1 distance between high-level VGG-19 features; it encourages color consistency, contrast and texture that align better with human perception. λ_{per} therefore controls the trade-off between pixel accuracy and perceptual realism.

Self-Paced Consistency Loss. For each unlabeled image x_j^u, the teacher model provides a pseudo-label y_p^j. SEA-PACE then modulates the corresponding consistency term by a dynamic weight ω_j:

$$\mathcal{L}_{\text{sp}} = \sum_{j=1}^{M} \omega_j \left\| f_{\theta_s}(x_j^u) - y_p^j \right\|_1, \tag{3}$$

where ω_j is a sample-wise reliability weight generated online by the model-aware reliability estimator.

Overview of SEA-PACE

Figure 2 sketches our framework, which extends the Mean Teacher paradigm with two new components: a *Model-Aware Reliability Estimator* (MARE) that quantifies sample-wise uncertainty via Gaussian Process regression, and a *Self-Paced Consistency Learning* (SPCL) strategy that injects these reliability cues into the consistency loss.

Vanilla Mean Teacher. Let the student model f_{θ_s} be optimized by gradient descent, while the teacher model f_{θ_t} tracks the EMA of the student model,

$$\theta_t \leftarrow \beta\, \theta_t + (1-\beta)\, \theta_s, \quad \beta \in (0,1). \tag{4}$$

For a labeled pair (x_i^l, y_i^l), a supervised loss $\mathcal{L}_{\text{sup}}$ is applied. For an unlabeled image x_j^u, the teacher model produces a pseudo-label $y_p^j = f_{\theta_t}(x_j^u)$ to regularize the student model via a consistency loss:

$$\mathcal{L}_{\text{cons}} = \sum_{j=1}^{M} \| f_{\theta_s}(x_j^u) - y_p^j \|_1. \tag{5}$$

Why Reliability Matters in UIE. Underwater degradations vary dramatically, so the teacher model often generates *high-uncertainty* pseudo-labels for unfamiliar scenarios. Without an explicit measure of reliability, the model treats all pseudo-labels equally in the consistency learning (Eq. (5)), which leads to noise propagation and slows convergence.

Self-Paced Upgrade. To discriminate reliable from risky pseudo-labels, we compute a reliability weight $\omega_j \in (0,1)$ for each x_j^u with our MARE. The SPCL reformulates Eq. (5) as

$$\mathcal{L}_{\text{sp}} = \sum_{j=1}^{M} \omega_j \, \| f_{\theta_s}(x_j^u) - f_{\theta_t}(x_j^u) \|_1, \tag{6}$$

where a higher ω_j is assigned when the GPR-based uncertainty is low. These weights are further modulated by SPCL throughout training. Unlike prior IQA-filtered methods that rely on external metrics (e.g., MUSIQ, URanker), SEA-PACE derives reliability *internally* from the model's uncertainty, enabling a principled, sample-adaptive training pace. The next two subsections detail the MARE component and the scheduling policy used in SPCL.

Model-Aware Reliability Estimator

The core challenge of semi-supervised UIE lies in defining and quantifying the reliability of the pseudo-label $y_p^j = f_{\theta_t}(x_j^u)$. Existing methods rely on model-agnostic IQA scores as proxies for reliability, but these proxies are disconnected from the model's own epistemic state.

In contrast, we characterize the "reliability" of a pseudo-label in terms of the epistemic uncertainty associated with its generation process, which reflects the teacher model's confidence in transferring supervised UIE knowledge to unlabeled data. We then estimate this uncertainty in two steps: (i) we quantify predictive uncertainty in the teacher's latent feature space via Gaussian Process Regression, and (ii) we convert the raw uncertainty values into stable, rank-based reliability weights $\omega_j \in (0,1)$.

