Content-Diverse Comparisons improve IQA

William Thong * ^{1,2} william.thong@sony.com Jose Costa Pereira³ jose.c.pereira@huawei.com Sarah Parisot³ sarah.parisot@huawei.com Ales Leonardis³ ales.leonardis@huawei.com Steven McDonagh³ steven.mcdonagh@huawei.com ¹ University of Amsterdam The Netherlands
 ² Sony Al Switzerland
 ³ Huawei Noah's Ark Lab United Kingdom

Abstract

Image quality assessment (IQA) forms a natural and often straightforward undertaking for humans, yet effective automation of the task remains highly challenging. Recent metrics from the deep learning community commonly compare image pairs during training to improve upon traditional metrics such as PSNR or SSIM. However, current comparisons ignore the fact that image content affects quality assessment as comparisons only occur between images of similar content. This restricts the diversity and number of image pairs that the model is exposed to during training. In this paper, we strive to enrich these comparisons with content diversity. Firstly, we relax comparison constraints, and compare pairs of images with differing content. This increases the variety of available comparisons. Secondly, we introduce listwise comparisons to provide a holistic view to the model. By including differentiable regularizers, derived from correlation coefficients, models can better adjust predicted scores relative to one another. Evaluation on multiple benchmarks, covering a wide range of distortions and image content, shows the effectiveness of our learning scheme for training image quality assessment models.

1 Introduction

This paper addresses image quality assessment, where the task is to assign quantitative scores to rank images by their perceptual quality. An ability to perform this task reliably and accurately is of critical importance to assess real-world applications such as image compression [21, 40, 45], restoration [29, 34] or synthesis [31, 52]. For example, visual artifacts can appear when applying an image compression algorithm, creating distortions that affect quality in comparison with an original pristine version. Towards robustly measuring image quality, the gold standard involves providing human observers the task of independently

^{*} Work done in part during an internship at Huawei Noah's Ark while a PhD student at University of Amsterdam. © 2022. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.



Figure 1: **Image Quality Assessment** by human judgement (columns) is affected by the image content despite a similar distortion severity and type (rows). Enabling models to incorporate this observation leads to richer training signals.

grading images on the axis of perceptual quality. A mean opinion score, per image, then provides relative rankings of images according to their respective average ratings and therefore perceived quality [39, 46].

Recent methods typically approach the ranking problem by learning to compare a pair of distorted images, originating from similar content [11, 28, 41, 62]. However, this ignores the fact that image content often plays an important role when determining human perceived image quality [47, 48, 49, 57]. As illustrated in Fig. 1, for a similar distortion severity and type, the perceived image quality varies with the image content. In other words, a given distortion affects perceptual quality differently, depending on the image content. For example, there can exist two images with differing content and differing distortions but with a similar mean opinion score; or alternatively two images with a similar distortion and differing content but with differing mean opinion scores. As such, we argue that models could gain richer training signals by going beyond comparing images with similar content, and start comparing images with differing content.

Relying on *pairwise* comparisons with similar image content simplifies the image quality ranking task [28, 41]. Indeed, models need only to compare two different distortions. Nevertheless, this formulation on image content can become a constraint as it doesn't reflect the observation in Fig. 1 and restricts the number of valid image pairs available to sample during training. Thus, we propose to generalize pairwise comparisons by considering all possible pairs, regardless of their image content, during model training. This relaxation provides more content diversity in the pairwise comparisons, which in return better leverages the available ground truth image quality annotations.

As comparisons are key to model training, we further explore *listwise* comparisons. While commonly employed in information retrieval settings [24, 26], such objectives remain challenging to optimize due to non-differentiable components. Recently, Brown *et al.* [8] importantly showcase their effectiveness for image retrieval. Indeed, when a task aims for a global ranking (*e.g.*, retrieval or quality assessment), it becomes logical to encourage such effect during training. Inspired by their success, we propose to enhance image quality assessment training signals through the introduction of listwise comparisons. We derive differentiable regularizers that encapsulate correlation scores between predicted outputs and ground truth mean opinion scores at the mini-batch level. Complementary to relaxing image content variance constraints, listwise comparisons provide the model with a holistic view of image ranking, which allows model training to adjust prediction scores in a relative manner.

Our main contributions introduce the concept of content-diverse comparisons when learning image quality assessment models, and can be summarized as:

1. We relax pairwise constraints to enable the comparison of image pairs with differing content. Pairs can be sampled on-the-fly and within one training dataset. In turn, this

better captures wider factors under consideration during human quality assessment.

- 2. We propose listwise comparisons at a mini-batch level, and derive three differentiable regularizers to encourage correlation with ground-truth image quality rankings. Such differentiable derivations are applicable to any model architecture without structural changes. Our regularization terms provide a holistic view during training, which enables the model to adjust image quality scores relative to one another.
- 3. We evaluate our proposed learning method on eight different image quality assessment datasets and benchmarks, covering a wide range of unseen distortions and image content. We also show the applicability of our learning method on multiple network architectures and its effectiveness to reach state-of-the-art performance.

2 Related Work

Image quality assessment assigns a scalar score, used to rank images in terms of their perceived quality. We consider the full-reference setting [54], where we have access to the reference image to score the quality of a distorted image. This is particularly useful for evaluating image restoration methods [29, 31, 34, 62]. Other settings can be employed, such as no-reference settings [54] for image aesthetics [36].

