

Consolidated Dataset and Metrics for High-Dynamic-Range Image Quality

Aliaksei Mikhailiuk, María Pérez-Ortiz, Dingcheng Yue, Wilson Suen, and Rafał K. Mantiuk

Abstract—Increasing popularity of high-dynamic-range (HDR) image and video content brings the need for metrics that could predict the severity of image impairments as seen on displays of different brightness levels and dynamic range. Such metrics should be trained and validated on a sufficiently large subjective image quality dataset to ensure robust performance. As the existing HDR quality datasets are limited in size, we created a Unified Photometric Image Quality dataset (UPIQ) with over 4,000 images by realigning and merging existing HDR and standard-dynamic-range (SDR) datasets. The realigned quality scores share the same unified quality scale across all datasets. Such realignment was achieved by collecting additional cross-dataset quality comparisons and re-scaling data with a psychometric scaling method. Images in the proposed dataset are represented in absolute photometric and colorimetric units, corresponding to light emitted from a display. We use the new dataset to retrain existing HDR metrics and show that the dataset is sufficiently large for training deep architectures. We show the utility of the dataset on brightness aware image compression.

Index Terms—High Dynamic Range, Image Quality Dataset, Image Quality Metric

I. INTRODUCTION

IMAGE quality assessment metrics, such as peak signal to noise ratio (PSNR) and structural similarity index measure (SSIM) are widely used in image compression, reconstruction, and enhancement [1], [2], [3], [4], [5]. However, most image quality assessment (IQA) metrics do not account for display characteristics such as the dynamic range and brightness of the display, influencing the perceived image quality. For example, compression artifacts are more visible on a bright HDR display, than on a dimmed mobile phone [6]. The plethora of display types motivates the need for a new, *photometric* IQA metric that accounts for absolute image luminance and can operate on both HDR and SDR images. Throughout this work we use the term *metric* to refer to image quality metric rather than to a distance function in a metric space.

The primary limitation to developing an HDR image quality metric has been the lack of a unified large-scale subjective image quality dataset. Although attempts have been made to adapt and verify performance of SDR metrics on HDR content [7], [8], [9], [10], [11] those have not been thoroughly tested due to the lack of a unified dataset. The absence of a large unified dataset also prevented the development of metrics based on machine learning for HDR images, which

require large amounts of versatile and heterogeneous data to train. While recent machine-learning-based SDR image quality metrics relied on large crowd-sourcing studies [12], [13], [14], these are not straight forward to conduct for HDR content as it requires an HDR display and controlled viewing conditions.

The available subjective image quality datasets [15], [16], [11], [17], [18], [19], [12], [13], [20], are insufficient in isolation, as they are limited in terms of the number of images, diversity of distortion types and image sizes. These datasets cannot be easily combined, due to the use of different experimental protocols and the relative nature of the quality scales, which precludes comparing quality scores across datasets. Moreover, incomparable quality scales across datasets prevent the use of absolute scores as a mean of benchmarking IQA metrics, forcing to rely on correlation coefficients computed individually on each dataset. This work addresses these issues. Instead of following the common practice of collecting a dataset from scratch, we argue for consolidation of existing datasets and focus on combining SDR and HDR image quality datasets to create the largest photometric subjective IQA dataset to date with a unified quality scale. We use the dataset to re-train existing full-reference metrics, including deep architectures.

Our contributions can be summarized as follows: (i) we perform a series of subjective image quality assessment experiments and construct the largest subjective HDR IQA dataset to date (UPIQ) using psychometric scaling [21]. The dataset contains 3779 SDR and 380 HDR images from four existing IQA datasets; (ii) we show the necessity and advantages of the psychometric scaling by comparing it to other strategies for merging datasets; (iii) we use the new dataset to retrain and benchmark existing HDR metrics. We show that the proposed dataset is sufficiently large for deep architectures by training a convolutional neural network (CNN)-based full-reference *photometric* image quality metric. The advantage of training on the unified dataset is shown in comparison with training on a single dataset and performing multi-task learning on disjoint datasets; (iv) the utility of training HDR metrics on the new dataset is shown in an application to image compression. The new dataset¹, code and metrics are available online².

II. RELATED WORK

In this section we set the context for the problem and provide a brief review on existing IQA datasets as well as the HDR and SDR objective quality metrics.

A. Mikhailiuk, D. Yue, W. Suen and R. Mantiuk are with the Department of Computer Science and Technology at the University of Cambridge (UK) (email: {am2442, dy276, wss28, rkm38}@cam.ac.uk).

M. Pérez-Ortiz is with the Department of Computer Science at the University College London (UK) (email: maria.perez@ucl.ac.uk)

¹UPIQ dataset: <https://doi.org/10.17863/CAM.68372>

²Project page: <https://www.cl.cam.ac.uk/research/rainbow/projects/upiq/>

A. Existing IQA Datasets

To train and validate image quality metrics, one requires a dataset where image quality scores are obtained from human observer judgements. Two most common experimental protocols for collecting such a dataset are ranking and rating experiments. In rating experiments, observers are asked to assign a numeric quality score to each image. All judgments are then averaged to produce Mean Opinion Scores (MOS). In ranking, observers are asked to compare two or more images and order them according to their quality. Ranking results can be then mapped to a one-dimensional quality scale using psychometric scaling [22], [23], [24], [25]. The most commonly used ranking protocol is pairwise comparison, where a pair of images is compared at a time. Advantages of pairwise comparisons over rating are the lower cognitive load on observers and a more accurate scale [26]. Unlike rating, scores produced by psychometric scaling are interpretable – distance between any two scores can be mapped to the probability of one condition being selected over another [25].

Although many subjective IQA datasets exist, they are far from ideal. For example, the largest currently available SDR dataset, BAPPS [13], offers only a single distortion type per content (where we define content as the scene depicted in the image) — therefore machine learning based metrics may struggle to learn how to scale the magnitude of a distortion. Moreover, image quality scores were not measured extensively, with only two judgments per 64×64 pixel patches rather than full-sized images. Another recently collected large-scale SDR dataset [12], contains pairwise preference probability, *i.e.* the likelihood of an image in a pair of being more similar to the reference. Even though authors collected a large number of comparisons per image (> 10), only within-content comparisons were performed, thus making the cross-content quality scores less reliable. Similarly authors in [27] collected a large scale SDR dataset with MOS. These datasets are not publicly available. Existing HDR IQA datasets [11], [18], [16] are significantly smaller than SDR datasets, more homogenous in the versatility of their contents and distortions. These datasets are thus insufficient for the applications outlined in this paper. To remedy these limitations, in this work we combine HDR datasets with much larger SDR datasets.

Since collecting large amount of IQA measurements is time consuming and costly, it is preferable to reuse existing datasets. The idea of combining subjective IQA datasets has been considered before. Authors in [11] align subjective scores of HDR datasets using objective quality metrics. The method assumes that the quality predictions from multiple objective metrics can be used to find the transformation of quality scores from one dataset to another. However, this approach is problematic when combining SDR and HDR datasets as very few metrics can reliably predict absolute quality of both SDR and HDR images. In contrast to that work, we conduct a set of subjective experiments to measure the relative cross-dataset quality and then use psychometric scaling [21] procedure to bring all datasets to a unified quality scale.

Datasets can be collected more efficiently using *active sampling* methods, which choose the measurements that maximize

the information gain [28], [29]. Another approach is to use quality metrics to find the conditions for which the metrics disagree the most and help to differentiate the performance of those metrics (MAD competition) [30], [31]. Those methods, however, are not intended for merging the existing large datasets.

B. Quality Assessment Criteria

There are at least three common criteria related to image quality: aesthetics, visibility and impairment. Aesthetic judgements deal with the quality of an image as judged by commonly established photographic rules — appropriate use of lighting, contrast, and image composition. Here, the quality may be perceived in terms of creative composition and execution of an image, rather than artifacts [32]. As an example, tone-mapping metrics, which assess the reproduction of HDR images on regular monitors, estimate the aesthetics of tones, brightness, details and colors [33], [34]. Visibility metrics predict whether a difference between a pair of images is going to be visible, but they do not assess the magnitude of a distortion [35], [36]. They also produce visual difference maps rather than a single quality score. The focus of this work are impairment metrics, which assess the quality of images distorted by noise, blur, compression, and other artifacts. Here, we only consider full-reference metrics, for which the original undistorted image is available.

C. Existing Approaches to IQA

Metrics vary in the number of trainable parameters and therefore the amount of data required to train them. The simplest metrics, such as PSNR and SSIM [37], are designed to capture image statistics that is deemed to be important for detecting distortions. These metrics do not have any trainable parameters. Other metrics involve modeling relevant characteristic of visual system, such as contrast masking or contrast sensitivity [38], [35]. Because those metrics rely on the existing psychophysical models, they have only a few parameters to train. Another common approach is to extract a number of hand-crafted features and then use a machine-learning model to predict quality based on those [39], [40], [41], [42]. The last group of methods relies on deep-learning methods to both learn the features relevant for quality and the function that would map those features into image quality [43], [44], [13], [12], [45], [46]. Deep learning methods have achieved the state-of-the-art performance on several benchmarks, however, they are susceptible to the quality and quantity of the training data. If data are scarce, which is common in IQA, given the difficulty in collecting the human judgements, the model will fail to generalize.

D. Training Robust Deep Learning models

Transfer learning is often used to alleviate the problem of insufficient and noisy data when training deep-learning models. For example, authors in [44], [47], [48], [49] pre-trained a CNN model on image classification tasks, arguing that learned features would capture image statistics important

for IQA. Others [13], [50] pre-trained the network on the quality predictions of the hand-crafted quality metrics, and then fine-tuned on the subjective image quality scores. Authors in [51] exploited yet another approach — they first pre-trained the network to classify distortions. The risk of this approach, however, is that it can overfit the model to the given set of distortions. Machine learning metrics can also be trained to rank pairs of images rather than rating them [46], [12], [13]. The advantage of such approach is that training can be performed directly on the pairwise comparison data. However, the shortcoming of this approach is that it discards meaningful information by converting the quality scores to a binary classification problem. Although all these approaches can improve the ability of the ML-metrics to generalize, they do not alleviate the need for a larger and diverse IQA dataset, an issue widely acknowledged in most works [52].

