

Contrastive Semi-supervised Learning for Underwater Image Restoration via Reliable Bank

Shirui Huang^{1*}, Keyan Wang^{1*†}, Huan Liu², Jun Chen², Yunsong Li¹

¹ Xidian University, ² McMaster University, * Equal Contribution, † Corresponding Author

srhuang@stu.xidian.edu.cn, {kywang, ysli}@mail.xidian.edu.cn, {liuh127, chenjun}@mcmaster.ca

Abstract

Despite the remarkable achievement of recent underwater image restoration techniques, the lack of labeled data has become a major hurdle for further progress. In this work, we propose a mean-teacher based **Semi-supervised Underwater Image Restoration (Semi-UIR)** framework to incorporate the unlabeled data into network training. However, the naive mean-teacher method suffers from two main problems: (1) The consistency loss used in training might become ineffective when the teacher’s prediction is wrong. (2) Using L1 distance may cause the network to overfit wrong labels, resulting in confirmation bias. To address the above problems, we first introduce a reliable bank to store the “best-ever” outputs as pseudo ground truth. To assess the quality of outputs, we conduct an empirical analysis based on the monotonicity property to select the most trustworthy NR-IQA method. Besides, in view of the confirmation bias problem, we incorporate contrastive regularization to prevent the overfitting on wrong labels. Experimental results on both full-reference and non-reference underwater benchmarks demonstrate that our algorithm has obvious improvement over SOTA methods quantitatively and qualitatively. Code has been released at <https://github.com/Huang-ShiRui/Semi-UIR>.

1. Introduction

Due to light refraction, absorption and scattering in underwater scenes, images taken in the water usually suffer severely from color distortion, low contrast and blur. Images with these defects tend to be less visually appealing and can potentially hinder the well-functioning of underwater robotic systems. Recently, many deep learning based methods [5–7, 24, 51] have been proposed to address image restoration problems. Numerous efforts have also been devoted to the specific domain of underwater image restora-



Figure 1. Examples from different benchmarks. (a) shows real-world underwater images from UIEB [22] with degraded images (first and second row). (b) shows the UWCNN training set [21] (synthesized based on the image formation model) and (c) shows the EUVP dataset [17] (synthesized by GAN). The ambient light and color cast of (b) and (c) are quite different from that of (a).

tion [11, 17, 20, 22, 49]. Compared with traditional methods that mostly rely on hand-crafted priors, deep learning based solutions are able to deliver superior restoration results due to their data-driven nature.

Despite their success, most of deep learning based methods are designed to learn the restoration mapping on paired datasets in a supervised manner. As is known, it is extremely hard, if not impossible, to acquire paired underwater images in real scenes. The existing datasets for underwater image restoration have several non-negligible issues: (1) *Lack of real data*. A popular way to construct paired datasets is to synthesize underwater images using some physical model [21] or GAN [17, 48]. However, there is a significant discrepancy between synthesized and real data. As is shown in Fig. 1, the ambient light and color cast of synthetic data are quite different from the real counterparts. Due to domain shift, models trained on synthetic datasets often exhibit poor generalization in real scenes. Another way [22] is to manually construct pseudo labels by selecting the best results among those produced by traditional al-

¹This work is supported by the Nature Science Foundation of Shaanxi Province of China (2021JM-125).

gorithms. Inevitably, the quality of the pseudo ground truth is restricted by the restoration capability of traditional algorithms. (2) *Limited data size*. The current benchmarks only provide a very limited amount of paired data. For example, in UIEB [22], there are only 890 paired underwater images. Models learned on such a small dataset may run the risk of overfitting. In comparison, the standard datasets for image classification, such as ImageNet [35], are several orders of magnitude larger in size.

On the other hand, unlabeled underwater images are relatively easy to collect. The challenge is how to make effective use of these unlabeled data. Semi-supervised learning, which capitalizes on both labeled and unlabeled data for model training, is best suited in this kind of scenarios. This motivates us to propose a semi-supervised scheme with the goal of improving the generalization of the resulting model on real-world underwater images. To be specific, we adopt the mean teacher method [40] as the basis. The mean teacher method finds a way to obtain pseudo labels for unlabeled data and utilizes a consistency loss to improve the accuracy and robustness of the network. Specifically, it constructs a teacher model with improved performance from a student model via the exponential moving average (EMA) strategy. The teacher’s prediction serves as the pseudo label to guide the training of the student. However, it is a non-trivial task to tailor the mean teacher method to the underwater image restoration problem. The reasons are as follows: (1) There is no guarantee that the teacher can consistently outperform the student. Wrong pseudo labels may jeopardize the training of the student network. (2) The commonly used consistency loss is based on L1 distance. The “strict” L1 loss can easily make the model overfit wrong predictions, resulting in confirmation bias.

To address the first issue, we construct a reliable bank to archive the best-ever outputs from the teacher as pseudo labels. The main challenge here is how to determine what are the “best-ever” outputs? Intuitively, non-reference image quality assessment (NR-IQA) can be leveraged to evaluate the quality of each output. However, as noted in [3, 10, 22], the current NR-IQA metrics for underwater images are, to some extent, inconsistent with human visual perception. To identify the right one for our purpose, we compare several NR-IQA metrics using the monotonicity property as the reliability criterion. Our empirical analysis suggests MUSIQ [19] best meets the criterion. For the second issue, we introduce contrastive learning as a supplementary regularization to alleviate overfitting. Unlike those conventional loss functions that are only concerned with how close the outputs and ground truths are, contrastive loss provides additional supervision to prevent the degradation of the outputs. In this sense, contrastive regularization is ideally suited to our semi-supervised learning framework since we only have access to the degraded images in the unlabeled dataset. It

enables the model to take advantage of unlabeled data.

