

Test Time Adaptation for Blind Image Quality Assessment

Subhadeep Roy* Shankhanil Mitra* Soma Biswas Rajiv Soundararajan
Indian Institute of Science, Bengaluru, India

subhadeepro2000@gmail.com, {shankhanilm, somabiswas, rajivs}@iisc.ac.in

Abstract

While the design of blind image quality assessment (IQA) algorithms has improved significantly, the distribution shift between the training and testing scenarios often leads to a poor performance of these methods at inference time. This motivates the study of test time adaptation (TTA) techniques to improve their performance at inference time. Existing auxiliary tasks and loss functions used for TTA may not be relevant for quality-aware adaptation of the pre-trained model. In this work, we introduce two novel quality-relevant auxiliary tasks at the batch and sample levels to enable TTA for blind IQA. In particular, we introduce a group contrastive loss at the batch level and a relative rank loss at the sample level to make the model quality aware and adapt to the target data. Our experiments reveal that even using a small batch of images from the test distribution helps achieve significant improvement in performance by updating the batch normalization statistics of the source model.

1. Introduction

The problem of image quality assessment (IQA) is extremely important in diverse image capture, processing, and sharing applications. However, a reference image is often not available for quality assessment. No reference (NR) or blind IQA primarily deals with the question of predicting image quality without using a reference image. Such NR IQA algorithms are often designed using machine learning approaches. More recently, deep learning based approaches have been extremely successful in achieving impressive performance. However, IQA applications are quite diverse and deal with several different distortions and distributional shifts. IQA models often have poor generalization ability and find it difficult to perform well under such shifts.

Test time adaptation (TTA) has emerged as an important approach to address distributional shifts at test time [8]. It has been shown that by modifying a few global parameters of the model using a suitable loss that does not require the

ground truth, one can significantly improve the performance of the model on the test data. Further, source-free adaptation, where the source data on which the original model was trained is not available while updating the model, is a realistic setting. While such approaches have been studied extensively in image classification literature [31, 33], there is hardly any literature on TTA for IQA.

There are multiple challenges in designing TTA for IQA. Typical losses used for TTA, such as entropy minimization [33], are not applicable for IQA. For example, IQA is often studied in the regression context. This makes it difficult to extend models based on class confidences [24] or class prototypes [10] for classification to IQA. Also, the relevance of other self-supervised tasks such as rotation prediction [14], context prediction [5], colorization [15], noise prediction [1], feature clustering [3] for adapting IQA models is not clear. While contrastive learning has also been employed for TTA [20], such a framework is not explicitly based on contrasting image quality, and its relevance is also not clear.

Our main contribution is in the design of auxiliary tasks to enable TTA for IQA. We start with a source model trained on a large IQA dataset and fine-tune the model on individual batches of test samples. The first task we introduce for adaptation is based on contrasting groups of low and high quality images in a batch. Thus, we exploit the initial knowledge of the source model and try to adapt it by enforcing quality relationships among the batch of samples. Such a group contrastive (GC) learning approach fits naturally to our setting to account for any errors on individual samples that the source model may be prone to.

In contrast to the GC learning that depends on the batch, our second auxiliary task is an image specific task based on distorted augmentations of different types. Here, our goal is to enable the model to rank the image quality of further distorted versions of each test sample. We explore the role of different distortion types to leverage the maximum benefit of this task. While GC learning is more effective when samples in a batch are diverse in quality, the rank order based learning is more effective when the quality of the images is not extremely poor. Thus, a combination of the tasks

* Authors contributed equally to this work

§ <https://github.com/subhadeepro2000/TTA-IQA>

helps overcome the shortcoming of both tasks and leads to an overall superior performance.

We study the TTA problem under different settings of source and target datasets for multiple state of the art IQA models. Our results show significant improvements of the source model and the importance of TTA for IQA. We summarize our main contributions as follows:

- We propose source-free test time adaptation techniques in the context of blind image quality assessment to mitigate distribution shifts between train and test data.
- We formulate two novel quality-aware self-supervised auxiliary tasks to adapt to the test data distribution. While group contrastive learning helps capture quality discriminative information among several images within a batch, rank ordering helps maintain the quality order between two different distorted versions of the same image.
- We show that our TTA method can significantly improve the performance of four different quality-aware source models, each on four different test IQA databases.

2. Related Work

2.1. Test Time Adaptation

One of the first pieces of work on TTA [31] introduces a joint training framework using a loss for the main task and a self-supervised auxiliary task loss. The choice of a relevant auxiliary task is a challenging part of the design. Sun *et al.* [31] use rotation prediction as a pre-text task for image classification applications. Researchers also explore simpler tasks that are highly correlated with the main task, such as feature alignment with the batch and a simple contrastive loss framework [20] for adaptation. However, these methods require the source data to train the source model all over again to enable TTA.

TENT [33] studies source-free TTA, where entropy minimization-based adaptation even outperforms some methods that use the source data. Moreover, only the batch normalization parameters of the model are adapted in TENT. There are several works on batch norm statistics adaptation [25, 28] to improve the robustness of the model at test time. SHOT [16] presents a clustering-based pseudo-labeling method to align features from the target domain to the source domain using an information maximization loss. While a plethora of methods exists as above, the auxiliary tasks in these methods are not relevant for IQA.