Traditional metrics for image quality assessment emulate human visual system responses, in order to visualize impairments and distortions [57], and rely on the assumption that human perceptual quality can be largely explained by the visibility of impairments. Handcrafted metrics attempt to capture well understood visual psychology phenomena such as spatial frequency dependent contrast sensitivity (*e.g.*, SSIM [56] and variants). PSNR and SSIM are still commonly used, yet often fall short compared to deep learning models able to capture more complex properties of human visual perception. Seminal works propose bespoke convolutional architectures trained from scratch [6, 18, 41], whereas Liu *et al.* [28] show that fine-tuning networks, pre-trained on ImageNet, are sufficient by following an underlying hypothesis that features extracted from image classification are meaningfully transferable. Zhang *et al.* [62] go one step further and highlight how to obtain quality scores in the feature space of these pre-trained networks without any fine-tuning. More recently, Cheon *et al.* [10] adopt a transformer architecture to predict perceptual quality. While effective, these works focus solely on characterizing distortions in the image.

In comparison, the broader field of quality of experience considers external factors, *e.g.*, user characteristics or interests [35, 64], as well as internal factors, *e.g.*, image content [47, 48, 49]. This is notably pertinent for videos where every frame comes with its own content and plays a role in the quality score [23, 27]. In a related direction, Ding *et al.* [11] incorporate texture similarity to the loss for image quality. We are inspired by this literature and propose to enrich supervision signals. Rather than proposing a new model architecture, we focus on the learning scheme to make comparisons more diverse and more holistic.

Pairwise comparisons transform the ranking problem into pairwise classification or regression [24, 26]. This turns the ranking problem into an easier variant, as the model only needs to compare two images at a time. Liu *et al.* [28] initially propose a margin-based ranking loss to compare a pair of distorted images with similar content. Prashnani *et al.* [41] learn to regress to the probability of preferring one image over another while Zhang *et al.* [62] simply treat the problem as a binary classification. Yet, pairwise comparisons are done between two distorted images coming from the same reference image, *i.e.*, similar image content, which does not reflect the true intent of image quality assessment. Closer to our work, Zhang *et al.*

al. [63] rely on pre-sampled image pairs, across multiple datasets. We relax current formulations, and enable our model to learn diverse pairwise comparisons of distorted images with differing image content, which are sampled on-the-fly within a single training dataset.

Listwise comparisons maintain a group structure of ranking, which is ignored in pairwise comparisons [24, 26]. Indeed, while pairwise comparisons focus on two images at a time, listwise comparisons have a more holistic view of the scores in a mini-batch. While Cao *et al.* [9] study the relevance of learning listwise over pairwise rankings, they don't tie the formulation to any correlation coefficients. Towards this goal, several works in quality assessment derive a differentiable regularizer for the Pearson coefficient [22, 23, 27, 32]. Some attempts have also been made to make Spearman coefficient differentiable through linear programming [5] or additional layers for sorting [12]. While effective, they require fundamental changes and cannot be directly be applicable to any model architectures. We then draw inspiration from the information retrieval literature [33], where the ranking operation can be approximated through a logistic function [42]. Brown *et al.* [8] importantly show the positive effect of optimizing such approximated objective for image retrieval. We build on this literature to derive differentiable listwise regularizers, which ensure holistic properties in a mini-batch and can be applied to any model architecture without any changes.

3 Method

Problem formulation. During training, we are given a full-reference image quality assessment dataset $\{x^i, x^i_{ref}, y^i\}_{i=1}^M$ of M distorted images x for which a mean opinion score y quantifies the quality with respect to its reference image x_{ref} . Note that it is common to have multiple distorted images corresponding to the same reference image in image quality assessment datasets. The objective is to learn a function f that produces a scalar value $\hat{y} = f(x, x_{ref})$, in order to predict the mean opinion score and therefore assess the quality of image x. In principle, f can be any statistical function and contemporary approaches typically employ a convolutional network. An effective learning scheme for the scoring function f relies on pairwise training [9, 62]. Given two images x^i and x^j with $i \neq j$, the idea is to learn to predict which image has the better quality. If the mean opinion score y^i is higher than y^j , $\hat{y}^i = f(x^i, x^i_{ref})$ should yield a higher output than $\hat{y}^j = f(x^j, x^j_{ref})$. The model then learns during training to produce a scalar output that quantifies which image in the pair has the better quality. Hence, what matters is to produce a faithful ranking c.f. regressing to y. Sec. 3.1 and 3.2 introduce our pair formation strategy and derive our differentiable regularizers, respectively.

3.1 Pairwise comparisons with differing content

Pair formation. A standard pair formation scheme [11, 62] is illustrated in Fig. 2(a). This first strategy relies on fixed pairs with similar image content, which heavily restricts the number of available pairs. Given N images in a mini-batch, N/2 pairs are formed such that every pair compares two distorted images that originate from the *same* reference image. Fig. 2(b) illustrates an alternative approach to increase the number of comparisons, by constructing mini-batches comprising a single image content (same reference image), and comparing *all* image pairs [28]. Indeed, $(N^2 - N)/2$ pairs can be formed from N images in a mini-batch. However, this imposes a strong constraint on the size of the mini-batch during training. Let D be the number of distortions available for a given reference image. Then, under this regime, the size of the mini-batch N cannot be larger than D, *i.e.*, $N \leq D$. Fig. 2(c) depicts our

tent



(a) *Fixed* pairs with *similar* content



(c) *All* pairs with *differing* content

Figure 2: **Image pair formation** in a mini-batch containing four distortions and respective reference images. (a) Fixed pairs that contrast only similar content limit comparison count. (b) Comparison count can be increased by considering all available distortions for fixed content. (c) We generalize previous assumptions by unconditionally comparing images with differing image content.