In summary, our main contributions are as follows: (1) We propose a mean teacher based semi-supervised underwater image restoration framework named Semi-UIR, which effectively leverages the knowledge from unlabeled data to improve the generalization of the trained model on real-world data. (2) We evaluate teacher outputs by a judiciously chosen NR-IQA metric and build a reliable bank to store best-ever teacher outputs, which ensures the reliability of pseudo-labels. (3) We adopt contrastive loss as a form of regularization to alleviate confirmation bias. (4) Extensive experimental results demonstrate the effectiveness of our proposed methods.

2. Related Work

2.1. Underwater Image Restoration Methods

Traditional underwater image restoration methods can be categorized into model-based and model-agnostic methods. Model-based methods [3, 8, 33] use hand-crafted priors to estimate unknown parameters of underwater imaging model [36], such as transmission and ambient light. In contrast, model-agnostic methods rely on the design of appropriate image enhancement techniques such as CLAHE [14], Retinex [52], fusion [2] and MMLE [53]. Despite their success, the traditional methods usually fail to cope with complex real scenes.

Most of early deep learning based underwater image restoration methods [18, 41] accomplish the goal by exploiting physical imaging models. Specifically, they make use of neural networks to estimate the transmission and ambient light. However, inaccurate estimation of these parameters hinders such methods from achieving good performance. Recent years have seen many deep learning based methods that directly learn the restoration mapping from the labeled dataset in a supervised manner without resorting to imaging models. [22] employs an effective network to fuse three feature maps enhanced by traditional model-agnostic methods. [20] designs a multi-color space encoder and a transmission-guided decoder by leveraging the ideas from traditional model-based approaches. [16] proposes a wavelet boost learning strategy, through which features in the frequency domain are utilized for fine detail restoration. [17] introduces an end-to-end network based on GAN for the purpose of real-time inference.

2.2. Semi-supervised Learning

In recent years, semi-supervised learning [55] has played an increasingly important role in tackling computer vision problems. It focuses on making effective use of both labeled and unlabeled data. Many semi-supervised methods have been developed, such as mean teacher [40], virtual adversarial learning [30] and FixMatch [38]. Among them, the

mean teacher method [40], which is based on consistency regularization, has achieved remarkable success in semi-supervised image recognition. This success also triggers its applications to other vision tasks such as semantic segmentation [15, 26, 44] and image restoration [27, 42]. Unfortunately, to the best of our knowledge, semi-supervised learning is rarely explored in underwater image restoration. [54] makes an initial attempt in this direction by training a single network with both supervised and unsupervised losses. In comparison, we adopt a more systematic approach and introduce several techniques in handling unlabeled data, including mean teacher, reliable bank and contrastive loss.

2.3. Contrastive Learning

Contrastive learning has emerged as an effective paradigm in self-supervised learning [4, 9, 12]. It enables visual representation to be learned empirically by instance discrimination, through which similar samples are kept close to each other whereas dissimilar samples are separated far apart. To take advantage of contrastive learning in image restoration, previous works largely focus on the construction of contrastive samples and feature space. For example, [23, 46] take clean images as positive instances and degraded images as negative instances and then project them into a new feature space by VGG [37]. Note that in the above works, contrastive loss is applied in a supervised manner, which is infeasible for unlabeled data. How to make contrastive loss applicable to unlabeled data remains an open problem. [11] is the first to employ contrastive learning in the context of underwater image restoration. However, it is still a supervised learning method in nature, and contrastive loss is used as a regularization term to boost the performance of supervised learning. Against this backdrop, the present work puts forward a systematic approach for making use of contrastive learning to exploit unlabeled data.

3. Method

3.1. Problem Formulation

Semi-supervised learning aims to enable a learning system to learn from both labeled and unlabeled data. The problem for underwater image restoration is defined as follows. Let $D_L = \{(x_i^l, y_i^l) | x_i^l \in \mathcal{I}_s^{LQ}, y_i^l \in \mathcal{I}_s^{HQ}\}_{i=1}^N$ denote the labeled dataset, where x_i^l and y_i^l are respectively the underwater image and clean ground truth from degraded set \mathcal{I}_s^{LQ} and ground truth set \mathcal{I}_s^{HQ} . Similarly, let $D_U = \{x_i^u | x_i^u \in \mathcal{I}_u^{LQ}\}_{i=1}^M$ denote the unlabeled dataset, where x_i^u is the underwater image sampled from the degraded set \mathcal{I}_u^{LQ} . It is worth mentioning that the data in D_L and D_U are disjoint, *i.e.* $D_L \cap D_U = \emptyset$. Our goal is to learn a mapping on $D = D_L \cup D_U$ that converts an underwater image x to its clean counterpart y .

3.2. Semi-supervised Underwater Restoration

Our semi-supervised learning framework follows the typical setup in semi-supervised learning [38, 40], as illustrated in Fig. 2. Specifically, our Semi-UIR consists of two networks of the same structure, called teacher and student respectively. The two networks differ mainly in how their weights are updated.

The teacher’s weights θ_t are updated by exponential moving average (EMA) of the student’s weights θ_s :

$$\theta_t = \eta\theta_t + (1 - \eta)\theta_s, \quad (1)$$

where $\eta \in (0, 1)$ is the momentum. Using this update strategy, the teacher model can aggregate previously learned weights immediately after each training step. As is noted in [34], temporal weight averaging can stabilize the training process and help improve the performance compared with standard gradient descent.

The weights of student network θ_s are updated using gradient descent. Usually, the optimization of the student network can be formulated as minimizing the following loss:

$$L_{total} = L_{sup} + \eta L_{un}, \quad (2)$$

where $L_{sup} = \sum_{i=0}^N |f_{\theta_s}(x_i^l) - y_i^l|$ denotes the supervised loss and $L_{un} = \sum_{i=0}^M |f_{\theta_s}(\phi_s(x_i^u)) - f_{\theta_t}(\phi_t(x_i^u))|$ represents the unsupervised teacher-student consistency loss. $|\cdot|$ refers to L1 distance. ϕ_s and ϕ_t are respectively the data augmentations of student’s inputs and teacher’s inputs.