2.2. No-Reference Image Quality Assessment

Most classical NR IQA algorithms are mainly based on natural scene statistics (NSS) [19, 22, 38]. In [23], the naturalness of the distorted image in the wavelet domain is modeled based on NSS. Saad *et al.* [27] also design the NSS model in the discrete cosine transform (DCT) domain. CORNIA [35], and HOSA [34] are among the earliest codebook learning based methods to predict quality. With the emergence of deep learning and the availability of large subject-rated IQA databases, various general-purpose NR IQA methods have been designed based on convolutional neural network (CNN) architectures [2, 18, 21, 36]. Some of these methods require end-to-end training [2, 21] of deep neural networks (DNN), while others [7, 12, 30, 39] are based on updating pre-trained models with some modifications.

More recently, Zhang *et al.* [39] propose a method to train bilinear DNNs to simultaneously model both authentic and synthetic distortions. Su *et al.* [30] design a self-adaptive hyper network to provide weights for the quality prediction module parameters and handle various types of distortions and content in the images. Several methods [7, 12] explore transformer-based architectures along with a CNN to capture dependencies between local and global features. MetaIQA [40] proposes meta-learning on synthetic distortions by using a shared quality prior knowledge model to adapt to any kind of distortion.

All the existing methods assume that the train and test data come from the same distribution. If there is a distribution shift across different databases, we need adaptation of the pre-trained model to learn about the target distributional information.

3. Methodology

We propose a novel self-supervised Test Time Adaptation technique for Image Quality Assessment (**TTA-IQA**) to adapt pre-trained quality models and mitigate distribution shifts between the source and target data. We consider source-free TTA, where we only have access to the pre-trained quality-aware models and no access to the source training data. When a batch of test data $D = \{x_j\}_{j=1}^n$ arrives, we adapt the model using the batch without knowledge of corresponding ground truth $\{y_j\}_{j=1}^n$.

3.1. Approach

Let the model trained on the source data be f_θ where $\theta = (\theta_e, \theta_c)$ corresponds to the parameters of the network, and θ_e and θ_c correspond to the parameters of the feature extractor and regression layers. Thus, for an input image x , we denote $f_\theta(x) = f_{\theta_c}(f_{\theta_e}(x))$. Since our goal is to learn the distribution shift between train and test data, we only update the parameters of the feature extractor, θ_e , to

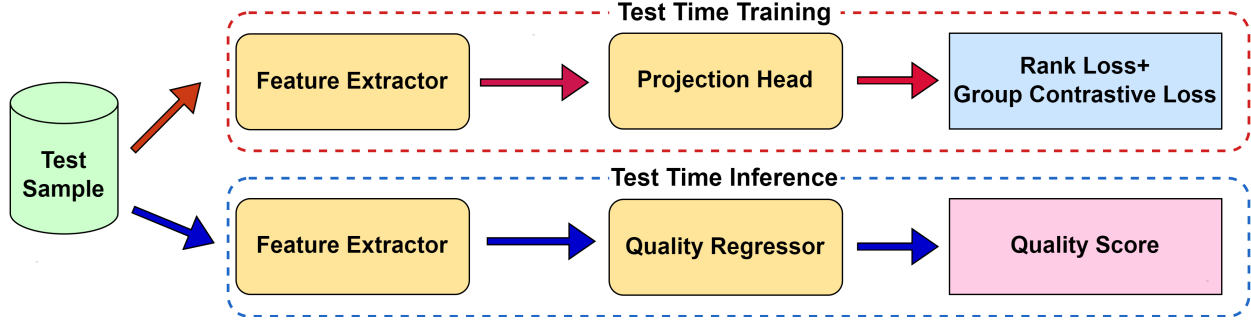


Figure 1. **Block diagram of a general architecture for test time adaptation.** At test time training, the normalization layers of the feature extractor are adapted by optimizing the combination of rank and GC loss. At inference time, we predict the quality scores of test images using the updated feature extractor and pre-trained quality regressor.

align the features between the train and test distributions in a lower dimensional space.

The key challenge in TTA is the choice of a self-supervised auxiliary task that is highly correlated with the main task of IQA. However, relying too much on the auxiliary task can affect the performance of the main task. To prevent loss of learnt information induced by the auxiliary task [4], we first project the features to a lower dimensional space using a non-linear projection head f_{θ_s} , parameterized by θ_s . The auxiliary task drives the adaptation of the feature extractor through the projection head. Thus our model now has three parts parameterized by $(\theta_e, \theta_s, \theta_c)$ to resemble a Y-shape as shown in Figure 1.

When a batch of test instances D arrives, we extract features from the feature extractor, followed by the projection head, and update the set of parameters (θ_e, θ_s) by optimizing a self-supervised objective function $\mathcal{L}_s(D)$. Updating all the model parameters of the feature extractor θ_e can cause the model to diverge too much from training, and the performance can drop drastically. Inspired by prior work on test time entropy minimization [33] and improving robustness for test data [28], we only update the linear and lower-dimensional feature modulation parameters. In a neural network, normalization layers satisfy these properties. So we only adapt the batch normalization layers by updating the affine parameters to mitigate distributional shifts. Thus we update these parameters by optimizing the auxiliary task loss given by

$$\theta_e^*, \theta_s^* = \underset{(\theta_e, \theta_s)}{\operatorname{argmin}} \mathcal{L}_s(D) \quad (1)$$

After optimizing the above loss function, we use the updated feature extractor parameters θ_e^* and pre-trained quality regressor parameters θ_c to predict the quality for a batch of test data. Since the distribution across batches can vary significantly, we ignore the earlier updated model and start the TTA based on source model weights for new test instances. Thus the updated target model always only depends

on the incoming test data.