GPR-Based Uncertainty Modeling. Let the teacher encoder be $\phi_{\theta_t} : \mathbb{R}^{3\times H\times W} \to \mathbb{R}^d$. For each sample we obtain a d-dimensional feature vector $z = \phi_{\theta_t}(x)$. We regard ϕ_{θ_t} as a random function drawn from a Gaussian Process with RBF kernel $k(\cdot,\cdot)$, which is capable of modeling the discrepancy information in latent space. Given the set of labeled features $\mathcal{F}_L = \{z_i^l\}_{i=1}^N$, we perform a *local GPR* using the K (K is set to 16) nearest neighbors of z_j^u in $\mathcal{F}_L$ to compute the conditional posterior for each unlabeled feature z_j^u:

$$p(z_j^u \mid \mathcal{F}_L) = \mathcal{N}\big(\mu_j, \Sigma_j\big), \tag{7}$$

Algorithm 1: Self-Paced Consistency Learning (one epoch)

Require: labeled set $\mathcal{D}_L$, unlabeled set $\mathcal{D}_U$, MARE, student f_{θ_s}, teacher f_{θ_t}

```
for each mini-batch {x_i^l, y_i^l}, {x_j^u} do
    Supervised loss: L_sup ← ℓ_1 + λ_per L_per
    Pseudo-labels: y_p^j ← f_θt(x_j^u)
    Uncertainty: u_j ← GPR(x_j^u)
    Rank weights: ω̃_j^(t) ← exp(−r_j / τ^(t))
    Momentum update: ω_j^(t) via Eq. (10)
    Self-paced loss: L_sp via Eq. (6)
    Update student: θ_s ← θ_s − η∇_θs(L_sup + γL_sp)
    Update teacher: θ_t ← βθ_t + (1 − β)θ_s
end for
Update τ with Eq. (11)
```

where closed-form μ_j, Σ_j depend on k. Then, We quantify epistemic uncertainty by the log-determinant of the posterior covariance,

$$u_j = \log\det\Sigma_j, \qquad u_j \geq 0. \tag{8}$$

A larger u_j indicates higher uncertainty, meaning that z_j^u lies in a region of the feature space insufficiently supported by the labeled data. Consequently, its pseudo-label is likely unreliable.

Rank-Based Reliability Transformation. Absolute uncertainty values are sensitive to the evolving feature distribution during training, and are therefore unsuitable for directly modulating image-level supervision (Yang et al. 2024; Mirzadeh et al. 2020). We thus transform u_j into a *relative rank*. Let $r_j \in \{1,\dots,M\}$ denote the rank of u_j in ascending order among all unlabeled samples, we then convert r_j into a reliability weight via:

$$\omega_j = \exp(-r_j/\tau), \qquad \omega_j \in (0,1), \tag{9}$$

where the sharpness factor τ is adaptively adjusted by the SPCL strategy. This rank-based transformation (i) stabilizes training, (ii) preserves the relative reliability ordering, and (iii) enables a progressively expanding range of samples to contribute to the consistency loss as τ increases.

Advantages over Prior Filtering. Unlike IQA-filtered frameworks that rely on external metrics, MARE is *task-intrinsic* and couples directly with the network's latent distribution, yielding calibrated and sample-adaptive weights.

Self-Paced Consistency Learning

Conventional consistency learning assigns every unlabeled image the same importance, which allows highly uncertain pseudo-labels to slow or even derail convergence, a situation that is especially acute in underwater scenes with diverse degradations.

To counter this, our SPCL strategy uses the MARE-derived weights $\omega_j \in (0,1)$ for each unlabeled image x_j^u to build a reliability-driven curriculum: the student model is first supervised by the most trustworthy pseudo-labels and gradually absorbs more challenging examples as its competence grows. (Algorithm 1 outlines one training epoch)

Figure 3: Qualitative comparison on UIEB, EUVP, SUIM-E, and U45 benchmarks. The first five rows present the PSNR scores, while the last five rows present the URanker scores. The best result for each metric is highlighted in red, and higher values indicate better performance.

Momentum-Smoothed Weight Update. To mitigate temporal fluctuations in sample reliability, we smooth each weight with an exponential-momentum update:

$$\omega_j^{(t)} = \lambda\, \omega_j^{(t-1)} + (1-\lambda)\, \widetilde{\omega}_j^{(t)}, \tag{10}$$

where $\widetilde{\omega}_j^{(t)}$ is the current reliability weight derived via Eq. (9) and λ controls the smoothing strength.