(b) All pairs with similar con-

proposal, where we consider *all* possible pairs yet allow image content to *differ* within a pair. This relaxation allows for a broader and more diverse definition of valid pairwise comparisons; a generalization of previous schemes. This content-diverse formulation no longer imposes a constraint on the mini-batch size. More importantly, it enables models to better learn the explored intent of the image quality assessment task without any constraint: images should be rankable outwith hard content restrictions.

Pairwise classification. We learn the scoring function f through pairwise classification. At every iteration, we present to the model a pair of images x^i and x^j with $i \neq j$, along with their respective reference images. If x^i has a better quality than x^j , this entails that the mean opinion score is higher $y^i > y^j$. Thus, f should learn to predict a similar ranking $\hat{y}^i > \hat{y}^j$. To learn this pairwise classification, we derive a probabilistic model of $y^i > y^j$ through the Bradley-Terry sigmoid [7]:

$$p(y^{i} > y^{j}) = 1/\left(1 + \exp\left(-(\hat{y}^{i} - \hat{y}^{j})/T\right)\right),\tag{1}$$

where *T* is a parametric temperature. The Bradley-Terry model applies a sigmoid function to the difference of predicted scores $\hat{y}^i - \hat{y}^j$, which results in a probabilistic model of $y^i > y^j$. With this formulation, we can treat the pairwise ranking task as a binary classification. We then rely on the binary cross-entropy for the loss function:

$$\mathcal{L}_{c} = \mathbb{1}[y^{i} > y^{j}] \cdot \log(p(y^{i} > y^{j})) + \mathbb{1}[y^{i} < y^{j}] \cdot \log(1 - p(y^{i} > y^{j})),$$
(2)

where $\mathbb{1}[\cdot]$ is the indicator function. Prashnani *et al.* [41] initially introduced the Bradley-Terry sigmoid for pairwise comparisons, however their loss formulation learns to regress to the ground truth probability $y^i > y^j$, by minimizing the mean squared error. Thus, the training dataset must include the pairwise probability of preferring x^i over x^j , which requires a different labeling as common datasets usually only assign a mean opinion score. Our binary classification formulation removes this constraint and enables the model to accommodate arbitrary measurement scales, for the scoring ground truth image quality.

3.2 Holistic listwise comparisons

Listwise comparisons. Pairwise comparisons only act as a proxy task for image quality assessment. Indeed, during training, the model has a limited view since it only ever compares

THONG ET AL.: CONTENT-DIVERSE COMPARISONS IMPROVE IQA



Figure 3: **Image ranking.** Three images A, B and C with decreasing image quality. (a) Under pairwise comparisons if the model makes one mistake, *e.g.* rating C as higher quality than B, scores cannot be adjusted relative to A. (b) Listwise comparisons enable a holistic view at the mini-batch level. For a similar mistake; richer feedback enables image scores to be adjusted in a relative manner.

two distorted images. Consider Fig. 3 with a mini-batch of three images A, B and C. If the model makes a mistake ranking B and C, it will be unable to dependably adjust the scores relative to A, as the two pairwise comparisons involving A are correct. As a result, the loss might over-adjust scores, which may damage the correct comparisons. We need a constraint to maintain the correct scores and ensure incorrect scores are not over-adjusted.

We propose to compute listwise comparisons through correlation coefficients. Let $\hat{Y} = {\hat{y}^1, \dots, \hat{y}^L}$ be the predicted scores output by f and $Y = {y^1, \dots, y^L}$ be the related ground truth mean opinion scores for a set of images of size L. Correlation coefficients measure the correlation between \hat{Y} and Y. Such measurements provide a holistic view during training as they encourage the model to predict scores relative to each other at the mini-batch level. We therefore derive regularizers for correlation coefficients at the mini-batch level to better mimic the true objective of image quality assessment.

Correlation coefficients. Listwise comparisons introduce helpful holistic properties on the predicted scores. For example, we can ensure a linearity of the scores to avoid producing outlier values; or preserve the ranking to avoid any over-adjustments. To achieve this, we rely on the Pearson coefficient [38] for linearity, and both Spearman [51] and Kendall [17] coefficients for ranking. We introduce regularizers to measure correlation coefficients at the mini-batch level during training: $R = ||1 - \operatorname{corr}(Y, \hat{Y})||_p$, where $|| \cdot ||_p$ is the *p*-norm and both *Y* and \hat{Y} are of size *N* (*i.e.*, mini-batch size). By minimizing regularizers in the loss function, the model learns to maximize the correlation between *Y* and \hat{Y} . While the Pearson coefficient is differentiable, the other two are not. We describe how to derive a regularizer for the Pearson coefficients:

Pearson linear correlation coefficient measures the linear relationship between *Y* and \hat{Y} , by computing the ratio between their covariance and the product of their standard deviations: $r(Y, \hat{Y}) = \operatorname{cov}(Y, \hat{Y})/\sigma_Y \sigma_{\hat{Y}}$. As the Pearson coefficient does not possess any non-differentiable components, and following previous works [22, 23, 27, 32], we can directly derive the regularizer: $R_r = ||1 - r(Y, \hat{Y})||_p$.