Ideally, since the teacher network is in general better than the student network, L_{un} could provide effective supervision to train the student network on the unlabeled dataset. We thus refer to the teacher’s output $\hat{y}_i^u = f_{\theta_t}(\phi_t(x_i^u))$ as pseudo label. However, it is not guaranteed that the outputs of the teacher are consistently better than those of the student. Wrong pseudo labels can potentially jeopardize the training of the student network.

3.3. Reliable Teacher-Student Consistency

To address the above issue, we shall select the reliable outputs of teacher as pseudo labels. In image classification and semantic segmentation [15, 38, 44], the reliability of the network’s outputs is usually measured by entropy and confidence. However, the extension to image restoration problems does not appear to be straightforward due to the presence of new challenges. In particular, as a regression task, underwater image restoration requires recovering fine textures and removing color cast.

To this end, we propose a reliable bank to store the best-ever outputs of the teacher network during the training process. To be specific, We first initialize our reliable bank to be an empty set, *i.e.* $\mathcal{B}_U = \emptyset$. In each training iteration, we compare the current output of teacher with both the student’s output and the pseudo label in the reliable bank. If

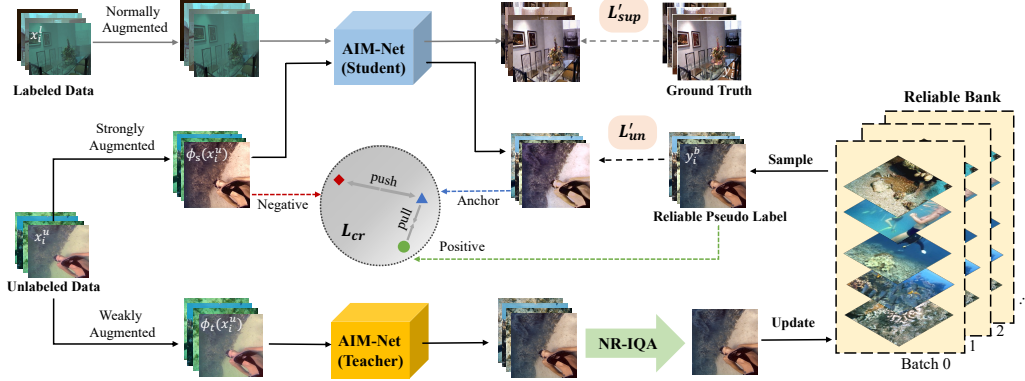


Figure 2. Illustration of our framework Semi-UIR. Semi-UIR is based on the mean teacher scheme with a student model and a teacher model. To guarantee the reliability of pseudo labels for unlabeled data, we build a reliable bank to archive best-ever teacher outputs measured by NR-IQA. Reliable pseudo labels guide the student’s training via the unsupervised teacher-student consistency loss L'_{un} and contrastive loss L_{cr} . The weights of the student are updated by minimizing the supervised loss (L'_{sup}) and unsupervised losses (L'_{un} and L_{cr}). The teacher is updated with EMA from the student.

the teacher’s output is the best in quality, then we replace the pseudo label in the reliable bank with the teacher’s current output. In this way, we could maintain a reliable bank $\mathcal{B}_U = \{y_i^b\}_{i=1}^M$. Note that $D' = D_U \cup \mathcal{B}_U = \{(x_i^u, y_i^b)\}_{i=1}^M$ is a pseudo labeled dataset. This reliable bank can keep track of the best pseudo labels and therefore avoid the wrong labels involved in the calculation of the unsupervised consistency loss L_{un} . Then, we can re-write the L_{un} in Eq. (2) as:

$$L'_{un} = \sum_{i=0}^M |f_{\theta_s}(\phi_s(x_i^u)) - y_i^b|. \quad (3)$$

Now arises the obvious question: How to determine the quality of a prediction without the true label?

3.4. Reliable Metric Selection

Intuitively, we could resort to non-reference image quality assessment (NR-IQA). Unfortunately, as is noted in [3, 22], the commonly used UCIQE [47] and UIQM [31] cannot accurately reflect the quality of restored underwater images. Therefore, building our reliable bank based on such metrics is questionable. To find the best possible NR-IQA for underwater images, we conduct an empirical analysis of several NR-IQA metrics.

Given a degraded underwater image x^l and a paired clean image y^l , we perform various linear combination of them to get a set of images with different quality. Specifically, let $\alpha_i = 0.1 \times i, i = 1, 2, \dots, 10$, we can obtain a set of ten images $\{\alpha_i x^l + (1 - \alpha_i) y^l\}_{i=1}^{10}$. With the increase of α_i , the visual quality of the corresponding image deteriorates, as shown in Fig. 3. It thus makes sense to evaluate the NR-IQA metrics based on how well they capture this monotonicity law. In particular, an NR-IQA metric is identified as reliable if its score on the $\alpha_i x^l + (1 - \alpha_i) y^l$ decreases

Algorithm 1 Update of Reliable Bank

Require: NR-IQA method $\Psi(\cdot)$;
Initialize $\mathcal{B}_U = \emptyset$;
Sample a batch of unlabeled images $\{x_i^u\}_{i=1}^b$ from D_U ;
for each x_i^u **do**
 Get teacher’s prediction: $\hat{y}_i^u = f_{\theta_t}(\phi_t(x_i^u))$;
 Get student prediction: $\tilde{y}_i^u = f_{\theta_s}(\phi_s(x_i^u))$;
 Compute NR-IQA scores of $\hat{y}_i^u, \tilde{y}_i^u$ and $y_i^b \in \mathcal{B}_U$:
 $z_t = \Psi(\hat{y}_i^u), z_s = \Psi(\tilde{y}_i^u), z_b = \Psi(y_i^b)$;
 if $z_t > z_s$ and $z_t > z_b$ **then**
 Replace the y_i^b in \mathcal{B}_U by \hat{y}_i^u ;
 end if
end for




Figure 3. Examples of image fusion based on different α .

with the increase of α_i . Following this rule, we conduct experiments with seven NR-IQA approaches on EUVP benchmark [17], as it covers a wide range of underwater scenes. The experimental results are shown in Fig. 4. We can observe that the deep learning based MUSIQ [19] is most in line with the monotonicity law. Therefore, it is selected to measure the reliability of the networks’ outputs.