Loss-Guided Sharpness Scheduling. We let the sharpness factor τ widen or narrow the curriculum according to the recent trend of the self-paced loss $\mathcal{L}_{\text{sp}}$:

$$\begin{aligned} \Delta\tau^{(t)} &= \eta \tanh\Big(-\alpha\,(\Delta\mathcal{L}_{\text{sp}}^{(t)} - \delta)\Big), \\ \tau^{(t)} &= \text{clip}\Big(\tau^{(t-1)} + \Delta\tau^{(t)},\, \tau_{\min}, \tau_{\max}\Big). \end{aligned} \tag{11}$$

We set $\eta = |\mathcal{D}_L|$, $\alpha = 0.125\,|\mathcal{D}_L|$, and derive δ from typical loss variation, while constraining τ within $[0.5\,|\mathcal{D}_L|,\, 2\,|\mathcal{D}_L|]$. When $\mathcal{L}_{\text{sp}}$ decreases or tends to converge, τ is progressively increased to incorporate more diverse and challenging samples. Otherwise, τ is reduced to narrow the sampling range and mitigate the impact of noisy supervision.

Experiments

Experimental Setup

Datasets. We conducted experiments on five widely used UIE datasets: UIEB (Li et al. 2019), EUVP (Islam, Xia, and Sattar 2020), SUIM-E (Qi et al. 2022), and U45 (Li, Li, and Wang 2019). For UIEB, 800 paired samples were used for training and 90 for testing. From the EUVP underwater scene subset, we utilized 2,185 paired images for training and 515 for evaluation. SUIM-E comprises 1,525 paired training images and 110 test images. U60 (60 challenging unpaired underwater images from UIEB) and U45 were used for non-reference evaluation only. Across all ex-

Method	UIEB			EUVP			SUIM-E			U45		
	PSNR	SSIM	MSE	PSNR	SSIM	MSE	PSNR	SSIM	MSE	URanker	UCIQE	UIQM
Ushape	20.89	0.810	1.02	24.41	0.823	0.60	22.46	0.835	0.70	1.67	0.568	0.692
NU2Net	23.69	0.913	<u>0.64</u>	19.58	0.798	1.36	23.52	**0.912**	0.60	**2.12**	0.597	0.777
HCLR-Net	22.96	<u>0.916</u>	0.78	<u>26.31</u>	0.838	0.51	22.97	0.876	0.61	1.78	0.598	0.809
SMDR-IS	23.71	**0.922**	0.75	19.79	0.798	1.33	21.75	0.862	0.91	1.74	0.588	0.767
ERD	22.75	0.881	0.82	24.95	**0.852**	0.58	<u>23.78</u>	<u>0.895</u>	<u>0.58</u>	1.95	0.607	0.798
SS-UIE	21.42	0.871	0.94	26.16	0.840	**0.45**	20.04	0.822	1.16	1.57	0.561	0.683
USUIR	20.98	0.876	1.06	18.09	0.742	1.83	20.19	0.852	1.19	1.09	**0.626**	<u>0.862</u>
UWFormer	23.52	<u>0.916</u>	0.67	24.82	0.841	0.50	22.88	0.872	0.65	1.82	0.603	0.767
SemiUIR	<u>24.26</u>	0.913	**0.63**	26.30	0.834	0.54	<u>23.78</u>	0.867	0.59	1.91	0.613	**0.885**
SEA-PACE	**24.57**	**0.922**	**0.63**	**26.39**	<u>0.851</u>	<u>0.48</u>	**23.96**	0.878	**0.53**	<u>2.07</u>	<u>0.614</u>	0.855

Table 1: Quantitative comparison on UIEB, EUVP, SUIM-E, and U45 benchmarks. The best and second-best results in each column are highlighted in bold and underlined, respectively.

periments, we adopted a supervision ratio of 1:2 (labeled-to-unlabeled). Specifically, for every N paired samples, $2N$ unlabeled images from the unpaired portion of the EUVP dataset were randomly selected to form the unlabeled set.

Implementation Details. Our method was implemented in PyTorch (Paszke et al. 2019) and was trained on a single NVIDIA RTX 4090 GPU. We used the AdamP (Heo et al. 2020) optimizer with an initial learning rate of 2×10^{-4}, which was decayed by a factor of 0.1 at epoch 100. The model was trained for 200 epochs with a batch size of 12, using 4 labeled and 8 unlabeled samples per iteration.