3.2. Self-Supervised Auxiliary Tasks

Our goal is to carefully choose self-supervised auxiliary tasks that capture quality-aware distributional information to adapt the feature encoder. We formulate two novel and complementary self-supervised learning techniques which help to learn the distribution shift between train and test data. These self-supervised objective functions are - 1) Group Contrastive Loss and 2) Rank Loss. While the GC loss works well when there is a reasonable separation of quality among the samples in a batch, the rank loss works better even when the quality of the batch samples is similar. On the other hand, the rank loss is meaningful only when the quality of the input image is not extremely low, while the GC loss is independent of the quality of a given image. Thus, a combination of the two losses renders our TTA extremely effective across various scenarios.

3.2.1 Group Contrastive Loss

While contrastive learning has been used for TTA of deep image classification, its direct application does not appear relevant to the task of IQA. Thus, we introduce group contrastive (GC) learning as an auxiliary task for TTA of IQA models. In particular, we make two groups of images from a single batch of N images based on the pseudo-labels given by the pre-trained source model. We sort the images in ascending order as $x_{(1)}, x_{(2)}, \dots, x_{(N)}$ based on the pseudo-labels, where $x_{(i)}$, $i = 1, 2, \dots, N$, corresponds to the i^{th} lowest quality image in the batch. We then segregate images with high quality scores (say, top p fraction of the data in a batch) and include them in a group of higher-quality images. Similarly, we separate out lower quality images (say, the lowest p fraction of the images in a batch) together to form another group. Here we assume that pN and $(1-p)N$ are integers for simplicity; else, they can be rounded off to the nearest integers.

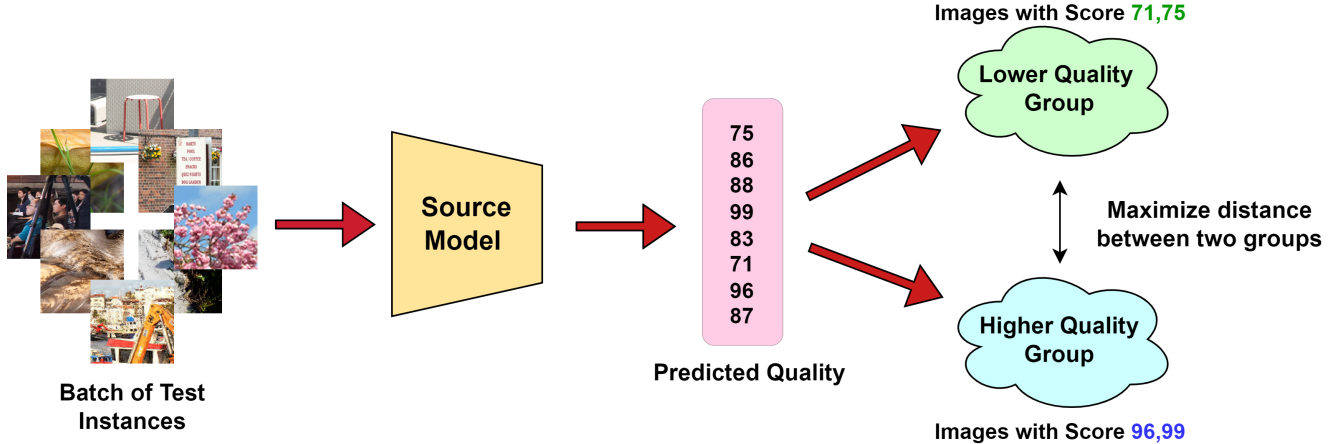


Figure 2. **The overview framework of group contrastive loss for test time adaptation.** For a batch of images, two different groups are formed based on pseudo-labels given by the source model. The group contrastive loss tries to minimize the distance between features extracted from images belonging to the same group while maximizing the distance between features from different groups.

The premise behind our loss for GC learning is that images from the same quality group should give similar feature representations in a lower dimensional space while features of images from different groups are separated out. Thus image pairs from the same group act as positive pairs, and image pairs from different groups act as negative pairs. By separating out these two groups, the model adapts itself by better separating the intermediate quality samples.