The overall procedure of our reliable bank construction is summarized in Algorithm 1.

3.5. Contrastive Regularization

Typically, many mean teacher based methods uses L1 distance as consistency loss as is shown by L_{un} in Eq. (2).

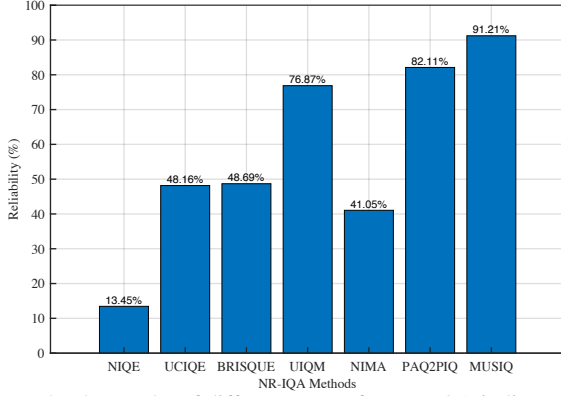


Figure 4. The results of different non-reference IQA indicators on EUVP benchmark, including UIQM [31], UCIQE [47], BRISQUE [28], NIQE [29], NIMA [39], PAQ2PIQ [50] and MUSIQ [19].

The simple consistency loss can easily make the student model overfit on wrong predictions, resulting in confirmation bias. To address this problem, we introduce contrastive loss in the training. Contrastive learning has emerged as an effective paradigm in the self-supervised domain [4, 9, 12]. The goal of contrastive learning is to enable a model to produce a similar representation of *positive* pairs and a dissimilar representation of *negative* pairs. Recently, it has been extended to address image restoration problems [23, 46]. Despite the tremendous success, they usually construct contrastive loss on paired datasets, where the *positive* and *negative* samples are respectively the labels and degraded images. In this section, we propose to incorporate contrastive loss in handling unlabeled data. To achieve this, we first need to construct positive and negative pairs. [43] provide an idea of directly using the teacher’s output as positive samples. However, due to the wrong label problem we have discussed in the previous sections, using the teacher’s outputs as positive samples might be harmful.

Thanks to our proposed reliable bank, where the samples are potentially of higher quality than the student’s outputs, we can take y_i^b as our positive sample. For the negative sample, we follow [23, 43, 46] to take the strongly augmented degraded image $\phi_s(x_i^u)$ as our negative sample. After constructing the positive and negative samples, we can calculate the contrastive loss as follows:

$$L_{cr} = \sum_{j=1}^K \sum_{i=1}^M \omega_j \frac{|\varphi_j(\tilde{y}_i^u), \varphi_j(y_i^b)|}{|\varphi_j(\tilde{y}_i^u), \varphi_j(\phi_s(x_i^u))|}, \quad (4)$$

where $\tilde{y}_i^u = f_{\theta_s}(\phi_s(x_i^u))$ is the student’s prediction on the unlabeled dataset D_U . $\varphi_j(\cdot)$ represents the j_{th} hidden layer of the pre-trained VGG-19 [37] and ω_j is the weight coefficient. We use L1 loss to measure the distance in feature space between the students’ outputs with the positive and negative samples.




Figure 5. An overview of the proposed Asymmetric Illumination-aware Multi-scale Network (AIM-Net).

3.6. Overall Optimization Objective

Similar to Eq. (2), our final optimization objective consists of supervised loss and unsupervised loss.

For the supervised loss, unlike the one defined in Eq. (2) that only calculate the L1 distance, we follow [16] to extend the original L_{sup} by adding perceptual loss L_{per} and gradient penalty L_{grad} :

$$L'_{sup} = L_{sup} + \beta_1 L_{per} + \beta_2 L_{grad}. \quad (5)$$

For the unsupervised loss, we replace the original L_{un} by a combination of the proposed reliable teacher-student consistency loss and contrastive loss:

$$L''_{un} = L'_{un} + \gamma L_{cr}. \quad (6)$$

Finally, we rewrite our overall optimization objective following Eq. (2):

$$L_{overall} = L'_{sup} + \lambda L''_{un}. \quad (7)$$

Due to the page limit, please refer to supplementary material for the detailed perceptual loss and gradient penalty.

4. Experimental Results

4.1. Implementation Details

Network Structure Our student and teacher model are based on the same structure, AIM-Net. To tackle the prominent issues with underwater images (e.g., low contrast, color distortion and blur), certain prior information of such images (e.g., illumination prior and gradient prior) are effectively exploited. As is shown in Fig. 5, the network consists of two branches: illumination-aware restoration branch and gradient branch. The restoration branch incorporates illumination prior to enhance the color and light source perception capabilities. The gradient branch is introduced to enhance the edge structure. Please refer to supplementary material for detailed network structure.

Training Details Our method is implemented using Pytorch library [32] and conducted on NVIDIA RTX 3090 GPUs. We use AdamP [13] as our optimizer. In consideration of its fast convergence to optimum, we select AdamP mainly

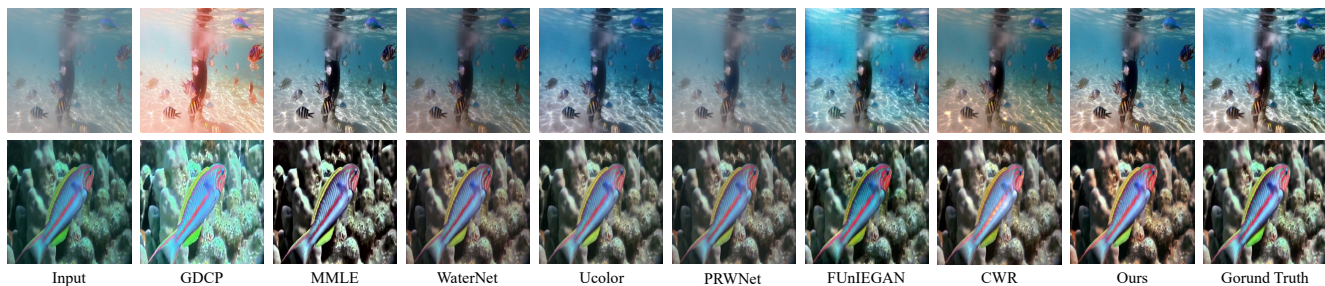


Figure 6. Visual comparisons of full-reference data from UIEB benchmark.