Comparison with State-of-the-Art Methods

Table 1 summarizes the quantitative results, and Figure 3 presents the qualitative comparisons. These results highlight the competitive performance of our approach in both quantitative metrics and visual fidelity. In terms of efficiency, our model achieves 44.5 FPS with 1.34M parameters and requires only 124.4 MB of memory during inference.

Comparison Methods. We compared SEA-PACE with a variety of UIE approaches, considering both quantitative and qualitative performance. These competing methods fall into three categories: 1) supervised Ushape (Peng, Zhu, and Bian 2023), NU2Net (Guo et al. 2023), HCLR-Net (Zhou et al. 2024), SMDR-IS (Zhang et al. 2024), ERD (Cao et al. 2025), SS-UIE (Peng and Bian 2025); 2) unsupervised USUIR (Fu et al. 2022); and 3) semi-supervised UWFormer (Chen et al. 2024) and SemiUIR (Huang et al. 2023).

Full-Reference Evaluation. We employed PSNR (Korhonen and You 2012), SSIM (Wang et al. 2004) and MSE for full-reference evaluation. Higher PSNR/SSIM and lower MSE indicate better enhancement quality. As shown in Table 1, the proposed SEA-PACE consistently ranks among the top three on all three benchmarks. Although the supervised method ERD ranks second on SUIM-E, SEA-PACE demonstrates strong performances across diverse datasets, showcasing the advantage of leveraging unlabeled data. Due to the lack of ground truth, the unsupervised USUIR exhibits inferior performance in both quantitative and qualitative evaluation. Compared to existing semi-supervised methods, our approach consistently yields superior results across all benchmarks and produces more stable and visually natural enhancements with better color fidelity. This performance gain stems from our MARE-assisted SPCL strategy, which enables a model-aware, progressive selection of reliable training samples throughout the training process.

Non-Reference Evaluation. We used URanker (Guo et al. 2023), UCIQE (Yang and Sowmya 2015) and UIQM (Panetta, Gao, and Agaian 2015) for non-reference evaluation. Higher values indicate better enhancement quality. As shown in Table 1, our proposed SEA-PACE achieves the second-best scores on both URanker and UCIQE, and the third-best on UIQM. Although the hand-crafted metrics UCIQE and UIQM capture low-level statistical properties, they often do not align with human perceptual quality, as reported by (Jiang et al. 2025). For example, although the unsupervised method USUIR achieves the highest UCIQE score and the second-highest UIQM score, it produces severe visual artifacts, such as obvious color distortions and blurring, as shown in Figure 3. In contrast, URanker is a data-driven quality assessment metric that more closely aligns with human perceptual preferences. NU2Net obtains the highest URanker score, which is expected since it explicitly uses URanker as a training objective. Notably, SEA-PACE surpasses all other methods by a clear margin on URanker, reflecting its ability to fully exploit both labeled and unlabeled data through a reliability-aware learning paradigm. This advantage is further reflected in the qualitative comparisons shown in Figure 3, where our method generates visually compelling results with balanced brightness, natural color tones, and rich details.

Ablation Study

As shown in Table 2 and Figure 4, we analyzed four ablation settings by progressively adding our key components to a baseline.

M1: Baseline. Our baseline is a vanilla Mean Teacher framework. Despite leveraging both labeled and unlabeled data, it lacks a mechanism to assess pseudo-label reliability. Consequently, this model yields suboptimal results and even underperforms supervised methods like SMDR-IS and

NU2Net (Table 1). Qualitatively, it often suffers from greenish color casts in seafloor scenes and fails to recover fine structures in severely blurred or low-visibility regions.

index	MT	MARE		SPCL	UIEB	U45
		GPR	RT		PSNR	URanker
M1	√	-	-	-	23.62	1.70
M2	√	√	-	-	24.28	1.89
M3	√	√	√	-	24.42	2.00
M4	√	√	√	√	**24.57**	**2.07**

Table 2: Ablation study of the components in our proposed SEA-PACE on the UIEB and U45 benchmarks. The best results are highlighted in bold.