Spearman rank correlation coefficient measures the linear relationship between the ranking of *Y* and \hat{Y} , by computing the Pearson coefficient between their rank values: $\rho(Y, \hat{Y}) = \rho(\operatorname{rank}_Y, \operatorname{rank}_{\hat{Y}}) = \operatorname{cov}(\operatorname{rank}_Y, \operatorname{rank}_{\hat{Y}}) / \sigma_{\operatorname{rank}_Y} \sigma_{\operatorname{rank}_{\hat{Y}}}$, where $\operatorname{rank}_Y = 1 + \sum_{i \neq j} \mathbb{1}[(Y_i - Y_j) < 0]$, and similarly for $\operatorname{rank}_{\hat{Y}}$. The indicator function in the ranking function makes the Spearman coefficient non-differentiable. Following works in image retrieval [8, 42], we approximate the indicator function with a temperature-based sigmoid:

$$\mathbb{1}[(Y^{i} - Y^{j}) < 0] \approx 1/(1 + \exp(-(Y^{i} - Y^{j})/T)),$$
(3)

which creates a differentiable Spearman rank $\tilde{\rho}(\cdot, \cdot)$. We can then derive a corresponding regularizer: $R_{\rho} = \|1 - \tilde{\rho}(Y, \hat{Y})\|_{p}$.

Doin formation	Content	I	LIVE[45	5]	0	CSIQ[2]]	TID2013[40]		
Pair Iormation		PLCC	SRCC	KRCC	PLCC	SRCC	KRCC	PLCC	SRCC	KRCC
fixed pairs	similar	0.958	0.964	0.828	0.951	0.953	0.807	0.889	0.878	0.690
all pairs	similar	0.961	0.966	0.836	0.948	0.952	0.805	0.902	0.892	0.709
all pairs	differing	0.963	0.968	0.842	0.950	0.954	0.809	0.908	0.897	0.717

Table 1: **Pair formation** for pairwise comparisons. We compare a *fixed* pair formation with *all* available pairs in a mini-batch. Pairs with no constraints on image content improve upon more constrained pair formations.

Kendall rank correlation coefficient measures the ordinal association between *Y* and \hat{Y} , by computing the number of concordant pairwise comparisons: $\tau(Y, \hat{Y}) = \frac{2}{n(n-1)} \sum_{i < j} \operatorname{sgn}(Y^i - Y^j) \operatorname{sgn}(\hat{Y}^i - \hat{Y}^j)$, where sgn is the sign function responsible for making the Kendall coefficient non-differentiable. Similar to the indicator function, we approximate sgn with a logistic curve. We propose a temperature-based hyperbolic tangent:

$$\operatorname{sgn}(Y^{i} - Y^{j}) \approx \operatorname{tanh}((Y^{i} - Y^{j})/T), \tag{4}$$

which creates a differentiable Kendall rank $\tilde{\tau}(\cdot, \cdot)$. We can derive a corresponding regularizer: $R_{\tau} = ||1 - \tilde{\tau}(Y, \hat{Y})||_p$. The final loss function then becomes:

$$\mathcal{L} = \mathcal{L}_c + \lambda (R_r + R_\rho + R_\tau), \tag{5}$$

where λ controls the contribution of the linearly combined regularizers.

4 Experiments

4.1 Datasets and Evaluation

Setup. Unless specified otherwise, we rely on DISTS [11] as the scoring function f, which builds on VGG16 [50], pre-trained on ImageNet [43] with L2 pooling [15] to extract features at several levels without aliasing. We train f by minimizing the loss function in Eq. 5 with the Adam [19] optimizer and cosine annealing [30], and set the hyper-parameters as follows: learning rate of 1×10^{-4} , batch size of 64, temperature T of 0.01. For regularization, we use the absolute value (*i.e.*, p = 1) and set λ to 1.0. During training, we randomly crop 256×256 image patches and apply random rotations ($\{0^{\circ}, 90^{\circ}, 180^{\circ}\}$) with horizontal flipping. During test, the model is evaluated using full images in all experiments.

Datasets. Following previous works [10, 11], we train our model on the KADID-10k [25] dataset, which contains 25 different distortions at five different levels of severity for a total of 10,125 distorted images. Distortion types include traditional artifacts, such as those from blurring or compression, but also in-the-wild artifacts, such as denoising or non-linear brightness artifacts. To evaluate generalization capabilities, we compare on three traditional image quality assessment datasets with broad distortion types: LIVE [45], CSIQ [21] and TID2013 [40]. They contain 779, 866 and 3000 distorted images, respectively. We additionally evaluate on five other benchmarks in the appendix.