E. Influence of a Display on the Perceived Image Quality

SDR metrics typically operate on 8-bit gamma-encoded pixel values and ignore display characteristics, such as its brightness and resolution and viewing conditions such as viewing distance. Such an approach was justifiable in the era of cathode ray tube (CRT) monitors with very similar characteristic. However, current display devices can vary widely. For example, peak display luminance can vary from 5 cd/m² for a dimmed mobile phone to 6 000 cd/m² for a bright HDR display. One way to account for display brightness is to represent both SDR and HDR images in absolute colorimetric units, so that, for example RGB=[100 100 100], corresponds to white color (D65) of 100 cd/m². This changes the paradigm of assessing image quality from device-independent measurements (e.g. PSNR on gamma-encoded pixels), to device-specific measurements, which require knowledge of the target display. While this introduces additional difficulty of selecting display parameters, it is a necessary step for quality assessment on modern displays, and especially those with HDR capabilities. The standardization of reference display parameters may simplify this step in the future.

1) *Display model*: Most modern displays can be modelled using a gain-gamma-offset model [53]. Such a model transforms gamma-encoded color values into absolute linear color values as follows:

$$C_{\text{lin}} = (L_{\text{peak}} - L_{\text{black}}) C_{\text{sRGB}}^{\gamma} + L_{\text{black}}, \quad (1)$$

where $\gamma = 2.2$, C_{lin} is a linear color value, C_{sRGB} is the gamma-encoded color value for one of the channels (R, G or B) and L_{peak} and L_{black} are peak and black luminance levels of the display, respectively.

2) *Photometric image quality metrics*: As the dynamic range of a display affects the visibility of distortions of a viewed image, a reliable quality metric should be able to account for it. We will refer to the metrics that operate on physical photometric/luminance values as *photometric quality metrics*. HDR quality metrics, such as HDR visual difference predictor (HDR-VDP) [35] or HDR video quality measure (HDR-VQM) [38]) are photometric and account for a large range of luminance produced by HDR displays.

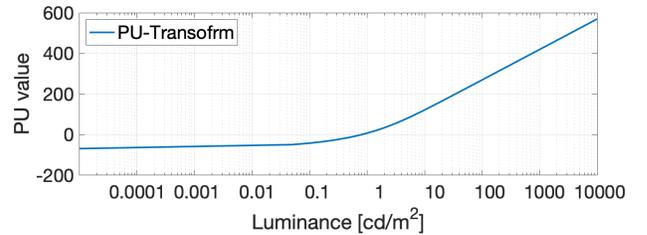


Fig. 1: PU-transform mapping physical luminance in cd/m² into approximately perceptually uniform units.

3) *Extending SRD metrics to HDR images*: SDR metrics can also be adapted to operate on photometric quantities [54]. For that, the absolute luminance values are converted into Perceptually Uniform (PU) or logarithmic values, with the former achieving better results [7], [18]. This transformation is necessary as the response of the human eye to luminance is not linear. Perceptually uniform values are then passed to an SDR image quality metric.

PU-transform [7] was derived from the contrast sensitivity function (CSF) that predicts detection thresholds of the human visual system across a broad range of luminance adaptation conditions. The transform was designed to ensure that the smallest perceivable change in luminance (just-noticeable-difference or detection threshold) is mapped to a constant change in the PU values. This was achieved by a numerical solution of:

$$PU(L) = \int_{L_{\min}}^L \frac{1}{T(l)} dl, \quad (2)$$

where L_{\min} is the minimum luminance to be encoded and $T(l)$ is the detection threshold at the absolute luminance l . The shape of PU-transform is shown in Figure 1. The transformation is typically stored as a look-up table. The transformation is further constrained to map luminance values typically reproduced on SDR displays (0.8-80 cd/m²) to 0 – 255 range (8-bit encoding). This ensures that metric predictions for SDR images in the PU domain are comparable to those computed on the original SDR pixel values.

Similar procedure can be applied to evaluate quality of SDR images when they are shown on displays with different brightness. Before passing through the PU-transform an SDR image is first transformed to luminance emitted from a display, assuming a model of that display, such as the one in Equation 1. The full pipeline of making SDR image quality metrics photometric is given in Figure 2.

Authors in [55] used the PU-transform to adapt a no-reference deep SDR IQA metric to operate on HDR images. The model in [55] had to rely on a metric trained on SDR images due to the absence of a sufficiently large HDR dataset. In this work, we provide such a HDR dataset enabling us to train a deep HDR IQA metric on both SDR and HDR images.

III. PSYCHOMETRIC SCALING

In this section we extend a probabilistic model from our previous work [21] to combine quality measurements from different datasets and from two different experimental protocols:

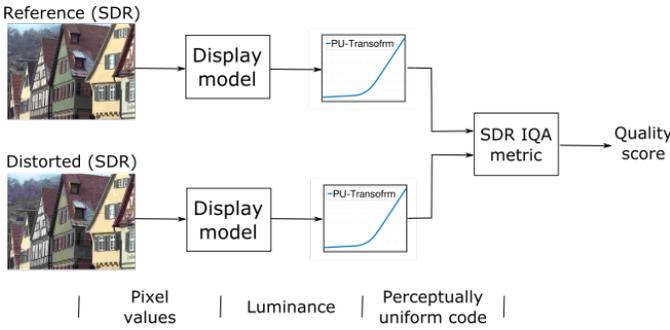


Fig. 2: Extending SDR metrics to photometric values. An SDR image is first converted from gamma-corrected pixel values to linear luminance values, via a display model. For HDR images the “Display model” is omitted, as those already store luminance values. Luminance values are then passed via the PU-transform to obtain perceptually uniform values. The images in PU domain are then passed to an SDR image quality metric.

ranking and rating. In the previous work, we demonstrated how mixing the scores from these protocols could lead to higher precision for a single dataset [21]. In this work, we extend the idea from our previous work to mix the scores from different datasets and realign them to a common unified scale. The scaling method and the experimental protocol are generalisable to mixing other subjective assessment datasets where rating and ranking preference aggregation methodologies are used.

A. Observer Model

To build a quality scale, certain assumptions need to be made about how observers respond and perceive quality. Such assumptions are encapsulated in the observer model. This model is needed because observers vary in their notions of quality (inter-observer variance), and their opinions are also likely to change when they repeat the same experiment (intra-observer variance).

Let \mathbf{q} be an N -dimensional variable whose individual components $q_i \in \mathbb{R}$ represent the underlying true quality of condition i , and N is the number of conditions. We use the term condition to refer to an object or an item to be compared; in our case a condition represents an image content with certain distortion type and at a certain level.

According to the widely used Thurstone’s model Case V [24], perceived quality is normally distributed and the standard deviation σ is known. We can then introduce r_i as the random variable associated with the measured quality, $r_i \sim \mathcal{N}(\hat{q}_i, \sigma)$. Here the mean of the distribution approximates the true quality q_i and the standard deviation is assumed to be the same for all conditions.

B. Scaling Fundamentals

Psychometric scaling aggregates pairwise comparisons into a (generally) uni-dimensional quality scale. The collected data can be represented in a count matrix \mathbf{C} , where element c_{ij} contains the number of times condition i was chosen over

condition j . Then, the probabilities of selecting one condition over another can be directly estimated from this matrix:

$$\hat{p}_{ij} = \frac{c_{ij}}{c_{ij} + c_{ji}}, i \neq j, \quad (3)$$

where \hat{p}_{ij} represents the probability that condition i is perceived as of better quality than j .

The main aim of psychometric scaling model is then to recover the true quality scores from these estimated probabilities. For the assumed model, the probability of choosing i over j should match the cumulative normal distribution Φ over the difference $r_i - r_j$:

$$P(r_i > r_j) = P(r_i - r_j > 0) = \Phi\left(\frac{q_i - q_j}{\sqrt{2}\sigma}\right), \quad (4)$$

where σ dictates the relationship between distances in the scale and probabilities of better quality. A common approach is to select σ so that a probability of 0.75 (in the midway between a random guess and being completely certain) is mapped to a distance of 1 unit in the quality scale. The difference of 2 units then corresponds to the probability of 0.91 and so on.

A common approach for scaling pairwise comparisons is to use maximum likelihood estimation. That is, we estimate the difference in quality scores maximizing the probability of observing collected data \mathbf{C} . The probability that i was selected over j in exactly c_{ij} trials from the total number of $n_{ij} = c_{ij} + c_{ji}$ trials is given by the Binomial distribution:

$$L(q_i, q_j | c_{ij}, c_{ji}) = \binom{c_{ij} + c_{ji}}{c_{ij}} P(r_i > r_j)^{c_{ij}} (1 - P(r_i > r_j))^{c_{ji}}. \quad (5)$$

To scale all compared conditions, we maximise the product of the likelihood for all pairs of conditions:

$$\operatorname{argmax}_{q_2, \dots, q_n} \prod_{i, j \in \Omega} L(q_i - q_j | c_{ij}, n_{ij}), \quad (6)$$

where Ω is the set of all pairs for which at least one comparison has been made. A more detailed discussion of the psychometric scaling can be found in [25].