Let a positive pair $x_{(i)}$ and $x_{(j)}$ come from the same group, i.e., either both $i, j \leq pN$ or $i, j > (1-p)N$. We use a modified NT-Xent contrastive loss [4] as our objective function. For a pair of images coming from the lower quality group where $i, j \leq pN$ and $i \neq j$, the GC loss is defined as

$$\mathcal{L}_{i,j}^{gc} = -\log \frac{\exp(\text{sim}(\mathbf{z}_{(i)}, \mathbf{z}_{(j)})/\tau)}{\sum_{k>(1-p)N} \exp(\text{sim}(\mathbf{z}_{(i)}, \mathbf{z}_{(k)})/\tau)}, \quad (2)$$

where $\mathbf{z}_{(i)} = f_{\theta_s}(f_{\theta_e}(x_{(i)}))$ represents the feature at the output of the projection head for sample $x_{(i)}$, sim refers to the cosine similarity between two features. and $\mathbf{z}_{(k)}$ is the feature extracted from an image from the higher quality group with $k > (1-p)N$. Also, τ represents temperature scaling parameter. While we define the loss above when $i, j \leq pN$, we can define a similar loss when $i, j > (1-p)N$. For all pairs of images within the same group, we obtain the GC loss and add them together to obtain the final loss, \mathcal{L}^{gc} as

$$\mathcal{L}^{gc} = \sum_{i=1}^{pN} \sum_{\substack{j=1 \\ j \neq i}}^{pN} \mathcal{L}_{i,j}^{gc} + \sum_{i=(1-p)N+1}^N \sum_{\substack{j=(1-p)N+1 \\ j \neq i}}^N \mathcal{L}_{i,j}^{gc}. \quad (3)$$

A block diagram of the GC loss is shown in Figure 2.

3.2.2 Rank Loss

Given a sample of images from test data, the rank loss helps adapt the features by capturing quality-aware distributional information at the sample level. We introduce a ranking objective to learn the quality orders between two distorted versions of the test images that are quality discriminable. We distort each test image $x_i, i = 1, 2, \dots, N$, in the minibatch of size N with two different degrees of degradation of a given distortion type. Let x_i^{high} and x_i^{low} denote the higher and lower distorted image, respectively. The degradation types include synthetic distortions such as blur, compression, and noise. We choose degradation levels randomly from two sets of parameters sufficiently farther apart. Ideally, the distance between the features extracted from the test image and the highly distorted image should be greater than that between the features extracted from the test image and the lower distorted image. Our rank loss tries to capture this order and adapt the model to the test data.

We project the triplet of images $(x_i, x_i^{high}, x_i^{low})$ in the feature space through the feature extractor and the projection head. We measure the distance between the distortion-augmented image and the original test data using the Euclidean distance in the projected feature space. Mathematically, let $\mathbf{z}_i = f_{\theta_s}(f_{\theta_e}(x_i))$ be the feature extracted for test sample x_i . Similarly we define \mathbf{z}_i^{high} and \mathbf{z}_i^{low} . Let d_i^{high} denote the distance between \mathbf{z}_i and \mathbf{z}_i^{high} . Similarly, d_i^{low} denotes the distance between \mathbf{z}_i and \mathbf{z}_i^{low} . The target model is now fine-tuned to obtain the correct ranking between the distances by achieving $d_i^{high} \geq d_i^{low}$ shown in Figure 3.

The probability of achieving this order is estimated by passing the difference in distances through a sigmoid func-

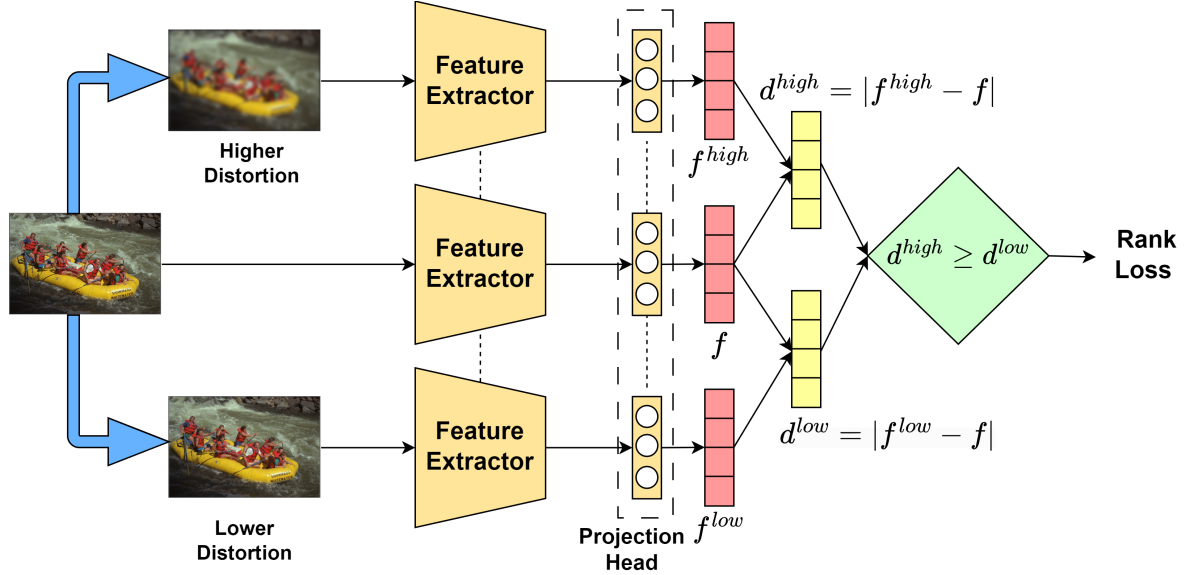


Figure 3. **The rank loss framework for test time adaptation.** We distort the test image with two different degrees of degradation. We then project the triplet of images in the features space through the shared feature extractor and projection head. Distances calculated between the features extracted from each distorted image and the original test image are used to compare their rank order.