Evaluation. We report the Pearson linear (PLCC) [38], Spearman rank (SRCC) [51] and Kendall rank (KRCC) [17] correlation coefficients. When computing the PLCC, we follow Ding *et al.* [11] and fit a four-parameter logistic regression (4LP) function: $\hat{y}_{4lp} = (\eta_1 - \eta_2)/(1 + \exp(-(\hat{y} - \eta_3)/\eta_4)) + \eta_2$, where $\{\eta_i\}_{i=1}^4$ are estimated through least squares.

Regularizer		LIVE[45]			CSIQ[21]			TID2013[40]					
	$R_r F$	Rρ	Rτ	PLCC	SRCC	KRCC	PLCC	SRCC	KRCC	PLCC	SRCC	KRCC	Table 2: Regularization of
				0.963	0.968	0.842	0.950	0.954	0.809	0.908	0.897	0.717	r o and τ Multiple regulari
	1			0.962	0.967	0.839	0.952	0.956	0.812	0.906	0.896	0.715	r, p una v. manupie regulari
		1		0.960	0.966	0.835	0.953	0.957	0.815	0.910	0.901	0.723	ers encourage linear properties ar
			1	0.962	0.968	0.840	0.950	0.955	0.811	0.910	0.900	0.722	rank preservation which consi
	1.	~		0.960	0.966	0.837	0.954	0.959	0.819	0.908	0.899	0.718	runk preservation, which const
		1	1	0.961	0.967	0.838	0.941	0.960	0.821	0.909	0.900	0.721	tently improve performance, se
	1		1	0.960	0.966	0.837	0.954	0.959	0.820	0.912	0.903	0.725	text for further detail
	1.	/	1	0.964	0.969	0.843	0.957	0.960	0.824	0.915	0.907	0.731	text for further detail.

4.2 Ablations

Pair formation. We report results of the investigated strategies in Tab. 1. Liu *et al.* [28] previously note that having an order of magnitude more comparative pairs, per image, yields improvement in the correlation metrics. We corroborate this experimentally and observe that, controlling for the number of images provided, the model learns to correctly rank more image pairs on average using *all pairs* compared with *fixed pairs*. By further removing the image content restriction and allowing comparisons between *differing* image content, our proposal leads to a further improvement. In contrast with Liu *et al.* [28], our proposal removes the limitation regarding mini-batch size. For example, if a training dataset contains a variable number of distortions per reference image, our method can gracefully handle such data without additional problem (*c.f. all pairs* with *same* content). Our approach enables models to learn from comparisons exhibiting highly diverse content, which better resembles the underdetermined problem of real-world image quality assessment.

Regularizers. Previous work has focused mainly on pairwise image comparisons (*e.g.* [11, 62]), and also explored a Pearson regularizer (*e.g.*, [22, 23, 27, 32]). Yet, none have assessed the interaction of multiple regularizers with complementary ranking properties. Consequently, we propose to provide models with a richer understanding of image quality, by encouraging the learning of a global ranking. Our regularization signals enable learning of listwise rankings at the mini-batch level, towards a holistic understanding of image quality.

Tab. 2 evaluates the effect of the regularizers at an individual level, as well as in combination. Regularizers can be considered to possess complementary components which we evaluate. The Pearson regularization R_r ensures linearity while the Spearman regularization R_{o} preserves the ranking of the predicted scores in the mini-batch. The Kendall regularization R_{τ} shares the same spirit with the classification loss function \mathcal{L}_c (Eq. 2), as it ensures the pairwise ranking of all available pairs in the mini-batch, and not only the magnitude of the predicted score. Overall, these different regularizers on top of the diverse pair formation provide complementary properties for the model to learn a linear global score for ranking images. In isolation, and also when combined with one other regularizer, we observe that R_r has a negligible effect. This suggests that \mathcal{L}_c already produces linear scores. Both R_{ρ} and R_{τ} share a broadly common goal and focus on the ranking, which is complementary to \mathcal{L}_c . Individually, we observe these two terms tend to provide similar benefit while combining both together may result in some redundancy. Interestingly, it is when all regularizers are included that the performance is best. In this setting, R_r does yield a positive effect, across all three considered benchmarks. Deriving multiple regularizers with complementary roles enhances the type of comparisons learned during training.

A limitation of the listwise regularizers concerns the batch size. When going from 64 to 8, we observe a drop in performance by up to 5.06% on TID2013. We provide more details in the supplementary materials (Sec. A.1.2). Thus, a lower batch size reduces the effectiveness of listwise comparisons.