C. Mixing Pairwise Comparisons and Rating

Our main assumption for mixing the two protocols, rating and ranking, is that different psychophysical experiments aim to recover the same latent ground truth variable \mathbf{q} by different means. For pairwise comparisons, quality scores are recovered by psychometric scaling and for rating, approximated by mean opinion scores (MOS). We further assume that the relationship between the MOS values (denoted as π) and the quality values resulting from psychometric scaling (denoted as r) is monotonic. We validated several polynomial relationships in supplementary. Similar to previous literature [11] we found that for image quality assessment the relationship was well described by a linear model:

$$r_i = a \cdot \pi_i + b, \quad (7)$$

where a and b are the parameters and $\pi_i \sim \mathcal{N}(q_i, c \cdot \sigma)$ with c compensating for varying measurement accuracy/observer models for both of these experimental protocols. Because

rating protocols often use different rating scales (1–5 or 0–100) and because different protocols result in different inter- and intra-observer variation, we optimize for separate set of parameters a, b and c for each dataset.

Given that MOS values are generally measured in a continuous scale and we are assuming the same randomly distributed observer model, the probability of observing m_{ik} (single rating measurement for i^{th} condition and k^{th} observer) can be expressed using the density function of the normal distribution:

$$f(m_{ik}|q_i, a, b, c) = \frac{1}{\sqrt{2\pi a^2 c^2 \sigma^2}} e^{-\frac{((a \cdot m_{ik} + b) - q_i)^2}{2a^2 c^2 \sigma^2}}. \quad (8)$$

Assuming independence between observers, the likelihood of observing ratings from J observers for N conditions aggregated in \mathbf{M} is:

$$P(\mathbf{M}|\mathbf{q}, \sigma, a, b, c) = \prod_{i=1}^N \prod_{k=1}^J f(m_{ik}|q_i, \sigma, a, b, c). \quad (9)$$

To recover latent scores \mathbf{q} from both measurements, the posterior probability can be factorised as:

$$P(\mathbf{q}, a, b, c | \mathbf{C}, \mathbf{M}, \sigma) \propto P(\mathbf{q}) \cdot P(\mathbf{C}|\mathbf{q}, \sigma) \cdot P(\mathbf{M}|\mathbf{q}, a, b, c, \sigma), \quad (10)$$

where $P(\mathbf{q} = \mathcal{N}(\frac{1}{N} \sum_{i=1}^N q_i, \sigma))$ is a Gaussian prior. As we found in our earlier work [25], the prior improves the precision of the scaling when the number of measured comparisons (per pair of conditions) is low. The prior introduces bounds on the score values in the presence of unanimous answers, which put no upper bound on the distance between the recovered scores. The conditional independence of \mathbf{C} and \mathbf{M} given \mathbf{q} is assumed. $P(\mathbf{C}|\mathbf{q}, \sigma)$ is computed as in Eq. 5 assuming independence between measurements. $P(\mathbf{M}|\mathbf{q}, a, b, c, \sigma)$ is computed using the density function of the normal distribution. We infer the quality scores \mathbf{q} , and the parameters a, b, c using maximum likelihood estimation. As scales are relative, we constrain the scores for all reference images (without any distortion) in all datasets to be zero ($q_i = 0$ for each i that represents a reference image).

Since likelihood functions are scale invariant, we can fix σ to any value without loss of generality. In our case we fix $\sigma = 1.048$, so that a distance of 1 unit between two conditions indicates that 75% of observers can see the difference between two conditions, allowing the interpretation of distances in the scale. These units are often referred to as Just-Objectable-Differences (JODs) [21].

IV. UNIFIED PHOTOMETRIC IQA DATASET (UPIQ)

Our goal is to create a large dataset consisting of both SDR and HDR images, with the image quality scores on a unified quality scale with JOD units. This is achieved by selecting existing SDR and HDR datasets, collecting additional cross and within-dataset comparisons, and scaling all the measurements together. We call our dataset UPIQ (“You Pick”) — Unified Photometric Image Quality. Before our work, the largest HDR IQA dataset contained only 240 conditions [18]. Our dataset, has 4159 images, making it the largest and the most diverse HDR image quality dataset to date. Unlike most IQA datasets,

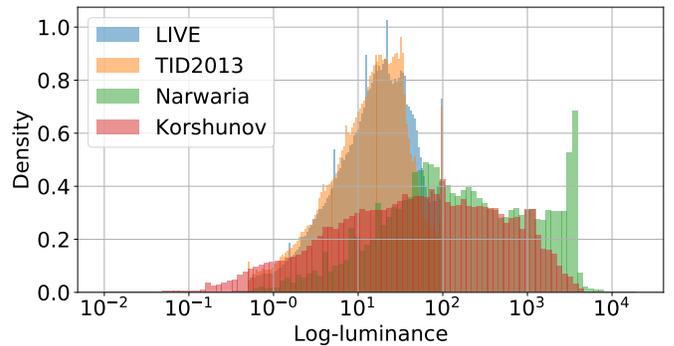


Fig. 3: Distribution of log-luminance per dataset.

images in our dataset are provided in absolute photometric units cd/m^2 and scores are provided in interpretable JOD units.

A. Selected Datasets

We selected four existing datasets—two SDR (TID2013 [15] and LIVE [17]) and two HDR (Korshunov [18] and Narwaria [16]), which we summarize in Table I. All four datasets span very large dynamic range, as shown in Figure 3. Despite a large number of available IQA datasets, only a few of them meet our criteria and could be included in UPIQ. Some datasets were constructed for the purpose of no-reference quality assessment [56], [57], [58] and do not contain reference images [20]. Other datasets contained a single distortion per content, thus they provided no means to scale the magnitude of a distortion [20], [19], [59]. For some datasets, the image size was too small for a proper judgement of image quality [60]. While we attempted to scale some datasets, we found their quality scores to be too inconsistent with our measurements to be included in UPIQ [11]. We provide additional details on the dataset selection in the supplementary.

B. Dataset Alignment Experiments

To align quality scores from different datasets, we need to perform several types of pairwise comparisons, illustrated in Figure 4. Comparisons within a single dataset (within-content, cross-content and with-reference) are needed to bring the quality values to a common scale of JOD units. This is especially important for the datasets with only MOS (rating) values as these are provided in an arbitrary scale. We need to find the relationship between MOS and JOD values by estimating the associated parameters (a, b and c in Eq. (7)). The cross-dataset comparisons are necessary to ensure that the quality values are comparable across the datasets. Because different datasets usually do not share the same content, cross-dataset comparisons tend to be also cross-content comparisons. Cross-content comparisons have been shown to be of the similar difficulty as within-content comparisons [21] and they significantly improve the accuracy of a quality scale [26], [61].

a) Displays and stimuli: The data necessary for alignment were collected on two different displays. Comparison of SDR to SDR images were performed on a color calibrated 32” SDR Samsung S32D850T display with 2560×1440

TABLE I: Characteristics of the chosen IQA datasets

Name	Dynamic range	Experiment	No. images	No. distortions	No. contents	Image sizes (h×w pixels)
LIVE [17]	SDR	MOS	779	5	29	512×768
TID2013 [15]	SDR	PWC	3000	24	25	348×512
Narwaria [16]	HDR	MOS	140	2	10	1080×1920
Korshunov [18]	HDR	MOS	240	3	20	1080×944

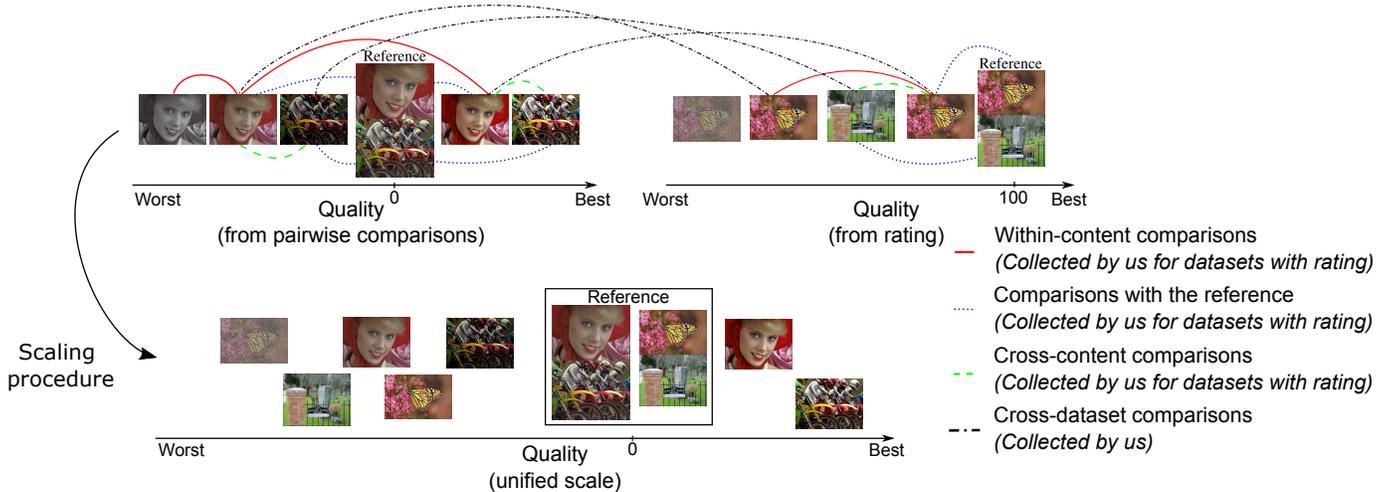


Fig. 4: Types of comparisons necessary for dataset alignment. The lines link pairs of images selected for pairwise comparisons. Within-content comparisons (red solid lines) are most commonly used in pairwise comparison experiments. However, such datasets often lack comparisons with reference (blue dotted lines), which are useful to provide an absolute anchor of quality. Cross-content comparisons (green dashed lines) are less common, but can substantially improved the quality of the scale [26]. Finally, cross dataset comparisons (black dash-dotted lines) are necessary to scale the datasets together.

pixels, 300cd/m^2 typical peak luminance and a black level of $\sim 0.3\text{cd/m}^2$. The comparisons involving HDR images were presented on a custom-built, color-calibrated 10" HDR display with 2048×1536 pixels, $15,000\text{cd/m}^2$ peak luminance and a black level below 0.01cd/m^2 [62].