tion as

$$P_i = \Pr(d_i^{high} \geq d_i^{low}) = \frac{\exp(d_i^{high} - d_i^{low})}{1 + \exp(d_i^{high} - d_i^{low})}. \quad (4)$$

We use binary cross-entropy loss between true label $\bar{P}_i = 1$ and the predicted probability P_i to obtain the rank loss as

$$\mathcal{L}_i^r(\bar{P}_i, P_i) = -\bar{P}_i \log P_i - (1 - \bar{P}_i) \log(1 - P_i). \quad (5)$$

For a batch of size N , the overall rank loss is given as

$$\mathcal{L}^r = \sum_{i=1}^N \mathcal{L}_i^r. \quad (6)$$

One of the challenges in the rank loss is selecting the distortion type for every image. The original test image may already be distorted, which poses difficulties in deciding the distortion type that can help TTA. For example, if the test image is extremely blurred and we blur it again with two different levels, both images may be visually indistinguishable, thereby limiting the extent to which adaptation can help. We exploit the source model’s knowledge to overcome this limitation in choosing the distortion type. We predict the quality scores of the pair (x_i^{high}, x_i^{low}) for each distortion type using the source model. We hypothesize that the pair with the maximal difference in predicted quality is sufficiently different in visual quality. Thus choosing such a pair in Equation (5) can help adapt the model.

Our overall self-supervised TTA objective function is a

combination of both the rank and GC loss given as

$$\mathcal{L}_s = \mathcal{L}^{gc} + \lambda \mathcal{L}^r, \quad (7)$$

where λ is a hyper-parameter used to combine the losses.

4. Experiments

4.1. Quality Models, Datasets and Metrics

We evaluate TTA on four popular IQA databases using four different quality-aware models. In particular, we consider state-of-the-art deep IQA models such as TReS [7], MUSIQ [12], HyperIQA [30], and MetaIQA [40]. These models contain a ResNet backbone with batch normalization layers, which we model as the feature extractor f_{θ_e} and only update its batch normalization parameters. We model the rest of the network as the quality regressor part, f_{θ_c} . TReS and MUSIQ use transformers as a part of their architecture, which we include as a part of the quality regressor in all our main experiments. We also explore the adaptation of transformers by including them as part of the feature extractor and updating their layer normalization statistics in the supplementary material.

We project the features extracted from the last layer of ResNet through a 256-dimensional fully connected (FC) layer corresponding to the self-supervised projection head f_{θ_s} . These lower dimensional features are used to adapt the model at test time. Three IQA models, TReS, MUSIQ, and HyperIQA are trained on the LIVEFB database [37] containing 39,810 images. MetaIQA is trained on two synthetically distorted databases, TID2013 [26] and KADID-10k [17].

Backbone	Database	KonIQ-10k		PIPAL		CID2013		LIVE-IQA	
	Method	SROCC	PLCC	SROCC	PLCC	SROCC	PLCC	SROCC	PLCC
TReS	Baseline	0.6520	0.6955	0.3845	0.4078	0.5272	0.6463	0.5435	0.4450
	Rotation	0.6506	0.6805	0.4061	0.4114	0.5706	0.6651	0.5866	0.5311
	TTA-IQA	0.6578	0.7074	0.4278	0.4204	0.6032	0.6710	0.6722	0.5963
MUSIQ	Baseline	0.6304	0.6802	0.3190	0.3414	0.5173	0.6032	0.2596	0.3351
	Rotation	0.6577	0.7154	0.3665	0.3693	0.5487	0.6164	0.3512	0.3976
	TTA-IQA	0.6693	0.7230	0.3743	0.3731	0.5499	0.6220	0.3649	0.4031
HyperIQA	Baseline	0.5861	0.6313	0.3037	0.3304	0.4895	0.6123	0.5143	0.4377
	Rotation	0.6033	0.6536	0.3268	0.3482	0.4902	0.6150	0.6268	0.5469
	TTA-IQA	0.5960	0.6495	0.3653	0.3767	0.5039	0.5988	0.6218	0.5438
MetaIQA	Baseline	0.5162	0.4460	0.3287	0.2955	0.7213	0.6817	0.7323	0.6732
	Rotation	0.5823	0.5311	0.3353	0.3042	0.7177	0.6745	0.7271	0.6868
	TTA-IQA	0.5838	0.5428	0.4073	0.3510	0.7809	0.7399	0.7999	0.7726

Table 1. Comparison of TTA-IQA with popular NR IQA methods and one popular auxiliary task - rotation prediction on authentically and synthetically distorted datasets. Bold entries imply the best performance for every individual quality-aware model on respective datasets.

We choose challenging databases to evaluate the generalization capability of our TTA-IQA method. The test datasets are described as follows:

KonIQ-10k [9] is a popular in-the-wild authentically distorted database consisting of 10,073 quality-scored images.

PIPAL [11] is a large-scale IQA database for evaluating perceptual image restoration. It contains 29k distorted images, including 19 different GAN-based algorithms.

CID2013 [32] consists of 480 images in six image sets captured by 79 imaging devices.

LIVE-IQA [29] consists of 779 synthetically distorted images from 29 reference images.