		LIVE	5[45]	I	CS	SIO[2]	1	TID2013[40]			
Model	PLC	C SRC	CC KR	CC F	LCC	RCC	KRCC	PLCC	SRCC	KRCC	
LPIPS	0.959 0.9		66 0.8	332 0	.952 ().958	0.815	0.895	0.887	0.700	
+L2 pool	0.95	6.0	64 0.8	328 0	.945	0.964	0.827	0.892	0.886	0.699	
DISTS	0.96	64 0.9	69 0.8	343 0	.957 (0.960	0.824	0.915	0.907	0.731	
Mathad		Ι	LIVE[4	5]		CSIQ[21]	TID2013[40]			
Method		PLCC	SRCC	KRCO	PLCC	SRC	C KRCO	C PLCC	SRCC	KRCC	
PSNR		0.865	0.873	0.680	0.819	0.81	0 0.601	0.677	0.687	0.496	
SSIM [56]		0.937	0.948	0.796	0.852	0.86	5 0.680	0.777	0.727	0.545	
MS-SSIM	[55]	0.940	0.951	0.805	0.889	0.90	6 0.730	0.830	0.786	0.605	
VSI [61]		0.948	0.952	0.806	0.928	0.94	2 0.786	0.900	0.897	0.718	
MAD [21]		0.968	0.967	0.842	0.950	0.94	0.797	0.827	0.781	0.604	
VIF [44]		0.960	0.964	0.828	0.913	0.91	1 0.743	0.771	0.677	0.518	
FSIMc [60)]	0.961	0.965	0.836	0.919	0.93	1 0.769	0.877	0.851	0.667	
NLPD [20]	0.932	0.937	0.778	0.923	0.93	2 0.769	0.839	0.800	0.625	
GMSD [58	8]	0.957	0.960	0.827	0.945	0.95	0 0.804	0.855	0.804	0.634	
WaDIQaM	1 [<mark>6</mark>]	0.940	0.947	0.791	0.901	0.90	9 0.732	0.834	0.831	0.631	
PieAPP [4	1]	0.908	0.919	0.750	0.877	0.89	2 0.715	0.859	0.876	0.683	
LPIPS [62]	0.934	0.932	0.765	0.896	0.87	6 0.689	0.749	0.670	0.497	
DISTS [1]]	0.954	0.954	0.811	0.928	0.92	9 0.767	0.855	0.830	0.639	
IQT [10]		-	0.970	0.849	-	0.94	3 0.799	-	0.899	0.717	
Ours		0.964	0.969	0.843	0.957	0.96	0 0.824	0.915	0.907	0.731	

Table 3: **Scoring functions** for perceptual similarity. Our learning strategy can be applied to different architectures without structural modifications, see text for further detail.

Table 4: **Comparison on traditional distortions** with top-2 results highlighted in bold. Only IQT and our method achieve a consistent performance across all three datasets. Our method based on the DISTS architecture improves upon the original training, confirming the benefits our proposed training scheme.

Scoring functions. Our contributions are not tied to a specific quality scoring function f(*i.e.*, model architecture). We evidence this claim by evaluating alternative full-reference image quality scoring functions, trained with our proposed learning strategy: (i) LPIPS [62] takes the difference between the distorted image and its reference at several levels in VGG16 [50] to produce a scalar predicted score; (ii) LPIPS with the addition of an antialiasing module. As image quality assessment requires to capture small differences in the image space and such terms are reported to be beneficial [15]; (iii) DISTS [11] differs by learning to weight the contribution of every feature level. Scoring functions, each employing our training contributions (Sec. 3), are evaluated in Tab. 3. We observe that our training scheme benefits all the selected models, as we improve upon their original scores on the benchmarks. This gain originates from comparison learning instead of training directly to regress mean opinion scores. The latter lacks a mechanism to induce score linearity and ordering, which our regularizers address. Additionally, we evaluate the effect of the patch size in the supplementary materials (Sec. A.1.1) as LPIPS and DISTS have been trained with different patch sizes. We select DISTS for our following experiments (Sec. 4.3) as it yields the best overall performance. We however highlight that our proposed training strategy is applicable to multiple state-of-the-art image quality scoring functions.

4.3 State-of-the-art comparisons

Quantitative evaluation. Tab. 4 compares approaches on LIVE [45], CSIQ [21] and TID 2013 [40]. Ideally, a generic image quality assessment method should yield high scores, regardless of the type of distortions and image content present in the data. In practice, we observe that methods tend to struggle to achieve good overall performance. For example, MAD correlates well with mean opinion scores on the LIVE and CSIQ datasets, yet struggles on TID2013 where there is a large gap with VSI. Disparity in dataset collection protocols (*e.g.*, controlled *vs.* uncontrolled environments) may introduce additional dataset annotation variance and challenges. Interestingly, similar behaviour can be observed on several deep learning based methods (*e.g.*, LPIPS and DISTS), which under-perform on TID2013. Only our method and the recent IQT [10] are able to perform consistently across all three datasets. IQT proposes a transformer, different from convolutional networks most commonly used in



Figure 4: **Qualitative application** of our image quality metric as a training objective for $\times 4$ super-resolution. We evaluate different objectives for ESRGAN [53] (b-d). All improve over the blurriness of bicubic interpolation (a). Ours (d) reduces subtle over-sharpening artifacts present in with Johnson *et al.* [16] (c), best viewed in color with digital zoom.

contemporary image quality assessment, which led the authors to win the NTIRE'21 challenge [14]. Still, our method achieves the best overall scores, which evidences generalization ability across a variety of datasets. Furthermore, we observe that our method, instantiated here using the DISTS architecture, outperforms their original training scheme. In principle, we expect similar gains when applying our training scheme to alternative transformer-based architectures (*e.g.*, IQT), which we consider a promising direction for future work. In the supplementary materials (Sec. A.2), we provide six additional benchmarks where our perceptual loss also yields state-of-the-art results.