We used the display model from Equation 1 to transform gamma-corrected sRGB colors to linear RGB values shown on the HDR display. Because we had no information on the displays used in the SDR image quality studies, we used the typical parameters of an SDR display: $\gamma = 2.2$, the peak luminance, $L_{\text{peak}} = 100\text{cd/m}^2$, and the black level, $L_{\text{black}} = 0.5\text{cd/m}^2$. For HDR images, we reproduced the absolute luminance values used in the original studies. The viewing distance was 90cm for both the SDR display (51 pixels per degree) and the HDR display. Both HDR and SDR images, shown on the HDR display, were upscald by a factor of 3.2 (50 pixels per degree), making the measurements taken for the original SDR and HDR studies comparable with ours.

b) Experimental procedure and participants: The observers were presented with pairs of images and were asked to select the image of better quality with respect to the reference. Observers could press and hold the space-bar to view the reference images. When the image size exceeded the size of our display, we provided a simple panning interface in which observers could use a trackball to inspect different portion of the image. Each participant viewed images in different

order. Each selected pair of images was compared by 6 participants, with each participant completing approximately 300 trials. Overall 6000 new comparisons were collected from 20 participants. Note that this required relatively moderate experimental effort as compared to collecting the data from scratch (3000 images in the TID2013 dataset required over 500,000 comparisons). To improve the information gain of the collected data and to exclude obvious comparisons [28], paired images were of similar quality.

We conduct two types of experiments. The first is cross-dataset comparisons. This type of comparisons was only necessary for rating-based datasets, which means we excluded TID2013 from this experiment since we used previously collected pairwise comparisons and rating measurements [61], [21]. We ensured that all three types of necessary comparisons were covered: to reference, within-content and cross-content. After the first experiment, all the data could be scaled, since we had comparisons to a common reference for all datasets.

For the second experiment we compared conditions exclusively from different datasets, connecting each dataset to the rest. Images were chosen to uniformly cover the whole quality scale. We performed several iterations of the pair selection. After conducting an experiment on a small batch of comparisons, we re-scaled the dataset with newly collected comparisons and selected the next batch from the new scale.

C. UPIQ Dataset Scaling

We combined the newly-collected comparisons with the original data from the four datasets and from the two follow-up studies on TID2013 [61] and LIVE [28]. In total, the combined dataset consists of 571,215 individual pairwise comparisons and 27,676 rating measurements, which were passed to the scaling procedure from Section III.

Figure 5 shows the relationships between original quality values of each dataset and the new JOD values from our unified dataset. The plots show substantial differences between the original and rescaled quality scores, suggesting that cross-dataset scaling and additional measurements have further improved the quality estimates. Note that the original scores of the TID2013 dataset were obtained with vote counts, reliant only on within-content comparisons. This approach has proven to be less accurate as compared to psychometric scaling [61].

D. UPIQ Dataset Validation

The qualitative comparison, showing images at a constant JOD level, is shown in supplementary. The figures demonstrate that images at the same JOD level contain comparable perceived magnitude of distortions. In the following subsections we compare our scaling with the metrics-based dataset alignment, and then demonstrate the improvement in pairwise accuracy.

1) *Comparison to previous re-scaling work:* Multiple IQA datasets can be merged together using an iterated nested least-squares (INLS) algorithm [63]. The algorithm uses existing objective quality metrics to find the relationship between conditions in different datasets. The assumption made is that a weighted combination of metrics should have high correlation with human judgments. The algorithm iteratively finds weights for the combination of objective quality metrics and aligns subjective quality scores from each of the datasets until convergence. Since no metric exists that has been exhaustively tested on both SDR and HDR images, we validate the results using two HDR datasets (Korshunov and Narwaria), which were aligned with INLS in the previous work [11]. Figure 6 shows that our scaling procedure and the one from [11] lead to substantially different scores. To determine which alignment is more consistent with the subjective judgements, we compute the rank-order correlation between the unprocessed human subjective measurements and scaled values. Since the collected human judgment data comes in the form of pairwise comparisons, we compute the correlation between empirical probability of selecting one condition over another and differences in quality scores.

The method proposed in [63] relies exclusively on rating and objective metrics to re-align datasets. Whereas our approach, uses psychometric scaling, which requires additional pairwise comparisons to build the unified quality scale. Thus, to ensure a fair comparison, we perform a five-fold cross-validation on the collected cross-dataset comparisons. We split cross-dataset comparisons into five equal-sized partitions. In each fold of the cross-validation we scale the data from four partitions and use the fifth partition for validation. The cross-validation results are given in Table II. For each fold, our model correlates

TABLE II: SROCC between scaled quality scores and empirical probabilities, for our and metric-based scaling. The values are reported for each fold of the cross-validation.

Validation Fold	1	2	3	4	5
Psychometric scaling (our)	0.77	0.72	0.62	0.74	0.71
Objective-metric-based [11]	0.67	0.60	0.52	0.52	0.53

better with the subjective judgments, with mean spearman rank order correlation coefficient (SROCC) of 0.71 versus 0.56 for the method from [11]. It should be noted that the correlation values computed in this manner cannot reach high values because of the measurement noise in the pairwise comparison data. Figure 7 also shows that the relationship is closer to the expected cumulative normal function for our method. This is further confirmed when JOD differences are converted into probabilities (Equation 4) and plotted in the right panel of Figure 7. The scaled probabilities are within the confidence interval of the measured probabilities, confirming that the scaling procedure leads to the quality values that well reflect the empirical quality differences.

2) *Measuring pairwise accuracy:* In this section we provide interpretation of the SROCC results from our validation experiment. We will demonstrate that the scale correctly ranks 97% of the pairs that are at least 1 JOD apart.

We first transform the collected data and the produced scale into pairwise rankings. This is, if the quality of i is higher than that of j (as measured in the collected pairwise comparison matrix C) then we set the binary target label t_{ij} to +1, otherwise we set the target t_{ij} to -1. This represents the ground truth pairwise rank averaged across the population. We then compare this ground truth binary label to our predicted binary labels \hat{t}_{ij} , following the same procedure but using the output of the scaling algorithm instead of probabilities. Having ground truth and predictions, we compute ranking accuracy. For this, we ran a 10-fold cross-validation. In each iteration we withheld 10% of the compared cross-dataset pairs of conditions for validation. The remaining 90% of compared pairs were for scaling. To compute the ranking accuracy, we assume the minimum threshold distance (in terms of JODs) that is required for a pair of conditions to be considered, then report the ratio of the number of correctly ranked considered pairs to the total number of considered pairs.

Figure 8 shows the accuracy scores for different thresholds of reliable JOD differences, for both our scale and that of [63]. For conditions > 0.75 JODs apart (where 63% of observers agreed on the highest quality image, only 13% more than random choice), our scale has 90% accuracy. That is, 90% of the pairs which are more than 0.75 JODs apart are correctly ranked by our psychometric scaling. The difference with [63] is very significant with our scale being consistently much more accurate across different thresholds.

3) *Cross-dataset scaling versus multi-task learning:* The final hypothesis we tested is whether a multi-task deep-learning network can implicitly infer the relation between the datasets and thus can be used to merge datasets without the need for additional data. In this approach, the neural network

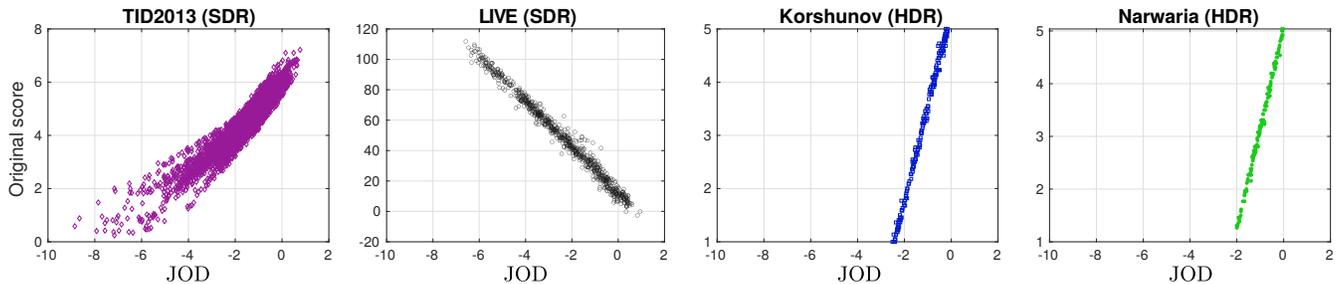


Fig. 5: Original quality scores versus the results of our joint dataset scaling in JODs (UPIQ dataset). Note that the original scores of the TID2013 are the most different from the new JOD scores. This is because the original dataset used a simplified scaling procedure and lacked cross-content comparisons [61].

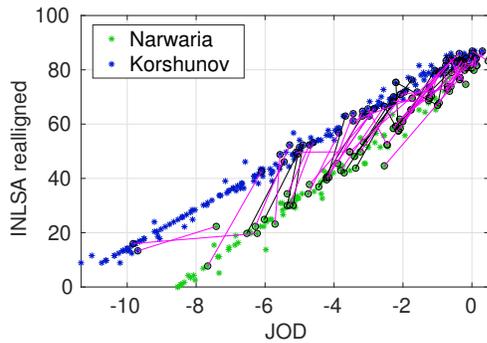


Fig. 6: Scatter plot of the two considered quality scales for the HDR datasets. The plot also shows an example of the data used for one of the cross-validation folds. Purple lines represent the training comparisons and black lines the test comparisons

is trained to predict the scores of the four individual datasets (tasks). The network consisted of two parts. The first part predicts a common score. The second part maps the common score via simple linear regressors, different for each dataset, to the original quality values collected in each dataset. Both parts were trained jointly end-to-end. The details of the architecture are provided in supplementary. This experiment, however, was unsuccessful. The average SROCC between the common score predicted by the multi-task network and the ground truth pairwise comparisons, was only 0.27 for cross-dataset pairs.