We evaluate our results using Spearman’s rank-order correlation coefficient (SROCC) and Pearson’s linear correlation coefficient (PLCC).

4.2. Implementation Details

We implement our setup in PyTorch and conduct all the experiments with an 16 GB NVIDIA RTX A4000 GPU. During TTA, we randomly select a patch of size 224×224 from the input image and apply quality preserving augmentations such as horizontal flip and vertical flip before passing it through the network for TTA. We use the ADAM [13] optimizer with a learning rate of 0.001 and set the number of iteration as 3. All the experiments were run on five different seeds using a batch size of 8, and the final results were obtained by averaging.

With regard to the self-supervised auxiliary tasks, we use a combination of rank loss and GC loss as our objective function using $\lambda = 1$ given in Equation (7). We choose $p = 0.25$ to determine the groups. We calculate the GC loss using $\tau = 1$. For the rank loss, the test image is synthetically blurred using a Gaussian blur filter of size 5×5 with two sets of standard deviations σ , for the Gaussian

blur kernel. We keep $\sigma \in [40, 80]$ for highly distorted images and $\sigma \in [1, 20]$ for less blurred images. For compression distortions, we specify the quality factor in [80, 95] for lower compression and [30, 60] for higher compression rates. Similarly, we add zero mean Gaussian white noise to the test image with variance in [0.05, 0.1] for higher distortion and in [0.005, 0.01] for lower distortion.

4.3. Performance Evaluation

Table 1 shows the performance of our method on all four test datasets using all four quality-aware models. In addition to the source model, we also compare with rotation prediction [31] as the auxiliary task during TTA. A comparison with such a task helps understand the role of quality-aware losses for the TTA of IQA models.

We observe that TTA-IQA using the combination of rank and GC loss outperforms the source models in all the cases. Note that the PIPAL dataset has a huge distribution shift from the authentically distorted LIVEFB dataset on which three of the models were trained. While most source models perform very poorly on PIPAL, we achieve 10%-20% improvement over the source models. On the KonIQ-10k dataset, TTA-IQA gives around 1%-10% improvement over the source model. As both KonIQ-10k and LIVEFB are authentically distorted datasets, the distribution shift is reasonably small, leading to smaller improvements using adaptation. CID2013 is also an authentically distorted dataset and gives similar improvements of around 2%-13% over the source models. Our experiments on the LIVE-IQA database provide a significantly greater improvement of 10%-33% owing to the shift from authentic to synthetic distortions for three models.

We also observe from Table 1 that our approach outperforms the rotation prediction task in most cases. Other TTA ideas, such as TENT [33] and training using masked autoen-

Backbone	Database		KonIQ-10k		PIPAL		CID2013		LIVE-IQA	
	Rank	GC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
TReS	✓	×	0.6562	0.6989	0.4171	0.4204	0.6016	0.6736	0.6705	0.5908
	×	✓	0.6516	0.6946	0.4666	0.5183	0.5366	0.6493	0.7160	0.6290
	✓	✓	0.6578	0.7074	0.4278	0.4204	0.6032	0.6710	0.6722	0.5963
MUSIQ	✓	×	0.6549	0.7149	0.3768	0.3718	0.5216	0.6034	0.3634	0.4024
	×	✓	0.6611	0.7176	0.3585	0.3642	0.5446	0.6159	0.3299	0.3914
	✓	✓	0.6693	0.7230	0.3743	0.3731	0.5499	0.6220	0.3649	0.4031
HyperIQA	✓	×	0.5928	0.6455	0.3616	0.3732	0.5039	0.5991	0.5505	0.5050
	×	✓	0.6094	0.6567	0.3333	0.3552	0.5120	0.6255	0.6331	0.5592
	✓	✓	0.5960	0.6495	0.3653	0.3767	0.5039	0.5988	0.6218	0.5438
MetaIQA	✓	×	0.5580	0.5118	0.3992	0.3437	0.7579	0.7282	0.7894	0.7534
	×	✓	0.5414	0.4636	0.3710	0.3287	0.7861	0.7067	0.7566	0.6927
	✓	✓	0.5838	0.5428	0.4073	0.3510	0.7809	0.7399	0.7999	0.7726

Table 2. Ablation study results on authentically and synthetically distorted datasets using popular quality aware source models. Bold entries imply the best performance among all three settings.




Figure 4. Scatter plot for predicting quality of similar quality images from LIVE-IQA dataset for the source model and TTA using only rank loss or GC loss. We mark two sets of images having similar human opinion score (DMOS) by ellipses in each case.

coders [6] are not applicable for the IQA task. In particular, the notion of entropy in regression tasks is not clear. On the other hand, masked reconstructions tend to lead to quality degradation. Thus, we do not compare with them.

4.4. Ablation Study & Other Experiments

Need for Rank Loss and GC Loss for TTA. We perform an ablation study with respect to the two losses in Table 2. We see from the results that the rank loss and the GC loss individually always improve on the source models. Further, the combination of the rank and GC loss provides an even better performance over the individual losses in most evaluation scenarios. Even in scenarios, where the combination loss achieves the second-best performance, its performance is very close to the best. For the rest of the experiments in this section, we present results on all four datasets using the TReS method alone.