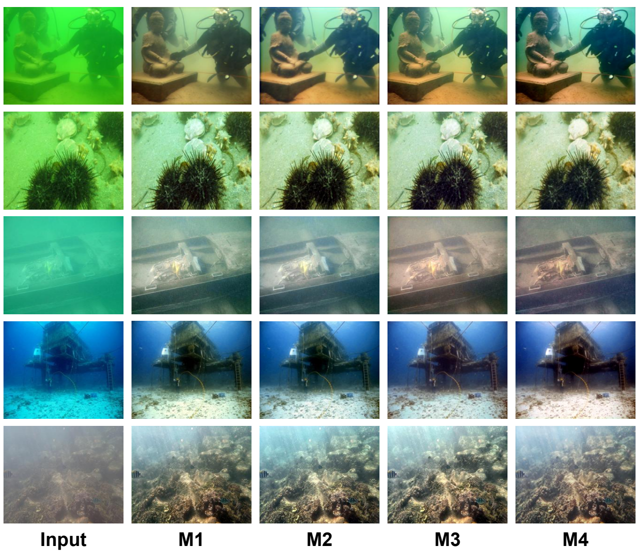

Figure 4: Visual results of ablation study.

M2: Baseline + GPR. Strikingly, introducing the Gaussian Process Regression (GPR) module yields a substantial performance gain, exceeding the baseline by 0.66 on PSNR and 0.19 on URanker. As shown in Figure 4, the greenish cast, particularly in seafloor scenes, is notably reduced, indicating that uncertainty-aware filtering encourages the model to focus on more trustworthy pseudo-labels and reduces the impact of noisy supervision that previously led to color distortions. During training, the GPR module introduces only a minor computational overhead of 1.15 ms per sample.

M3: Baseline + MARE. Building upon M2, M3 incorporates the rank-based transformation, completing our Model-Aware Reliability Estimator (MARE) and yielding further improvements across all metrics. Figure 4 shows that M3 produces sharper edges with better color correction. Moreover, Figure 5 showcases a visual comparison between pseudo-labels with varying reliability weights. As demonstrated, samples assigned high reliability weights consistently correspond to high-quality enhancements. Conversely, samples with low weights exhibit color casts or blurred details. This strong correlation validates that our MARE can effectively distinguish reliable predictions from ambiguous ones, serving as an excellent proxy for perceptual quality.

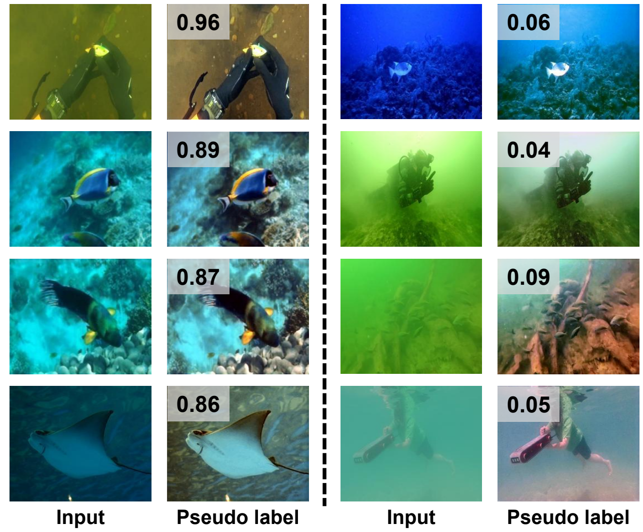

Figure 5: Visual results of the reliability assessment performed by MARE on pseudo-labels. The first two columns show samples with high reliability weights and correspondingly superior visual quality, whereas the last two columns show samples with low reliability weights that exhibit color casts.

M4: Full Model. Finally, adding Self-Paced Consistency Learning (SPCL) strategy creates the complete SEA-PACE model. Instead of relying on static IQA-guided pesudo-label selection, SPCL gradually incorporates more challenging unlabeled samples under a loss-guided scheduler, preventing overfitting to easy data and enhancing generalization. As a result, SEA-PACE achieves the best overall performance, outperforming the baseline in PSNR (24.57 vs. 23.62) and a remarkable 21.8% in URanker (2.07 vs. 1.70). As illustrated in Figure 4, our full model offers both high-fidelity structure preservation and perceptually pleasing enhancement quality.