V. APPLICATIONS

In this section we show how our UPIQ dataset is useful to train CNN-based metrics and benchmark existing HDR quality metrics. We further show how metrics trained on our dataset can be used for brightness-aware image compression.

A. Training Data-driven HDR Metric

UPIQ is a sufficiently large image quality dataset to enable us to train from scratch a CNN-based image quality metric to predict quality of both SDR and HDR images. The metric combines the ideas behind PU encoding [7] (Section II-E) and a recently proposed CNN architecture for image quality assessment (PieAPP, [12]). We will refer to this metric as PU-PieAPP.

a) *Architecture*: The diagram of the deep metric architecture is shown in Figure 9. The metric takes as input a pair of test and reference images and produces a single quality score s_A in JODs. To account for the dynamic range of the displayed images, the input images need to be transformed into the display domain. This is achieved by a *display model* from Equation 1 for SDR images, or by scaling color values according to the presentation conditions from the original papers for HDR images. Then, the resulting trichromatic color values (with Rec. 709 primaries [64]) are converted into approximately perceptually uniform units using the PU-transform (Section II-E), which is applied individually to each color channel. Such encoded images are fed into the PieAPP architecture, which combines a pair of feature extraction networks with shared weights with the score computation network, both identical to the one used in [12]. The detailed architecture is provided in supplementary.

We train the network on 64×64 patches. To densely cover the whole image, the image is stratified by a uniform grid and patches are sampled at random positions in each grid cell (jittered sampling). The grid size is selected to give approximately square cells. In training, we extract 1024 patches per image. We found that 1024 was the largest number of patches that we could process on our GPU. When testing, we sampled twice the number of 64×64 patches needed to cover the image. This number was optimal in terms of the time vs. performance trade off.

b) *Training*: In contrast to [12], we train the network as regression rather than learning-to-rank. Our scaling procedure achieves the same goals as learning-to-rank, but offers a more accurate observer model and allows us to split the problem into two separate steps of scaling and learning. We train the network from scratch, using Adam optimizer on 4 NVIDIA P100 GPUs. Every tested architecture was run for 500 epochs and the model with the best performance on the validation set was saved (using 60-20-20 split into training, validation and test sets).

B. Benchmark of HDR Quality Metrics

Although HDR image quality metrics have been compared in many studies [18], [11], [9], none of them could test the metrics on an extensive dataset such as UPIQ. Therefore, we use UPIQ to test existing HDR metrics.

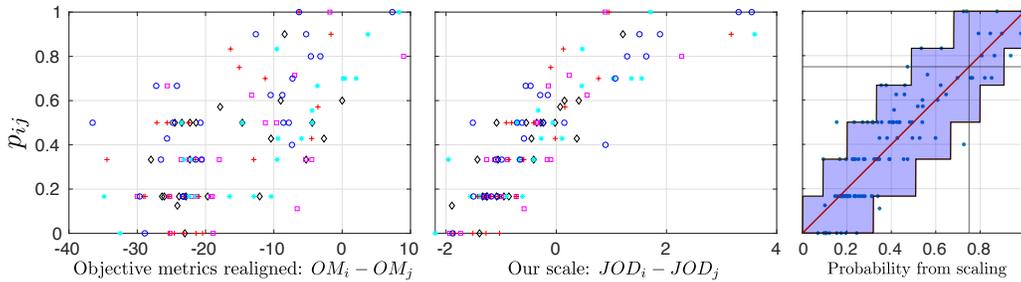


Fig. 7: Right two plots: scatter plots for five folds (distinguished by colors and shape); difference in quality scores in the scale constructed with objective quality metrics [11] $OM_i - OM_j$ versus empirical probability, p_{ij} , of one image i being selected over image j (left); and difference in quality scores in our scale $JOD_i - JOD_j$ versus p_{ij} (center). Our scale is clearly better aligned with the measurements. The right plot: the same as the centre plot but for the entire dataset (no folds) and plotted as the probability instead of JOD difference. The shaded area denotes a 90% confidence interval for the measurements (p_{ij}). The 75% preference (1 JOD) is denoted by the black vertical and horizontal lines. Our scaling procedure brings the quality differences within the confidence interval of the measurements.

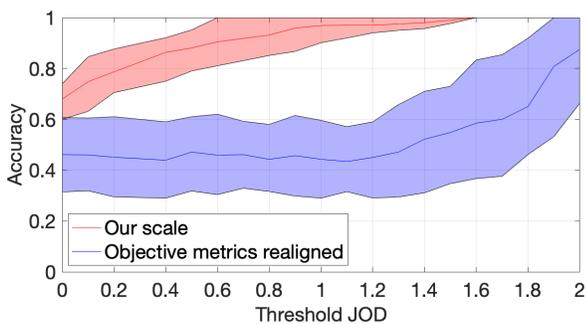


Fig. 8: Accuracy of classifying cross-dataset conditions into better/worse after alignment with the proposed method. Higher value of Threshold JODs means that more conditions are excluded from training and testing sets. Shaded region is 95% confidence interval.

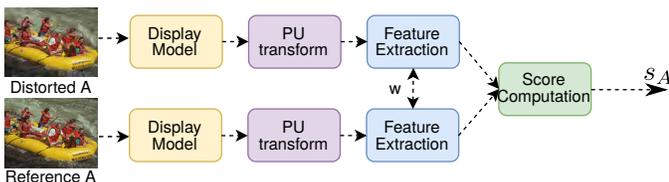


Fig. 9: The pipeline used to train PU-PieAPP on absolute scores. Images first pass through the display model and are then fed to the PU-transform. The feature extraction network with shared weights (w) extract representations, which are passed to the score computation network.

Here we consider full-reference metrics, which are either adapted to HDR content using PU-transform: PU-PSNR, PU-SSIM [37], PU-FSIM [65], or are designed to work with HDR data: HDR-VQM [38], HDRVDP-2.2 [35], [66] and NLP [67]. We also evaluate no-reference metrics, adapting them to the HDR content with PU-transform: PU-BRISQUE [40], PU-PIQE [41] and PU-NIQE [42], due to their widespread use

and competitive performance. Finally, we adapted existing SDR CNN-based metrics to HDR content using the PU-transform: PU-KonCept512 [14] (no-reference) and original PU-PieApp (original) [12] (full-reference). We did not re-train deep metrics on UPIQ, but used weights provided by the authors. For comparison, we also include full reference PSNR and FSIM metrics, not adapted to the HDR content.

Most objective metrics predict values that are non-linearly related to absolute quality in JOD units. The scatter plot of the considered metrics predictions versus those of the JOD's is provided in Figure 10. Since our goal is to predict the absolute quality, we need to map metric predictions to JODs. We follow a standard approach [17] and fit a logistic function mapping objective quality o into absolute JOD units $q(o)$:

$$q(o) = \frac{a_1}{1 + e^{a_2(o-a_3)}} + a_4 o + a_5, \quad (11)$$

where a_1, \dots, a_5 are fitted parameters. Fitting a logistic function is necessary for computing performance measures, RMSE and PLCC, but it also helps to scale objective metric results into interpretable and comparable units of JODs. For example, while the result of PU-SSIM of 0.98 is difficult to interpret, the result of -1 JODs tell us that 75% of the population will notice the difference.

For fair comparison, we use the same 5-fold split into 80-20% training and testing dataset when fitting psychometric function for the tested metrics. In each fold a different portion of the entire dataset is tested while ensuring that no content is shared between training and testing sets. We also make sure that each subset (TID2013, LIVE, Narwaria and Korshunov) was split in the same 80-20 ratio. Note that since PU-PieAPP (re-trained) is trained on the quality scores from UPIQ dataset, we do not need to fit the logistic function into its prediction.

a) Cross-content validation: The most common approach to the validation of learning-based quality metrics is the split into training and testing sets that contain different content but share distortion types. Note that we took extra care to isolate the same content in LIVE and TID datasets as those share some of the reference images. The results for