We now discuss scenarios where the GC loss is particularly more useful than the rank loss. If the test image is extremely distorted, none of the three kinds of distortion types can create much difference in the perceptual quality of the two degraded versions. To understand this further, we consider highly distorted images (corresponding to a dif-

ferential mean opinion score (DMOS) greater than 70) of different distortion types in the LIVE-IQA dataset. We apply TTA only for these images using different losses. The SROCC performance for these images in Figure 5 reveal that the rank loss is not very effective. However, the GC loss works well in such scenarios.

Conversely, we also discuss scenarios where the rank loss is more useful than the GC loss. To illustrate this idea, we select multiple images having similar quality with DMOS in [28, 44] from the LIVE-IQA dataset. We adapt the source model for this set of images using both the GC and rank losses. From Figure 4, we observe that the rank loss leads to much better adaptation in terms of the resulting model correlating with the ground truth scores. Intuitively, when the input images have a similar quality, the pseudo-labels given by the source model may be noisy, leading to an inaccurate grouping of the images into the lower and higher quality groups. This can adversely affect the GC loss. So the rank loss is more effective than the GC loss for test batches with similar quality images.

Effect of Selecting Multiple Distortion Types in the Rank Loss. We experiment with the impact of different kinds of distortion types while incorporating the rank loss.




Figure 5. TTA-IQA on different distortion-specific images from LIVE-IQA dataset with respect to individual losses

In particular, we consider three scenarios: one where only one of the distortion types is used in the rank loss, another where all the three types are used to obtain three loss terms which are summed together to update the model, and a third, where only one of the distortion types is chosen based on the discussion in Section 3.2.2. We observe from Table 3 that choosing the distortion type with the maximal difference in quality gives the best performance for the rank loss. This experiment validates our hypothesis that only a rank loss that can discriminate between the degraded image pairs is useful for TTA.

Database	Blur	Comp	Noise	All	Best
KonIQ-10k	0.656	0.656	0.667	0.615	0.656
PIPAL	0.384	0.415	0.389	0.415	0.417
CID2013	0.527	0.609	0.606	0.560	0.609
LIVE-IQA	0.578	0.605	0.618	0.634	0.671

Table 3. Impact of selecting multiple distortion types on rank loss

Selection of Group Size by Varying p . We explore different values of p for constructing the GC loss. For a batch size of 8, possible choices of the group size are 2, 3, 4 corresponding to $p = 0.25, 0.375, 0.5$ respectively. From Table 4, we observe that the performances are roughly similar for different values of p with a slightly superior performance for $p = 0.25$. We note that as p increases, the groups are larger and probably less discriminative in terms of quality, leading to a slightly poorer performance.

Database	$p=0.25$	$p=0.375$	$p=0.5$
KonIQ-10k	0.6516	0.6497	0.6467
PIPAL	0.4666	0.4602	0.4616
CID2013	0.5366	0.5486	0.5536
LIVE-IQA	0.7160	0.7171	0.6937

Table 4. Impact of varying group size using p for GC loss

Selection of Number of Groups for GC Loss. In our formulation, we define the GC loss only between two con-

trastive groups. In principle, we can extend our idea of GC loss to multiple groups. In particular, we can cluster the images in a batch into multiple groups and apply the GC loss between every pair of groups and sum the loss terms. Thus, images coming from the same group act as positive pairs, and images from two different groups are considered as negative pairs. From Table 5, we observe that the number of groups G does not impact the performance much.

Database	$G = 2$	$G = 3$	$G = 4$
KonIQ-10k	0.6516	0.6524	0.6509
PIPAL	0.4666	0.4668	0.4663
CID2013	0.5366	0.5327	0.5388
LIVE-IQA	0.7160	0.7011	0.6883

Table 5. Impact of varying number of groups G for GC loss

Choice of Number of Iterations for Learning Auxiliary Task. We also examine the effect of the number of iterations of parameter updates during TTA. Figure 6 shows that the performance on CID2013 dataset is fairly robust until 4 iterations. Beyond that, the model overfits the auxiliary task and leads to poor performance at test time.




Figure 6. Effect of increasing number of iterations

5. Conclusion

Test time adaptation has become very popular due to its simplicity and the lack of need for end-to-end training. Our work is perhaps one of the first attempts to design the method of TTA in the context of blind IQA. While most IQA methods focus on making the models robust enough to perform well for cross-database experiments, our TTA-IQA method can outperform existing state-of-the-art methods because of the adaptation at test time. We formulate novel self-supervised auxiliary tasks using the rank and group contrastive losses, which can learn quality-aware information from the test data. While primarily explored TTA for IQA, it would be interesting to understand the role of TTA for video quality assessment as well.