Table 1. Evaluations of different methods on full-reference benchmarks in terms of PSNR and SSIM. Best results are in **bold** and the second best results are with underline.

Method	testS		testR	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
Input	14.64	0.641	18.23	0.746
GDCP [33]	12.89	0.576	15.78	0.757
MMLE [53]	12.76	0.651	20.01	0.781
WaterNet [22]	15.44	0.706	21.58	0.858
Ucolor [20]	<u>23.32</u>	0.853	<u>22.92</u>	<u>0.881</u>
PRWNet [16]	17.27	0.723	20.98	0.848
FGAN [17]	18.54	0.743	19.41	0.824
CWR [11]	14.79	0.697	21.87	0.815
Semi-UIR	23.40	<u>0.821</u>	24.59	0.901

to reduce the training time. During training, we use a mini-batch size of 16, where 8 samples are labeled and 8 samples are unlabeled. The initial learning rate is set to $2e^{-4}$. We train for 200 epochs with the learning rate multiplied by 0.1 at 100 epochs. The training images are all cropped to a size of 256×256 . For the data augmentation on unlabeled data, we only apply resize on the teacher’s inputs and impose strong data augmentation on the student’s inputs. The strong augmentation includes resize, color jitter, gaussian blur and gray scale. Labeled data is normally augmented, including resize, random crop and rotation. The weights of different loss components are set as follows: $\beta_1 = 0.3$, $\beta_2 = 0.1$, $\gamma = 1$ and λ is updated with training epoch t following an exponential warming up function [27]: $\lambda(t) = 0.2 \times e^{-5(1-t/200)^2}$.

4.2. Datasets

Our training set contains 1600 labeled image pairs and 1600 unlabeled images. The labeled image pairs are randomly sampled from [21] and UIEB [22] with a ratio of 1:1. [21] provides a synthesized underwater image dataset in indoor scene. The UIEB [22] dataset contains 890 real underwater images with corresponding ground truths. The unlabeled images are sampled from the unpaired data in the EUVP benchmark [17], which cover a variety of underwater scenes, water types and lighting conditions.

Test set is built with full-reference and non-reference

benchmarks. Full-reference test set includes 110 pairs from [21] and 90 pairs from UIEB [22], namely testS and testR. Non-reference test set includes nearly 700 real world underwater images without ground truths from benchmarks such as UIEB, EUVP, RUIE [25] and Seathru [1].

4.3. Comparison with the State-of-the-Arts

We compare our proposed Semi-UIR with seven state-of-the-art underwater restoration methods, including two traditional methods (GDCP [33], MMLE [53]) and five deep learning based methods (WaterNet [22], Ucolor [20], FUnIE-GAN [17], PRWNet [16] and CWR [11]). All the compared methods are re-trained on our training dataset. For testS and testR, we conduct full-reference evaluations using PSNR, SSIM [45]. For non-reference test set, we provide evaluation results using UIQM [31], UCIQE [47] and MUSIQ [19].

Results on full-reference datasets. The quantitative results on testS and testR are shown in Table 1. On testS, our method performs the best in PSNR, but slightly worse than Ucolor in terms of SSIM. One potential reason is that incorporating unlabeled real underwater images in training might emphasize the network to pay more attention to the real underwater scene. This can be confirmed by checking the quantitative results on testR. On testR, our method outperforms the other methods by a significant margin (outperforms the second best by 1.67dB in PSNR). In addition to the quantitative results, qualitative results are shown in Fig. 6. Our results are visually pleasant, while the compared methods suffer from color cast and over-enhancement.

Results on non-reference datasets. The quantitative results on UIEB, EUVP, RUIE and Seathru are shown in Table 2. By quickly checking throughout the table, we can observe that our method significantly outperform the compared method in MUSIQ. Besides, we also achieve competitive performance in terms of UIQM and UCIQE. However, as is noted in [3, 10, 22], UIQM and UCIQE might be biased to some characteristics and thus cannot accurately reflect the true visual quality of restored images. Similarly, the performance in MUSIQ is also for a reference. Therefore, the quantitative results might be insufficient to indicate the quality of restored underwater images due the underde-

Table 2. Evaluations of different methods on non-reference benchmarks in terms of UIQM, UCIQE and MUSIQ. Best results are in **bold** and the second best results are with underline.

Method	UIQM (higher, better)				UCIQE (higher, better)				MUSIQ (higher, better)			
	UIEB	EUVP	RUIE	Seathru	UIEB	EUVP	RUIE	Seathru	UIEB	EUVP	RUIE	Seathru
Input	3.066	4.729	3.948	5.925	0.509	0.517	0.490	0.537	41.70	42.73	33.53	60.25
GDCP [33]	3.401	4.738	4.509	5.343	0.564	0.599	0.565	0.590	40.07	42.49	34.63	60.54
MMLE [53]	4.283	4.723	4.967	5.555	<u>0.578</u>	<u>0.596</u>	<u>0.571</u>	0.620	<u>40.33</u>	<u>47.55</u>	<u>36.80</u>	<u>66.16</u>
WaterNet [22]	4.118	<u>5.317</u>	4.568	<u>6.829</u>	0.572	0.595	0.572	0.610	40.32	43.07	32.23	64.38
Ucolor [20]	3.894	<u>5.286</u>	4.426	<u>6.752</u>	0.542	0.566	0.534	0.594	40.08	41.81	33.66	64.44
PRWNet [16]	<u>4.371</u>	5.330	4.395	6.778	0.518	0.543	0.518	0.572	40.30	43.52	33.12	62.82
FGAN [17]	4.315	4.469	4.519	4.853	0.541	0.561	0.527	0.564	40.95	43.36	34.48	64.25
CWR [11]	4.133	5.152	4.469	6.067	0.587	<u>0.596</u>	0.565	<u>0.624</u>	38.46	41.46	31.25	64.21
Semi-UIR	4.598	5.291	<u>4.671</u>	6.846	0.587	0.593	0.557	0.632	43.77	51.66	37.87	66.61