Data Scalability. To evaluate the data scalability of SEA-PACE, we conduct experiments with different supervision ratios. Using 1:1, 1:2, and 1:4 ratios yields consistent PSNR gains (24.47, 24.57, and 24.69) as the amount of unlabeled data increases, at the expense of increasing the training cost. Therefore, we adopt the 1:2 ratio as a balanced trade-off between performance and efficiency.

Conclusion

In this work, we propose SEA-PACE, a novel semi-supervised learning framework for underwater image enhancement, which integrates a model-aware reliability estimator (MARE) and a self-paced consistency learning (SPCL) strategy. Extensive experiments show SEA-PACE outperforms existing methods. Beyond our UIE task, the proposed uncertainty estimation mechanism and reliability-aware self-paced learning strategy offer flexible solutions applicable to other semi-supervised or weakly-supervised vision tasks challenged by noisy or uncertain supervision.

Acknowledgments

This work was supported by the Natural Science Foundation of Shandong Province of China (ZR2023QF145, ZR2023MF012, ZR2024ZD04), the Key R&D Program of Shandong Province (2024ZLGX06), and the Natural Science Foundation of China (62531022).

References

Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N. E.; and McGuinness, K. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International joint conference on neural networks (IJCNN)*, 1–8. IEEE.

Bi, H.; Chen, L.; Cao, J.; Wang, J.; Sun, J.; Rao, Y.; and Dong, J. 2025. SeaDiff: Underwater Image Enhancement with Degradation-Aware Diffusion Model. *IEEE Transactions on Circuits and Systems for Video Technology*.

Cao, J.; Peng, W.; Liu, Y.; Dong, J.; Le Callet, P.; and Kwong, S. 2025. ERD: Encoder-Residual-Decoder Neural Network for Underwater Image Enhancement. *IEEE Transactions on Circuits and Systems for Video Technology*.

Cao, J.; Wang, R.; Jia, Y.; Zhang, X.; Wang, S.; and Kwong, S. 2021. No-reference image quality assessment for contrast-changed images via a semi-supervised robust PCA model. *Information Sciences*, 574: 640–652.

Chen, W.; Lei, Y.; Luo, S.; Zhou, Z.; Li, M.; and Pun, C.-M. 2024. Uwformer: Underwater image enhancement via a semi-supervised multi-scale transformer. In *2024 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.

Chiang, J. Y.; and Chen, Y.-C. 2011. Underwater image enhancement by wavelength compensation and dehazing. *IEEE transactions on image processing*, 21(4): 1756–1769.

Cong, R.; Yang, W.; Zhang, W.; Li, C.; Guo, C.-L.; Huang, Q.; and Kwong, S. 2023. Pugan: Physical model-guided underwater image enhancement using gan with dual-discriminators. *IEEE Transactions on Image Processing*, 32: 4472–4485.

Cong, X.; Zhao, Y.; Gui, J.; Hou, J.; and Tao, D. 2024. A comprehensive survey on underwater image enhancement based on deep learning. *arXiv preprint arXiv:2405.19684*.

Fu, Z.; Lin, H.; Yang, Y.; Chai, S.; Sun, L.; Huang, Y.; and Ding, X. 2022. Unsupervised underwater image restoration: From a homology perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 643–651.

Guo, C.; Wu, R.; Jin, X.; Han, L.; Zhang, W.; Chai, Z.; and Li, C. 2023. Underwater ranker: Learn which is better and how to be better. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 702–709.

Heo, B.; Chun, S.; Oh, S. J.; Han, D.; Yun, S.; Kim, G.; Uh, Y.; and Ha, J.-W. 2020. Adamp: Slowing down the slow-down for momentum optimizers on scale-invariant weights. *arXiv preprint arXiv:2006.08217*.

Huang, S.; Wang, K.; Liu, H.; Chen, J.; and Li, Y. 2023. Contrastive semi-supervised learning for underwater image restoration via reliable bank. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18145–18155.

Huo, F.; Li, B.; and Zhu, X. 2021. Efficient wavelet boost learning-based multi-stage progressive refinement network for underwater image enhancement. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1944–1952.

Islam, M. J.; Xia, Y.; and Sattar, J. 2020. Fast underwater image enhancement for improved visual perception. *IEEE Robotics and Automation Letters*, 5(2): 3227–3234.