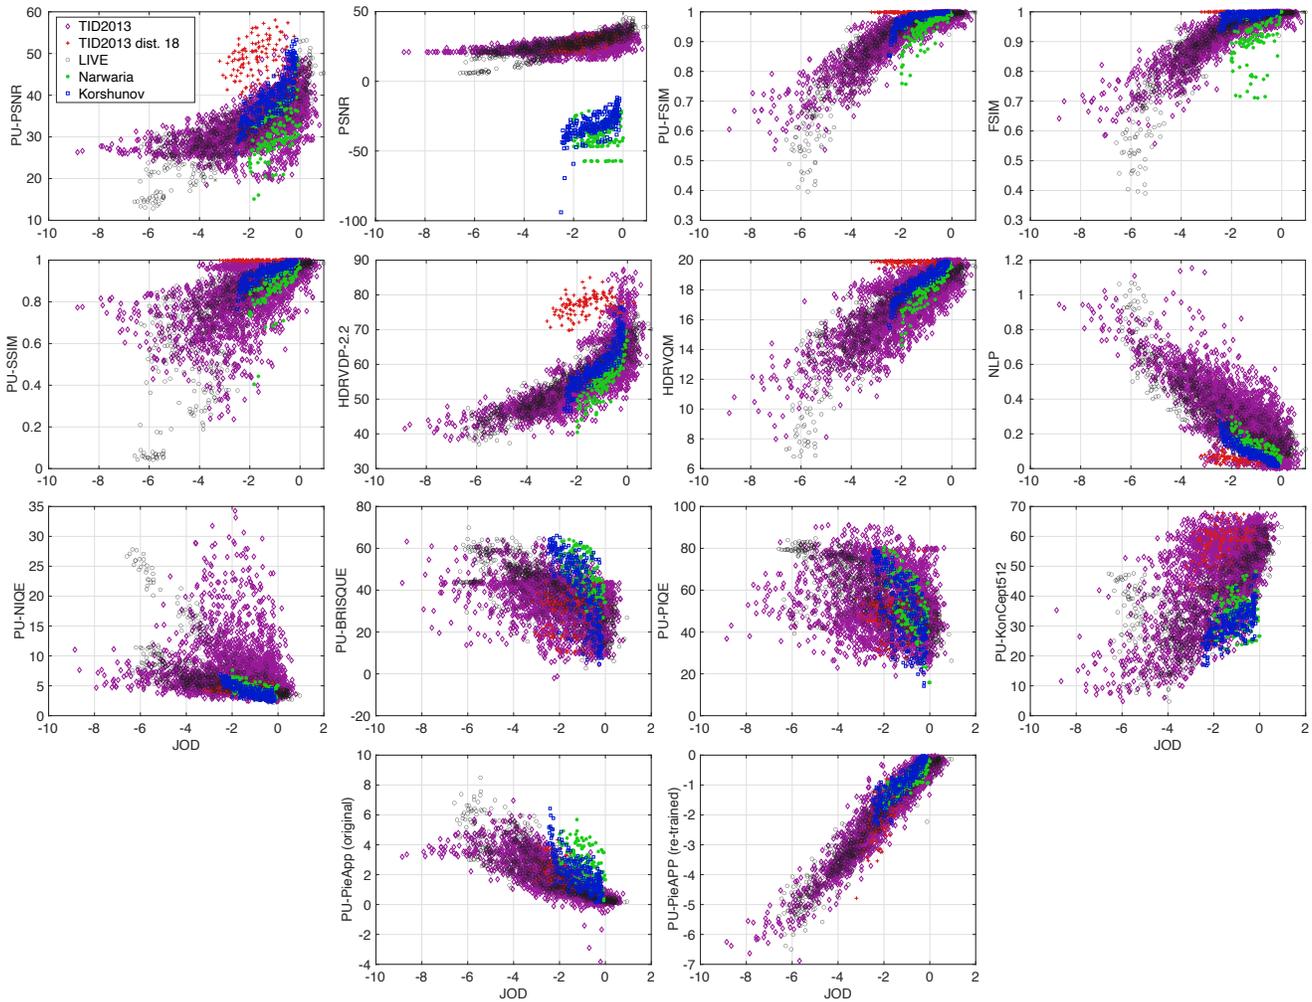


Fig. 10: Objective metrics predictions vs. JOD quality values. We separately label distortion 18 from TID2013 (change of color saturation) as it introduces the biggest prediction error for the metrics that operate on luma/luminance values and ignore color information.

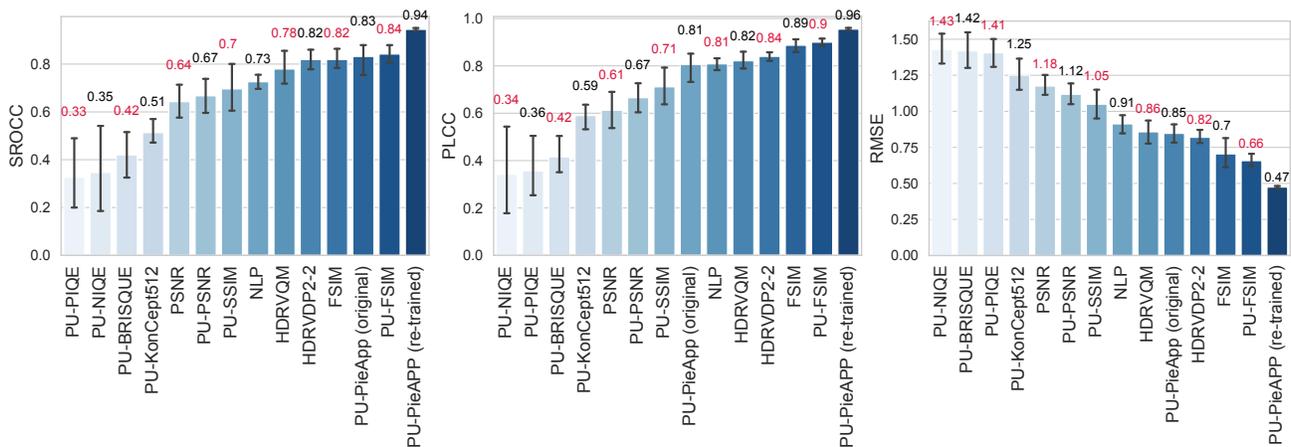


Fig. 11: Cross-validation results for all trained metrics, expressed as SROCC, pearson linear correlation coefficient (PLCC) and root-mean-squared error (RMSE). Error bars denote 95% confidence intervals.

the 5-fold cross-validation on such cross-content splits, shown in Figure 11, indicate that PU-PieAPP (re-trained) outper-

forms existing hand-crafted metrics. PU-PieAPP shows 30% improvement to the second-best performing metric, PU-FSIM,

TABLE III: Test RMSE, SROCC and PLCC for different data partitioning schemes and the best performing metrics. (C-C – cross-content, C-D – cross-dataset, C-DR – cross-dynamic-range). We remove the listed test portion of the UPIQ from training, and test on it. Note that the first column (C-C) shows the same data as Figure 11.

Metric	C-C Test: sel. cont.	C-D Test: TID2013	C-D Test: LIVE	C-D Test: Narwaria	C-D Test: Korshunov	C-DR Test: HDR	C-DR Test: SDR
RMSE							
PU-PieAPP	0.47	0.92	0.70	0.68	0.62	0.72	1.29
PU-FSIM	0.66	0.65	0.50	0.26	0.29	0.68	1.39
FSIM	0.70	0.65	0.51	0.45	0.52	1.17	1.61
HDRVDP	0.82	0.88	0.64	0.24	0.21	0.78	1.34
HDRVQM	0.86	1.04	0.68	0.23	0.20	0.39	1.43
SROCC							
PU-PieAPP	0.94	0.78	0.87	0.82	0.79	0.74	0.65
PU-FSIM	0.90	0.80	0.96	0.87	0.93	0.71	0.77
FSIM	0.89	0.80	0.96	0.54	0.52	0.45	0.54
HDRVDP	0.84	0.78	0.94	0.94	0.94	0.81	0.82
HDRVQM	0.82	0.71	0.92	0.95	0.95	0.87	0.60
PLCC							
PU-PieAPP	0.96	0.78	0.89	0.78	0.75	0.73	0.67
PU-FSIM	0.90	0.89	0.96	0.87	0.90	0.66	0.77
FSIM	0.89	0.89	0.96	0.53	0.66	0.34	0.51
HDRVDP	0.84	0.83	0.93	0.89	0.95	0.72	0.78
HDRVQM	0.82	0.78	0.92	0.89	0.95	0.86	0.62

which is followed by FSIM without the PU-transform. We later show that the performance difference between PU-metrics and original metrics is much higher when tested on HDR datasets (UPIQ is dominated by images from SDR datasets). No-reference metrics, based on hand-crafted features, exhibit the worst performance — the PU-transformation applied to the images distort the statistics that these metrics rely on. Deep learning based no-reference metric PU-Koncept512 does not perform well either.

Original PieApp adapted to our dataset with PU-transform, performs reasonably well on SDR images (SROCC: 0.8764). However, exhibits poor performance on both HDR datasets (SROCC: 0.5791). This is expected, as the metric was trained on SDR images, and the range of PU-transformed HDR images is larger than that of SDR.

b) Cross-validation schemes: To understand what mixture of data is required to robustly train quality metrics, we experiment with different data partitioning schemes. For this experiment, we selected 5 best performing metrics from Figure 11. Table III lists the training and test data combinations we tested and the corresponding results.

PU-PieAPP generalizes well when trained cross-content (C-C), i.e. the training and test set overlap in distortion types but not in content. However, the performance of this deep-learning metric drops significantly if one or more datasets are missing from the training set. This, and the poor performance of no-reference metrics in Figure 11, show that learning-based metrics are prone to overfitting when the training dataset is not sufficiently large.

As expected, SDR metrics exhibit better performance when tested on SDR datasets. The same holds for metrics aimed for HDR content – they perform better on HDR datasets. PU-

FSIM and FSIM have similar performance when tested on SDR datasets. However, when tested on HDR, PU-FSIM performs significantly better compared to FSIM, clearly demonstrating the need for the PU-transform.

C. Maximum differentiation competition

We use the gMAD [31] procedure to find the pairs of images that differ the most according to one metric, but are similar according to another metric. Figure 13 shows a set of failure cases for the two best performing metrics (PU-PieAPP (re-trained) and PU-FSIM) when paired against each other in the gMAD competition. PU-PieAPP tends to correctly capture the ranking of the images, however underestimates the quality of images with JPEG artifacts. PU-FSIM fails to account for color change of the image. PU-FSIM is also too sensitive to contrast change and not sensitive enough to the structural distortions.