Qualitative application. While image quality assessment primarily targets the scoring of distortions arising from image restoration methods, it may alternatively be used to train the respective restoration models. We consider the super-resolution task, and select ESR-GAN [53] which includes a PSNR-oriented *L*2 loss, that optimizes image reconstruction, and a further perceptual loss function component. Following Wang *et al.* [53], we train the models on the DIV2K [3] dataset. Fig. 4 provides qualitative results for a ×4 SR task on a sample from Set14 [59]. While optimizing for the *L*2 loss reduces the blurriness of a simple bicubic interpolation, it still misses high-frequency textures (Fig. 4(a) *vs.* Fig. 4(b)). Adding a perceptual component to the loss function helps. In lieu of the Johnson *et al.* [16] perceptual loss, in the original ESRGAN [53] formulation, we substitute our model as the perceptual component. We observe that over-sharpening artifacts are reduced when our model is used (Fig. 4(c) *vs.* Fig. 4(d)). In the supplementary materials (Sec. A.3), we provide additional examples where our perceptual loss also benefits for reducing the presence of other artifacts. Our brief exploration illustrates the potential of our approach as an objective for downstream tasks where perceptual quality is of prime importance.

5 Conclusion

We present a learning scheme with content-diverse comparisons for learning an image quality assessment model. Firstly, we formulate comparisons in a manner that relaxes image content constraints. This makes pairwise comparisons content-diverse and better leverages the available ground truth annotations. Second, we include listwise comparisons to provide a holistic view to the training signal. We derive differentiable regularizers from correlation coefficients to ensure the linearity and rankings of predicted scores at the mini-batch level, which enables models to adjust scores in a relative manner. Finally, we validate the effectiveness of our learning scheme, where we report state-of-the-art performance and explore the potential of the trained image quality model for downstream imaging tasks.

References

- [1] MindSpore. https://www.mindspore.cn/. Accessed: 2022-10-07.
- [2] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. TensorFlow: a system for Large-Scale machine learning. In *12th USENIX symposium on operating* systems design and implementation (OSDI 16), pages 265–283, 2016.
- [3] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image superresolution: Dataset and study. In *CVPRw*, 2017.
- [4] Sangnie Bhardwaj, Ian Fischer, Johannes Ballé, and Troy Chinen. An unsupervised information-theoretic perceptual quality metric. arXiv preprint arXiv:2006.06752, 2020.
- [5] Mathieu Blondel, Olivier Teboul, Quentin Berthet, and Josip Djolonga. Fast differentiable sorting and ranking. In *ICML*, 2020.
- [6] Sebastian Bosse, Dominique Maniry, Klaus-Robert Müller, Thomas Wiegand, and Wojciech Samek. Deep neural networks for no-reference and full-reference image quality assessment. *TIP*, 27(1):206–219, 2017.
- [7] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [8] Andrew Brown, Weidi Xie, Vicky Kalogeiton, and Andrew Zisserman. Smooth-ap: Smoothing the path towards large-scale image retrieval. In *ECCV*, 2020.
- [9] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to listwise approach. In *ICML*, 2007.
- [10] Manri Cheon, Sung-Jun Yoon, Byungyeon Kang, and Junwoo Lee. Perceptual image quality assessment with transformers. In CVPRw, 2021.
- [11] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *PAMI*, 2021.
- [12] Martin Engilberge, Louis Chevallier, Patrick Pérez, and Matthieu Cord. Sodeep: a sorting deep net to learn ranking loss surrogates. In *CVPR*, 2019.
- [13] Jinjin Gu, Haoming Cai, Haoyu Chen, Xiaoxing Ye, Jimmy S Ren, and Chao Dong. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. In ECCV, 2020.
- [14] Jinjin Gu, Haoming Cai, Chao Dong, Jimmy S Ren, Yu Qiao, Shuhang Gu, and Radu Timofte. NTIRE 2021 challenge on perceptual image quality assessment. In *CVPRw*, 2021.
- [15] Olivier J Hénaff and Eero P Simoncelli. Geodesics of learned representations. In *ICLR*, 2016.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In ECCV, 2016.

- [17] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [18] Jongyoo Kim and Sanghoon Lee. Deep learning of human visual sensitivity in image quality assessment framework. In *CVPR*, 2017.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [20] Valero Laparra, Johannes Ballé, Alexander Berardino, and Eero P Simoncelli. Perceptual image quality assessment using a normalized laplacian pyramid. *Electronic Imaging*, 2016(16):1–6, 2016.
- [21] Eric Cooper Larson and Damon Michael Chandler. Most apparent distortion: fullreference image quality assessment and the role of strategy. *Journal of electronic imaging*, 19(1):011006, 2010.
- [22] Dingquan Li, Tingting Jiang, and Ming Jiang. Norm-in-norm loss with faster convergence and better performance for image quality assessment. In ACM Multimedia, 2020.
- [23] Dingquan Li, Tingting Jiang, and Ming Jiang. Unified quality assessment of in-the-wild videos with mixed datasets training. *IJCV*, 129(4), 2021.
- [24] Hang Li. Learning to rank for information retrieval and natural language processing. *Synthesis lectures on human language technologies*, 7(3):1–121, 2014.
- [25] Hanhe Lin, Vlad Hosu, and Dietmar Saupe. Kadid-10k: A large-scale artificially distorted iqa database. In *QoMEX*, 2019.
- [26] Tie-Yan Liu. Learning to rank for information retrieval. Springer, 2011.
- [27] Wentao Liu, Zhengfang Duanmu, and Zhou Wang. End-to-end blind quality assessment of compressed videos using deep neural networks. In ACM Multimedia, 2018.
- [28] Xialei Liu, Joost van De Weijer, and Andrew D Bagdanov. Rankiqa: Learning from rankings for no-reference image quality assessment. In *ICCV*, 2017.
- [29] Yiming Liu, Jue Wang, Sunghyun Cho, Adam Finkelstein, and Szymon Rusinkiewicz. A no-reference metric for evaluating the quality of motion deblurring. *TOG*, 32(6):175–1, 2013.
- [30] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.
- [31] Chao Ma, Chih-Yuan Yang, Xiaokang Yang, and Ming-Hsuan Yang. Learning a noreference quality metric for single-image super-resolution. *CVIU*, 158:1–16, 2017.
- [32] Kede Ma, Zhengfang Duanmu, and Zhou Wang. Geometric transformation invariant image quality assessment using convolutional neural networks. In *ICASSP*, 2018.
- [33] Christopher D Manning and Prabhakar Raghavan. Introduction to Information Retrieval. Cambridge University Press, 2008.