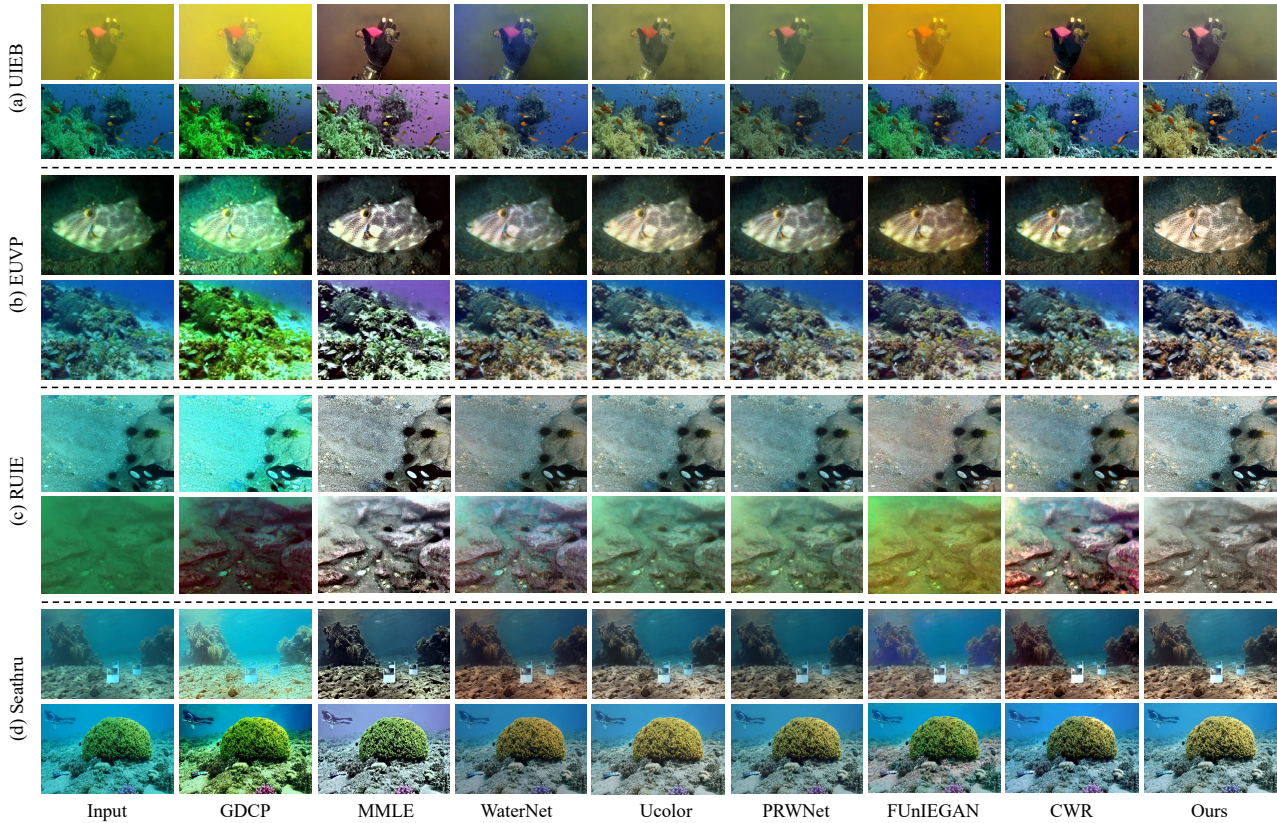


Figure 7. Visual comparisons on non-reference benchmarks UIEB [22], EUVP [17], RUIE [25] and Seathru [1]

veloped NR-IQA metrics. We further show the qualitative results on the four benchmarks in Fig. 7. Compared with other methods, our approach can robustly restore various types of underwater images with natural color and rich details. Under the guidance of reliable pseudo labels and powerful contrastive regularization, our framework generalizes well on various underwater scenes.

4.4. Ablation Study

To analyze the effectiveness of Semi-UIR, we conduct ablation studies to reveal the influence of the key compo-

nents in our method. They are presented as follows: (a) **Sup-base**: where we train the network AIM-Net without semi-supervised learning and unlabeled data. (b) **Semi-base**: Base semi-supervised training with consistency loss L_{un} . (c) **Semi-base+RB***: using reliable bank based on **Semi-base** without contrastive loss. (d) **Semi-base+CL***: adding contrastive loss to **Semi-base**, without using reliable bank. (e) **Semi-UIR**: our proposed Semi-UIR.

The quantitative results of above methods are shown in Table 3. We can observe that our full solution performs best. In addition, by comparing **Semi-base+RB*** with **Semi-base**




Figure 8. Results of ablation study about Semi-UIR.

Table 3. Ablation studies on testR, EUVP and UIEB benchmarks in terms of PSNR or MUSIQ. MT denotes mean teacher framework, CL represents contrastive loss, and RB is reliable bank.

Method	MT	RB	CL	testR	EUVP	UIEB
Sup-base				24.38	42.66	40.70
Semi-base	✓			23.11	42.48	40.33
Semi-base+RB*	✓	✓		24.53	43.27	42.16
Semi-base+CL*	✓		✓	23.97	46.59	40.64
Semi-UIR	✓	✓	✓	24.59	51.66	43.77

and Semi-UIR with **Semi-base+CL***, it is easy to verify the effectiveness of incorporating the reliable bank.