D. Brightness-Adaptive Coding

As HDR metrics account for absolute luminance levels, they are not only useful for testing the quality of HDR images, but also open opportunities for new applications such as brightness-adaptive coding. Such coding adapts the required bandwidth to the brightness of a display; saving bits on a dimmed displays and using higher bandwidth when higher quality is needed for a bright display. We investigate how PU-PieAPP and PU-FSIM (the two best-performing metrics) predict the quality of images shown on displays of different brightness. Figure 12a shows the quality predictions for JPEG distortion as a function of peak display luminance. As expected, both metrics predict that JPEG distortions are less noticeable when the display is darker. However, the predictions

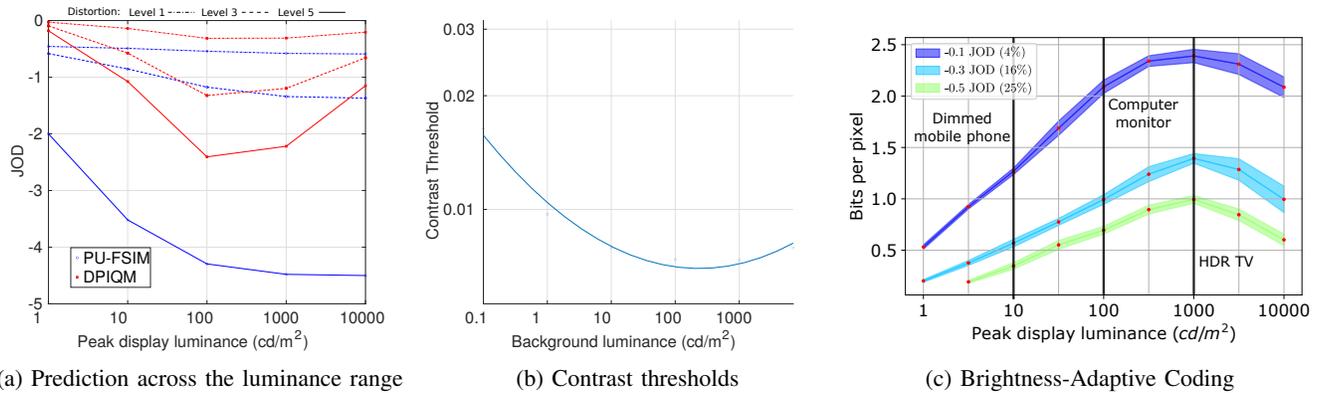


Fig. 12: (a) The quality predictions as the function of display peak luminance. The predictions are shown separately for three distortion levels of JPEG and averaged across contents from TID2013 dataset. (b) Contrast threshold function for a varied display brightness; (c) Bits per pixel for JPEG compression to achieve constant perceived quality at different luminance levels. Different colors represent different quality levels. Shaded regions are 75% confidence intervals.

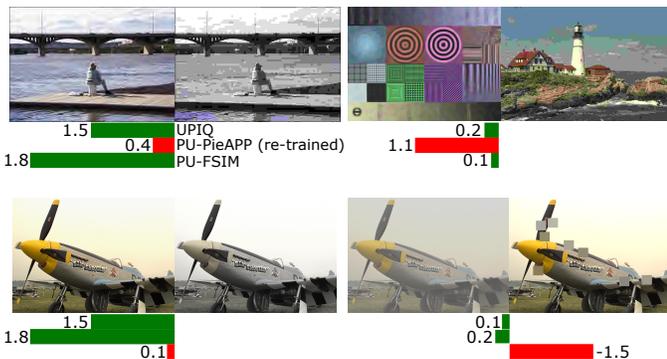


Fig. 13: Selection of image pairs from the gMAD competition where PU-PieAPP (re-trained) and PU-FSIM disagree the most. In the top row we show successful attack and successful defence of the PU-FSIM and in the bottom row the same for PU-PieAPP (re-trained). Below each pair we report the difference in the JOD scores between the left and the right image found in the UPIQ dataset and assigned by each of the tested metrics.

diverge at luminance levels above 100 cd/m^2 : PU-FSIM predicts a decrease in image quality, whereas PU-PieAPP predicts an improvement. Interestingly, PU-PieAPP’s U-shaped curve is consistent with the recent measurements [62] of human contrast detection thresholds. We show an example of these measurements in Figure 12a.

Next, we use PU-PieApp to control the compression rate of a standard JPEG codec (the “quality” parameter) in order to achieve a distortion at a desirable JOD level. Figure 12c shows the distribution of the required bit-rate to compress 200 pristine test images at the desired JOD level. The selected levels signify that about 4% (-0.1 JOD), 16% (-0.3 JOD) or 25% (-0.5 JOD) will correctly indicate a compressed image from a test and reference pair (discounting 50% guess rate). The vertical bars in the plot denote the peak luminance levels of three typical displays: an HDR TV, computer monitor, and a dimmed mobile phone. The plot shows that the bit-rate

could be substantially reduced when images are shown on a dimmed mobile phone, but it should be increased for HDR TV. Furthermore, the difference between brightness levels is larger for images encoded with high quality. Such information could be useful, for example, for internet caches that attempt to reduce the amount of data sent to mobile web browsers. It is only meaningful to use photometric metrics, trained on both SDR and HDR images, for such applications as they can capture the effect of absolute luminance on image quality. A more detailed validation of such brightness-adaptive image coding can be found in [68].

VI. CONCLUSIONS

A large scale photometric image quality dataset would enable the development of deep learning based image quality metric. However, existing HDR image quality datasets are small in size and expensive to collect. We remedy this limitation and increase their size by merging together a mixture of both HDR and SDR datasets. Our merging procedure requires collecting additional data (cross-dataset comparisons), however, the experimental effort is much smaller compared to collecting the dataset from scratch. The accuracy of the resulting dataset is much higher than that of alternative procedures [63], [11]. The proposed dataset merging procedure can be applied to other quality domains, such as the quality of high-frame-rate, omni-directional or foveated video.

Another major contribution of this work is Unified Photometric Image Quality dataset (UPIQ), which is the first large-scale dataset that can be used for training and testing HDR image quality metrics. Images in our dataset are represented in absolute photometric and colorimetric units and their quality scores are represented in the interpretable JOD units [21]. We use the dataset to (a) train a deep-learning based quality metric for HDR images (PU-PieApp); (b) benchmark the state-of-the-art HDR image quality metrics; and (c) demonstrate how trained quality metrics can be used for brightness-adaptive image coding. All those applications demonstrate the benefits of UPIQ and the need for photometric image quality metrics.

ACKNOWLEDGMENT

This project has received funding from EPSRC research grants EP/P007902/1 and EP/R013616/1, from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement N° 725253 (EyeCode)), and from the Marie Skłodowska-Curie grant agreement N° 765911 (RealVision).

REFERENCES

- [1] H. Talebi and P. Milanfar, "Nima: Neural image assessment," *IEEE Transactions on Image Processing*, vol. 27, no. 8, pp. 3998–4011, 2018.
- [2] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Transactions on Computational Imaging*, vol. 3, no. 1, pp. 47–57, March 2017.
- [3] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, Aug 2007.
- [4] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017.
- [5] A. Bovik, *Handbook of Image and Video Processing*. Academic Press, 2010.
- [6] N. Ye, K. Wolski, and R. K. Mantiuk, "Predicting visible image differences under varying display brightness and viewing distance," *Conference on Computer Vision and Pattern Recognition*, 2019.
- [7] T. O. Aydin, R. Mantiuk, and H.-P. Seidel, "Extending quality metrics to full luminance range images," in *Human Vision and Electronic Imaging*. Spie, 2008, pp. 68 060B–10.
- [8] G. Valenzise, F. De Simone, P. Lauga, and F. Dufaux, "Performance evaluation of objective quality metrics for HDR image compression," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 9217, 09 2014.
- [9] M. Azimi, A. Banitalebi Dehkordi, Y. Dong, M. Pourazad, and P. Nasiopoulos, "Evaluating the performance of existing full-reference quality metrics on High Dynamic Range (HDR) video content," *International Conference on Multimedia Signal Processing*, 03 2018.
- [10] P. Hanhart, M. V. Bernardo, M. Pereira, A. M. G. Pinheiro, and T. Ebrahimi, "Benchmarking of objective quality metrics for HDR image quality assessment," *EURASIP Journal on Image and Video Processing*, vol. 2015, no. 1, p. 39, Dec 2015. [Online]. Available: <https://doi.org/10.1186/s13640-015-0091-4>
- [11] E. Zerman, G. Valenzise, and F. Dufaux, "An extensive performance evaluation of full-reference HDR image quality metrics," *Quality and User Experience*, vol. 2, no. 1, p. 5, Apr 2017.
- [12] E. Prashnani, H. Cai, Y. Mostofi, and P. Sen, "Pieapp: Perceptual image-error assessment through pairwise preference," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [13] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [14] V. Hosu, H. Lin, T. Sziranyi, and D. Saupe, "Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment," *IEEE Transactions on Image Processing*, vol. 29, p. 4041–4056, 2020.
- [15] N. Ponomarenko, L. Jin, O. Ieremeiev, V. Lukin, K. Egiazarian, J. Astola, and Benoit, "Image database TID2013: Peculiarities, results and perspectives," *Signal Processing: Image Communication*, vol. 30, pp. 57 – 77, 2015.
- [16] M. Narwaria, M. P. Da Silva, P. Le Callet, and R. Pepion, "Tone mapping-based high-dynamic-range image compression: study of optimization criterion and perceptual quality," *Optical Engineering*, vol. 52, no. 10, 2013.
- [17] H. Sheikh, M. Sabir, and A. Bovik, "A Statistical Evaluation of Recent Full Reference Image Quality Assessment Algorithms," *IEEE Transactions on Image Processing*, vol. 15, no. 11, pp. 3440–3451, 2006.
- [18] P. Korshunov, P. Hanhart, T. Richter, A. Artusi, R. Mantiuk, and T. Ebrahimi, "Subjective quality assessment database of HDR images compressed with JPEG XT," in *2015 Seventh International Workshop on Quality of Multimedia Experience (QoMEX)*, May 2015, pp. 1–6.
- [19] D. Jayaraman, A. Mittal, A. K. Moorthy, and A. C. Bovik, "Objective quality assessment of multiply distorted images," in *2012 Conference Record of the Forty Sixth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, Nov 2012, pp. 1693–1697.
- [20] D. Ghadiyaram and A. C. Bovik, "Massive online crowdsourced study of subjective and objective picture quality," *IEEE Transactions on Image Processing*, vol. 25, no. 1, pp. 372–387, Jan 2016.
- [21] M. Pérez-Ortiz, A. Mikhailiuk, E. Zerman, V. Hulusic, G. Valenzise, and R. K. Mantiuk, "From pairwise comparisons and rating to a unified quality scale," *IEEE Transactions on Image Processing*, vol. 29, pp. 1139–1151, 2019.
- [22] R. Herbrich, T. Minka, and T. Graepel, "Trueskill™: A bayesian skill rating system," in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. C. Platt, and T. Hoffman, Eds. MIT Press, 2007, pp. 569–576.
- [23] R. A. Bradley and M. E. Terry, "The rank analysis of incomplete block designs — I. The method of paired comparisons," *Biometrika*, vol. 39, pp. 324–345, 1952.
- [24] L. L. Thurstone, "A law of comparative judgment," *Psychological Review*, vol. 34, no. 4, pp. 273–286, 1927.
- [25] M. Perez-Ortiz and R. K. Mantiuk, "A practical guide and software for analysing pairwise comparison experiments," *CoRR*, 2017.
- [26] E. Zerman, V. Hulusic, G. Valenzise, R. K. Mantiuk, and F. Dufaux, "The relation between MOS and pairwise comparisons and the importance of cross-content comparisons," *Proc. of Human Vision and Electronic Imaging*, 2018.
- [27] H. Lin, V. Hosu, and D. Saupe, "Kadid-10k: A large-scale artificially distorted iqa database," in *2019 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2019, pp. 1–3.
- [28] P. Ye and D. Doermann, "Active sampling for subjective image quality assessment," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 4249–4256.
- [29] A. Mikhailiuk, C. Wilmot, M. Perez-Ortiz, D. Yue, and R. Mantiuk, "Active sampling for pairwise comparisons via approximate message passing and information gain maximization," in *International Conference on Pattern Recognition*, 2020.
- [30] W. Zhou and E. P. Simoncelli, "Maximum differentiation (mad) competition: a methodology for comparing computational models of perceptual quantities," in *Journal of Vision*, vol. 8, no. 8.1-13., 2008, pp. 586–595.
- [31] K. Ma, Z. Duanmu, Z. Wang, Q. Wu, W. Liu, H. Yong, H. Li, and L. Zhang, "Group maximum differentiation competition: Model comparison with few samples," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 4, pp. 851–864, 2020.
- [32] Y. Deng, C. C. Loy, and X. Tang, "Image aesthetic assessment: An experimental survey," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 80–106, July 2017.
- [33] H. Hadizadeh and I. V. Bajic, "Full-Reference Objective Quality Assessment of Tone-Mapped Images," *IEEE Transactions on Multimedia*, vol. 20, no. 2, pp. 392–404, feb 2018.
- [34] L. Krasula, K. Fliegel, and P. Le Callet, "FFTMI: Features Fusion for Natural Tone-Mapped Images Quality Evaluation," *IEEE Transactions on Multimedia*, vol. 22, no. 8, pp. 2038–2047, aug 2020.
- [35] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich, "HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions," *ACM Trans. Graph.*, vol. 30, no. 4, pp. 40:1–40:14, Jul. 2011.
- [36] K. Wolski, D. Giunchi, N. Ye, P. Didyk, K. Myszkowski, R. Mantiuk, H.-P. Seidel, A. Steed, and R. K. Mantiuk, "Dataset and Metrics for Predicting Local Visible Differences," *ACM Transactions on Graphics*, vol. 37, no. 5, pp. 1–14, nov 2018.
- [37] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.
- [38] M. Narwaria, M. P. D. Silva, and P. L. Callet, "HDR-VQM: An objective quality measure for high dynamic range video," *Signal Processing: Image Communication*, vol. 35, pp. 46 – 60, 2015.
- [39] Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy, and M. Manohara, "Toward A Practical Perceptual Video Quality Metric," The NETFLIX Tech Blog, Tech. Rep., 2016.
- [40] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, Dec 2012.
- [41] N Venkatanath, D Praneeth, Bh Maruthi Chandrasekhar , S. S. Channappayya, and S. S. Medasani, "Blind image quality evaluation using perception based features," in *2015 Twenty First National Conference on Communications (NCC)*, Feb 2015, pp. 1–6.
- [42] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "completely blind" image quality analyzer," *IEEE Signal Processing Letters*, vol. 20, pp. 209–212, 2013.