- [34] Xiongkuo Min, Guangtao Zhai, Ke Gu, Yucheng Zhu, Jiantao Zhou, Guodong Guo, Xiaokang Yang, Xinping Guan, and Wenjun Zhang. Quality evaluation of image dehazing methods using synthetic hazy images. *TMM*, 21(9):2319–2333, 2019.
- [35] Sebastian Möller and Alexander Raake. *Quality of experience: advanced concepts, applications and methods.* Springer, 2014.
- [36] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In CVPR, 2012.
- [37] 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.
- [38] Karl Pearson. Vii. note on regression and inheritance in the case of two parents. *proceedings of the royal society of London*, 58(347-352):240–242, 1895.
- [39] Nikolay Ponomarenko, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Lina Jin, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, et al. Color image database tid2013: Peculiarities and preliminary results. In *EUVIP*, 2013.
- [40] Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, and Kuo C.-C. Jay. Image database tid2013: Peculiarities, results and perspectives. *Signal processing: Image communication*, 30:57–77, 2015.
- [41] Ekta Prashnani, Hong Cai, Yasamin Mostofi, and Pradeep Sen. Pieapp: Perceptual image-error assessment through pairwise preference. In *CVPR*, 2018.
- [42] Tao Qin, Tie-Yan Liu, and Hang Li. A general approximation framework for direct optimization of information retrieval measures. *Information retrieval*, 13(4):375–397, 2010.
- [43] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [44] Hamid R Sheikh and Alan C Bovik. Image information and visual quality. *TIP*, 15(2): 430–444, 2006.
- [45] Hamid R Sheikh, Zhou Wang, Lawrence Cormack, and Alan C Bovik. Live image quality assessment database. http://live.ece.utexas.edu/research/quality, 2003.
- [46] Hamid R Sheikh, Muhammad F Sabir, and Alan C Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *TIP*, 15(11):3440–3451, 2006.
- [47] Ernestasia Siahaan, Alan Hanjalic, and Judith A Redi. Augmenting blind image quality assessment using image semantics. In ISM, 2016.

- [48] Ernestasia Siahaan, Alan Hanjalic, and Judith A Redi. Does visual quality depend on semantics? a study on the relationship between impairment annoyance and image semantics at early attentive stages. *Electronic Imaging*, 2016(16):1–9, 2016.
- [49] Ernestasia Siahaan, Alan Hanjalic, and Judith A Redi. Semantic-aware blind image quality assessment. *Signal Processing: Image Communication*, 60:237–252, 2018.
- [50] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for largescale image recognition. In *ICLR*, 2015.
- [51] C Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.
- [52] Shishun Tian, Lu Zhang, Luce Morin, and Olivier Déforges. A benchmark of dibr synthesized view quality assessment metrics on a new database for immersive media applications. *TMM*, 21(5):1235–1247, 2018.
- [53] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *ECCVw*, 2018.
- [54] Zhou Wang and Alan C Bovik. *Modern image quality assessment*, volume 2. Morgan & Claypool Publishers, 2006.
- [55] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *IEEE Asilomar Conference on Signals, Systems & Computers*, 2003.
- [56] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, 2004.
- [57] Andrew B Watson and Cynthia H Null. Digital images and human vision. In *Electronic Imaging Science and Technology Conference*, 1997.
- [58] Wufeng Xue, Lei Zhang, Xuanqin Mou, and Alan C Bovik. Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *TIP*, 23(2):684–695, 2013.
- [59] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711– 730. Springer, 2010.
- [60] Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. Fsim: A feature similarity index for image quality assessment. *TIP*, 20(8):2378–2386, 2011.
- [61] Lin Zhang, Ying Shen, and Hongyu Li. VSI: A visual saliency-induced index for perceptual image quality assessment. *TIP*, 23(10):4270–4281, 2014.
- [62] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [63] Weixia Zhang, Kede Ma, Guangtao Zhai, and Xiaokang Yang. Uncertainty-aware blind image quality assessment in the laboratory and wild. *TIP*, 30, 2021.
- [64] Yi Zhu, Alan Hanjalic, and Judith A Redi. Qoe prediction for enriched assessment of individual video viewing experience. In *ACM Multimedia*, 2016.