Besides, the qualitative results are shown in Fig. 8, where special attention should be paid on **Semi-base+CL*** **Semi-base+RB***. (1) In **Semi-base+CL***, without the reliable positive samples, the contrastive loss pushes the network to produce extremely different results than the negative samples (inputs). However, this unfortunately results in over-enhancement. (2) On the contrary, in **Semi-base+RB***, without the help of contrastive loss, the restored images still suffer from color distortion and are close to the degraded inputs. The two ablation studies verify the utility of the reliable bank and contrastive regularization.

4.5. Breakdown of the Training

To further illustrate that the teacher’s outputs can be used to train student network, we here provide some intermediate results during training. The results are shown in Fig. 9. At the beginning of the training (10 epochs), the teacher’s prediction is much better than that of student. As the training processes, the student’s outputs and teacher’s outputs are improved simultaneously.

4.6. Influence of Non-reference Metric

We here conduct experiments to show the influence of using different NR-IQA approaches in building our reliable bank. We conclude in Sec. 3.4 that MUSIQ is the most reliable one. To further demonstrate the correctness of this selection, we here show the final performance of using NIMA, PAQ2PIQ and MUSIQ on the labeled dataset, *i.e.* testS and testR. Table 4 shows the results. It can be observed that

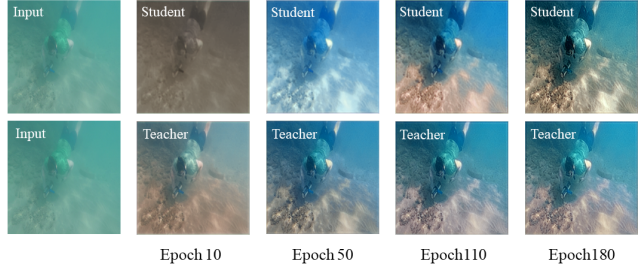


Figure 9. Examples of intermediate predictions of the teacher model and student model.

we can achieve the best performance by using MUSIQ. It also shows that using PAQ2PIQ is better than using NIMA, which is consistent with their reliability shown in Fig. 4.

Table 4. Evaluation the influence of adopting different NR-IQA metrics on testS and testR.

Method	Reliability	PSNR		SSIM	
		testS	testR	testS	testR
NIMA	41.05%	23.01	23.88	0.815	0.888
PAQ2PIQ	<u>82.11%</u>	<u>23.08</u>	<u>24.28</u>	<u>0.818</u>	<u>0.893</u>
MUSIQ	91.21%	23.40	24.59	0.821	0.901

4.7. Influence of Data Augmentation

We finally show the influence of using different data augmentations in Table 5. By comparing the method using data augmentation with baseline, it is easy to conclude that adopting any of the data augmentations is beneficial. Besides, using a mixture of the three strategies achieves the best performance.

Table 5. Evaluation of using different data augmentation. Baseline is our full solution without using the three strong data augmentations. Numbers are either in SSIM or MUSIQ.

Strategy	testR	UIEB	EUVP	RUIE	Seathru
Baseline	0.880	40.12	46.06	31.14	64.71
Color Jitter	0.889	40.31	49.16	33.66	64.87
Gaussian Blur	0.896	41.23	49.27	36.88	64.88
Gray Scale	0.895	40.61	47.57	32.51	65.19
All	0.901	43.77	51.66	37.87	66.61

5. Conclusion

We propose an efficient semi-supervised underwater image restoration method named Semi-UIR. As demonstrated the ablation experiments, the superior performance of the proposed method over other SOTA algorithms can be attributed to reliable teacher-student consistency and contrastive regularization. The follow-up research can be carried out in two directions: 1) extend the semi-supervised framework to cover other restoration tasks, 2) optimize memory usage during training and improve performance via memory management.