Jamadandi, A.; and Mudenagudi, U. 2019. Exemplar-based underwater image enhancement augmented by wavelet corrected transforms. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 11–17.

Jiang, Q.; Gu, Y.; Wu, Z.; Li, C.; Xiong, H.; Shao, F.; and Wang, Z. 2025. Deep underwater image quality assessment with explicit degradation awareness embedding. *IEEE Transactions on Image Processing*.

Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5148–5157.

Korhonen, J.; and You, J. 2012. Peak signal-to-noise ratio revisited: Is simple beautiful? In *2012 Fourth international workshop on quality of multimedia experience*, 37–38. IEEE.

Li, C.; Guo, C.; Ren, W.; Cong, R.; Hou, J.; Kwong, S.; and Tao, D. 2019. An underwater image enhancement benchmark dataset and beyond. *IEEE transactions on image processing*, 29: 4376–4389.

Li, H.; Li, J.; and Wang, W. 2019. A fusion adversarial underwater image enhancement network with a public test dataset. *arXiv preprint arXiv:1906.06819*.

Liu, Y.; Zhu, L.; Pei, S.; Fu, H.; Qin, J.; Zhang, Q.; Wan, L.; and Feng, W. 2021. From synthetic to real: Image dehazing collaborating with unlabeled real data. In *Proceedings of the 29th ACM international conference on multimedia*, 50–58.

Mirzadeh, S. I.; Farajtabar, M.; Li, A.; Levine, N.; Matsukawa, A.; and Ghasemzadeh, H. 2020. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 5191–5198.

Panetta, K.; Gao, C.; and Agaian, S. 2015. Human-visual-system-inspired underwater image quality measures. *IEEE Journal of Oceanic Engineering*, 41(3): 541–551.

Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

Peng, L.; and Bian, L. 2025. Adaptive Dual-domain Learning for Underwater Image Enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 6461–6469.

Peng, L.; Zhu, C.; and Bian, L. 2023. U-shape transformer for underwater image enhancement. *IEEE Transactions on Image Processing*, 32: 3066–3079.

Qi, Q.; Li, K.; Zheng, H.; Gao, X.; Hou, G.; and Sun, K. 2022. SGUIE-Net: Semantic attention guided underwater image enhancement with multi-scale perception. *IEEE Transactions on Image Processing*, 31: 6816–6830.

Rasmussen, C. E. 2003. Gaussian processes in machine learning. In *Summer school on machine learning*, 63–71. Springer.

Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30.

Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.

Wang, Z.; Shen, L.; Xu, M.; Yu, M.; Wang, K.; and Lin, Y. 2023. Domain adaptation for underwater image enhancement. *IEEE Transactions on Image Processing*, 32: 1442–1457.

Xu, S.; Wang, J.; He, N.; Xu, G.; and Zhang, G. 2024. Optimizing underwater image enhancement: integrating semi-supervised learning and multi-scale aggregated attention. *The Visual Computer*, 1–19.

Yan, S.; Chen, X.; Wu, Z.; Tan, M.; and Yu, J. 2023. Hybrur: A hybrid physical-neural solution for unsupervised underwater image restoration. *IEEE Transactions on image processing*, 32: 5004–5016.

Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything v2. *Advances in Neural Information Processing Systems*, 37: 21875–21911.

Yang, M.; and Sowmya, A. 2015. An underwater color image quality evaluation metric. *IEEE Transactions on Image Processing*, 24(12): 6062–6071.

Yin, J.; Wang, Y.; Guan, B.; Zeng, X.; and Guo, L. 2024. Unsupervised underwater image enhancement based on disentangled representations via double-order contrastive loss. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–15.

Zhang, D.; Zhou, J.; Guo, C.; Zhang, W.; and Li, C. 2024. Synergistic multiscale detail refinement via intrinsic supervision for underwater image enhancement. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 7033–7041.

Zhou, J.; Sun, J.; Li, C.; Jiang, Q.; Zhou, M.; Lam, K.-M.; Zhang, W.; and Fu, X. 2024. HCLR-net: Hybrid contrastive learning regularization with locally randomized perturbation for underwater image enhancement. *International Journal of Computer Vision*, 1–25.