Acknowledgement: This work was supported in part by a grant from Kotak IISc AI-ML Centre and Department of

References

- [1] Piotr Bojanowski and Armand Joulin. Unsupervised learning by predicting noise. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 517–526. PMLR, 06–11 Aug 2017. [1](#)
- [2] Sebastian Bosse, Dominique Maniry, Klaus-Robert Müller, Thomas Wiegand, and Wojciech Samek. Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Transactions on Image Processing*, 27(1):206–219, 2018. [2](#)
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 139–156, Cham, 2018. Springer International Publishing. [1](#)
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, ICML’20. JMLR.org, 2020. [3, 4](#)
- [5] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. [1](#)
- [6] Yossi Gandelsman, Yu Sun, Xinlei Chen, and Alexei A. Efros. Test-time training with masked autoencoders. *arXiv preprint arXiv:2209.07522*, 2022. [7](#)
- [7] S Alireza Golestaneh, Saba Dadsetan, and Kris M Kitani. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3209–3218, 2022. [2, 5](#)
- [8] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019. [1](#)
- [9] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *Trans. Img. Proc.*, 29:4041–4056, jan 2020. [6](#)
- [10] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 2427–2440. Curran Associates, Inc., 2021. [1](#)
- [11] Gu Jinjin, Cai Haoming, Chen Haoyu, Ye Xiaoxing, Jimmy S. Ren, and Dong Chao. Pipal: A large-scale image quality assessment dataset for perceptual image restoration. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 633–651, Cham, 2020. Springer International Publishing. [6](#)
- [12] 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 (ICCV)*, pages 5148–5157, October 2021. [2, 5](#)
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [14] Nikos Komodakis and Spyros Gidaris. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations (ICLR)*, Vancouver, Canada, Apr. 2018. [1](#)
- [15] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [1](#)
- [16] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6028–6039. PMLR, 13–18 Jul 2020. [2](#)
- [17] Hanhe Lin, Vlad Hosu, and Dietmar Saupe. Kadid-10k: A large-scale artificially distorted iqa database. In *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–3, 2019. [5](#)
- [18] Xialei Liu, Joost van de Weijer, and Andrew D. Bagdanov. Rankiqa: Learning from rankings for no-reference image quality assessment. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [2](#)
- [19] Yutao Liu, Ke Gu, Yongbing Zhang, Xiu Li, Guangtao Zhai, Debin Zhao, and Wen Gao. Unsupervised blind image quality evaluation via statistical measurements of structure, naturalness, and perception. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(4):929–943, 2020. [2](#)
- [20] Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 21808–21820. Curran Associates, Inc., 2021. [1, 2](#)
- [21] Kede Ma, Wentao Liu, Kai Zhang, Zhengfang Duanmu, Zhou Wang, and Wangmeng Zuo. End-to-end blind image quality assessment using deep neural networks. *IEEE Transactions on Image Processing*, 27(3):1202–1213, 2018. [2](#)
- [22] 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. [2](#)
- [23] Anush Krishna Moorthy and Alan Conrad Bovik. Blind image quality assessment: From natural scene statistics to perceptual quality. *IEEE Transactions on Image Processing*, 20(12):3350–3364, 2011. [2](#)

- [24] Chaithanya Kumar Mummadi, Robin Huttmacher, Kilian Rambach, Evgeny Levinkov, Thomas Brox, and Jan Hendrik Metzen. Test-time adaptation to distribution shift by confidence maximization and input transformation, 2021. [1](#)
- [25] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. [2](#)
- [26] Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, and C.-C. Jay Kuo. Image database tid2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication*, 30:57–77, 2015. [5](#)
- [27] Michele A. Saad, Alan C. Bovik, and Christophe Charrier. Blind image quality assessment: A natural scene statistics approach in the dct domain. *IEEE Transactions on Image Processing*, 21(8):3339–3352, 2012. [2](#)
- [28] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11539–11551. Curran Associates, Inc., 2020. [2](#), [3](#)
- [29] H.R. Sheikh, M.F. Sabir, and A.C. Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing*, 15(11):3440–3451, 2006. [6](#)
- [30] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#), [5](#)
- [31] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9229–9248. PMLR, 13–18 Jul 2020. [1](#), [2](#), [6](#)
- [32] Toni Virtanen, Mikko Nuutinen, Mikko Vaahteranoksa, Pirkko Oittinen, and Jukka Häkkinen. Cid2013: A database for evaluating no-reference image quality assessment algorithms. *IEEE Transactions on Image Processing*, 24(1):390–402, 2015. [6](#)
- [33] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#), [6](#)
- [34] Jingtao Xu, Peng Ye, Qiaohong Li, Haiqing Du, Yong Liu, and David Doermann. Blind image quality assessment based on high order statistics aggregation. *IEEE Transactions on Image Processing*, 25(9):4444–4457, 2016. [2](#)
- [35] Peng Ye, Jayant Kumar, Le Kang, and David Doermann. Unsupervised feature learning framework for no-reference image quality assessment. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1098–1105. IEEE, 2012. [2](#)
- [36] Peng Ye, Jayant Kumar, Le Kang, and David Doermann. Unsupervised feature learning framework for no-reference image quality assessment. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1098–1105, 2012. [2](#)
- [37] 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 (CVPR)*, June 2020. [5](#)
- [38] Guangtao Zhai, Xiongkuo Min, and Ning Liu. Free-energy principle inspired visual quality assessment: An overview. *Digital Signal Processing*, 91:11–20, 2019. [2](#)
- [39] Weixia Zhang, Kede Ma, Jia Yan, Dexiang Deng, and Zhou Wang. Blind image quality assessment using a deep bilinear convolutional neural network. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(1):36–47, 2020. [2](#)
- [40] Hancheng Zhu, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. MetaIqa: Deep meta-learning for no-reference image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#), [5](#)