- [43] J. Kim and S. Lee, "Deep learning of human visual sensitivity in image quality assessment framework," *Computer Vision and Pattern Recognition (CVPR)*, pp. 1969–1977, 2017.
- [44] F. Gao, Y. Wang, P. Li, M. Tan, J. Yu, and Y. Zhu, "Deepsim: Deep similarity for image quality assessment," *Neurocomputing*, vol. 257, pp. 104 – 114, 2017.
- [45] J. Kim, H. Zeng, D. Ghadiyaram, S. Lee, L. Zhang, and A. C. Bovik, "Deep convolutional neural models for picture-quality prediction: Challenges and solutions to data-driven image quality assessment," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 130–141, Nov 2017.
- [46] X. Liu, J. van de Weijer, and A. D. Bagdanov, "RankIQA: Learning from rankings for no-reference image quality assessment," *CoRR*, vol. abs/1707.08347, 2017.
- [47] S. A. Amirshahi, M. Pedersen, and S. X. Yu, "Image quality assessment by comparing CNN features between images," *Journal of Imaging Science and Technology*, vol. 60, no. 6, 2016.
- [48] Y. Li, L. M. Po, L. Feng, and F. Yuan, "No-reference image quality assessment with deep convolutional neural networks," in *2016 IEEE International Conference on Digital Signal Processing (DSP)*, Oct 2016, pp. 685–689.
- [49] S. Bianco, L. Celona, P. Napoletano, and R. Schettini, "On the use of deep learning for blind image quality assessment," *Signal, Image and Video Processing*, vol. 12, no. 2, pp. 355–362, Feb 2018.
- [50] H. Lin, V. Hosu, and D. Saupe, "DeepFL-IQA: Weak Supervision for Deep IQA Feature Learning," *arXiv preprint arXiv:2001.08113*, 2020.
- [51] K. Ma, W. Liu, K. Zhang, Z. Duanmu, Z. Wang, and W. Zuo, "End-to-end blind image quality assessment using deep neural networks," *IEEE Transactions on Image Processing*, vol. 27, no. 3, pp. 1202–1213, March 2018.
- [52] A. Aldahdooh, E. Masala, O. Janssens, G. Van Wallendael, M. Barkowsky, and P. Le Callet, "Improved performance measures for video quality assessment algorithms using training and validation sets," *IEEE Transactions on Multimedia*, vol. 21, no. 8, pp. 2026–2041, 2019.
- [53] R. R. Hainich, *Perceptual display calibration*. CRC Press, 2016.
- [54] R. K. Mantiuk, "Practicalities of predicting quality of high dynamic range images and video," in *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, sep 2016, pp. 904–908.
- [55] S. Jia, Y. Zhang, D. Agrafiotis, and D. Bull, "Blind high dynamic range image quality assessment using deep learning," in *2017 IEEE International Conference on Image Processing (ICIP)*, Sept 2017, pp. 765–769.
- [56] G. Yue, C. Hou, K. Gu, N. Ling, and B. Li, "Analysis of structural characteristics for quality assessment of multiply distorted images," *IEEE Transactions on Multimedia*, vol. 20, no. 10, pp. 2722–2732, 2018.
- [57] B. Yan, B. Bare, and W. Tan, "Naturalness-aware deep no-reference image quality assessment," *IEEE Transactions on Multimedia*, vol. 21, no. 10, pp. 2603–2615, 2019.
- [58] X. Min, K. Gu, G. Zhai, J. Liu, X. Yang, and C. W. Chen, "Blind quality assessment based on pseudo-reference image," *IEEE Transactions on Multimedia*, vol. 20, no. 8, pp. 2049–2062, 2018.
- [59] X. Min, G. Zhai, K. Gu, Y. Zhu, J. Zhou, G. Guo, X. Yang, X. Guan, and W. Zhang, "Quality evaluation of image dehazing methods using synthetic hazy images," *IEEE Transactions on Multimedia*, vol. 21, no. 9, pp. 2319–2333, 2019.
- [60] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," *arXiv*, 2018.
- [61] A. Mikhailiuk, M. Pérez-Ortiz, and R. K. Mantiuk, "Psychometric scaling of TID2013 dataset," in *International Conference on Quality of Multimedia Experience (QoMEX)*, 2018.
- [62] S. Wuerger, M. Ashraf, M. Kim, J. Martinovic, M. Pérez-Ortiz, and R. K. Mantiuk, "Spatio-chromatic contrast sensitivity under mesopic and photopic light levels," *Journal of Vision*, vol. 20, no. 4, p. 23, apr 2020.
- [63] S. Voran, "An iterated nested least-squares algorithm for fitting multiple data sets," *NASA STI/Recon Technical Report N*, 10 2002.
- [64] ITU-R, "Parameter values for the hdtv standards for production and international programme exchange," ITU-R Recommendation BT.500-13, Jun 2015.
- [65] L. Zhang, L. Zhang, X. Mou, and D. Zhang, "FSIM: A feature similarity index for image quality assessment," *IEEE Transactions on Image Processing*, vol. 20, no. 8, pp. 2378–2386, Aug 2011.
- [66] M. Narwaria, R. K. Mantiuk, M. P. Da Silva, and P. Le Callet, "HDR-VDP-2.2: a calibrated method for objective quality prediction of high-dynamic range and standard images," *Journal of Electronic Imaging*, vol. 24, no. 1, p. 010501, jan 2015.
- [67] V. Laparra, A. Berardino, J. Ballé, and E. P. Simoncelli, "Perceptually optimized image rendering," *J. Opt. Soc. Am. A*, vol. 34, no. 9, pp. 1511–1525, Sep 2017.
- [68] A. Mikhailiuk, N. Ye, and R. Mantiuk, "The effect of display brightness and viewing distance: a dataset for visually lossless image compression," in *Human Vision and Electronic Imaging (HVEI)*, 2021.