References

- [1] Derya Akkaynak and Tali Treibitz. Sea-thru: A method for removing water from underwater images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1682–1691, 2019. 6, 7
- [2] Codruta O Ancuti, Cosmin Ancuti, Christophe De Vleeschouwer, and Philippe Bekaert. Color balance and fusion for underwater image enhancement. *IEEE Transactions on image processing*, 27(1):379–393, 2017. 2
- [3] Dana Berman, Deborah Levy, Shai Avidan, and Tali Treibitz. Underwater single image color restoration using haze-lines and a new quantitative dataset. *IEEE transactions on pattern analysis and machine intelligence*, 43(8):2822–2837, 2020. 2, 4, 6
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3, 5
- [5] Zhixiang Chi, Xiao Shu, and Xiaolin Wu. Joint demosaicking and blind deblurring using deep convolutional neural network. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2169–2173. IEEE, 2019. 1
- [6] Zhixiang Chi, Yang Wang, Yuanhao Yu, and Jin Tang. Test-time fast adaptation for dynamic scene deblurring via meta-auxiliary learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9137–9146, 2021. 1
- [7] Minghan Fu, Huan Liu, Yankun Yu, Jun Chen, and Keyan Wang. Dw-gan: A discrete wavelet transform gan for nonhomogeneous dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 203–212, 2021. 1
- [8] Adrian Galdran, David Pardo, Artzai Picón, and Aitor Alvarez-Gila. Automatic red-channel underwater image restoration. *Journal of Visual Communication and Image Representation*, 26:132–145, 2015. 2
- [9] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 3, 5
- [10] Chunle Guo, Ruiqi Wu, Xin Jin, Linghao Han, Zhi Chai, Weidong Zhang, and Chongyi Li. Underwater ranker: Learn which is better and how to be better. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 2, 6
- [11] Junlin Han, Mehrdad Shoeiby, Tim Malthus, Elizabeth Botha, Janet Anstee, Saeed Anwar, Ran Wei, Mohammad Ali Armin, Hongdong Li, and Lars Petersson. Underwater image restoration via contrastive learning and a real-world dataset. *Remote Sensing*, 14(17):4297, 2022. 1, 3, 6, 7
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3, 5
- [13] Byeongho Heo, Sanghyuk Chun, Seong Joon Oh, Dongyoon Han, Sangdoo Yun, Gyuwan Kim, Youngjung Uh, and Jungwoo Ha. Adamp: Slowing down the slowdown for momentum optimizers on scale-invariant weights. In *International Conference on Learning Representations (ICLR)*, 2021. 5
- [14] Muhammad Suzuri Hitam, Ezmahamrul Afreen Awalludin, Wan Nural Jawahir Hj Wan Yussof, and Zainuddin Bachok. Mixture contrast limited adaptive histogram equalization for underwater image enhancement. In *2013 International conference on computer applications technology (ICCAT)*, pages 1–5. IEEE, 2013. 2
- [15] Hanzhe Hu, Fangyun Wei, Han Hu, Qiwei Ye, Jinshi Cui, and Liwei Wang. Semi-supervised semantic segmentation via adaptive equalization learning. *Advances in Neural Information Processing Systems*, 34:22106–22118, 2021. 3
- [16] Fushuo Huo, Bingheng Li, and Xuegui Zhu. 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*, pages 1944–1952, 2021. 2, 5, 6, 7
- [17] 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. 1, 2, 4, 6, 7
- [18] Aupendu Kar, Sobhan Kanti Dhara, Debashis Sen, and Prabir Kumar Biswas. Zero-shot single image restoration through controlled perturbation of koschmieder’s model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16205–16215, 2021. 2
- [19] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5148–5157, 2021. 2, 4, 5, 6
- [20] Chongyi Li, Saeed Anwar, Junhui Hou, Runmin Cong, Chunle Guo, and Wenqi Ren. Underwater image enhancement via medium transmission-guided multi-color space embedding. *IEEE Transactions on Image Processing*, 30:4985–5000, 2021. 1, 2, 6, 7
- [21] Chongyi Li, Saeed Anwar, and Fatih Porikli. Underwater scene prior inspired deep underwater image and video enhancement. *Pattern Recognition*, 98:107038, 2020. 1, 6
- [22] 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, 2020. 1, 2, 4, 6, 7
- [23] Dong Liang, Ling Li, Mingqiang Wei, Shuo Yang, Liyan Zhang, Wenhan Yang, Yun Du, and Huiyu Zhou. Semantically contrastive learning for low-light image enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1555–1563, 2022. 3, 5
- [24] Huan Liu, Zijun Wu, Liangyan Li, Sadaf Salehkalaibar, Jun Chen, and Keyan Wang. Towards multi-domain single image dehazing via test-time training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2022. 1
- [25] Risheng Liu, Xin Fan, Ming Zhu, Minjun Hou, and Zhongxuan Luo. Real-world underwater enhancement: Chal-

- lenges, benchmarks, and solutions under natural light. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(12):4861–4875, 2020. 6, 7
- [26] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4258–4267, 2022. 3
- [27] Ye Liu, Lei Zhu, Shunda Pei, Huazhu Fu, Jing Qin, Qing Zhang, Liang Wan, and Wei Feng. From synthetic to real: Image dehazing collaborating with unlabeled real data. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 50–58, 2021. 3, 6
- [28] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012. 5
- [29] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. 5
- [30] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2
- [31] 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. 4, 5, 6
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [33] Yan-Tsung Peng, Keming Cao, and Pamela C Cosman. Generalization of the dark channel prior for single image restoration. *IEEE Transactions on Image Processing*, 27(6):2856–2868, 2018. 2, 6, 7
- [34] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992. 3
- [35] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 2
- [36] Yoav Y Schechner and Nir Karpel. Clear underwater vision. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. 2
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3, 5
- [38] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 2, 3
- [39] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE transactions on image processing*, 27(8):3998–4011, 2018. 5
- [40] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [41] Keyan Wang, Yan Hu, Jun Chen, Xianyun Wu, Xi Zhao, and Yunsong Li. Underwater image restoration based on a parallel convolutional neural network. *Remote sensing*, 11(13):1591, 2019. 2
- [42] Lin Wang and Kuk-Jin Yoon. Semi-supervised student-teacher learning for single image super-resolution. *Pattern Recognition*, 121:108206, 2022. 3
- [43] Yanbo Wang, Shaohui Lin, Yanyun Qu, Haiyan Wu, Zhizhong Zhang, Yuan Xie, and Angela Yao. Towards compact single image super-resolution via contrastive self-distillation. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1122–1128. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track. 5
- [44] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4248–4257, 2022. 3
- [45] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [46] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Contrastive learning for compact single image dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10551–10560, 2021. 3, 5
- [47] Miao Yang and Arcot Sowmya. An underwater color image quality evaluation metric. *IEEE Transactions on Image Processing*, 24(12):6062–6071, 2015. 4, 5, 6
- [48] Tian Ye, Sixiang Chen, Yun Liu, Yi Ye, Erkang Chen, and Yuche Li. Underwater light field retention: Neural rendering for underwater imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 488–497, June 2022. 1
- [49] Xiangyu Yin, Xiaohong Liu, and Huan Liu. Fmsnet: Underwater image restoration by learning from a synthesized dataset. In *Artificial Neural Networks and Machine Learning—ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part III 30*, pages 421–432. Springer, 2021. 1
- [50] Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, and Alan Bovik. From patches to

- pictures (paq-2-piq): Mapping the perceptual space of picture quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3575–3585, 2020. [5](#)
- [51] Yankun Yu, Huan Liu, Minghan Fu, Jun Chen, Xiyao Wang, and Keyan Wang. A two-branch neural network for non-homogeneous dehazing via ensemble learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 193–202, 2021. [1](#)
- [52] Shu Zhang, Ting Wang, Junyu Dong, and Hui Yu. Underwater image enhancement via extended multi-scale retinex. *Neurocomputing*, 245:1–9, 2017. [2](#)
- [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. [2](#), [6](#), [7](#)
- [54] Huabo Zhu, Xu Han, and Yourui Tao. Semi-supervised advancement of underwater visual quality. *Measurement Science and Technology*, 32(1):015404, 2020. [3](#)
- [55] Xiaojin Jerry Zhu. Semi-supervised learning literature survey. 2005. [2](#)