Quality-Aware Image-Text Alignment
for Opinion-Unaware Image Quality Assessment

Lorenzo Agnolucci   Leonardo Galteri   Marco Bertini

University of Florence

[name.surname]@unifi.it

Abstract

No-Reference Image Quality Assessment (NR-IQA) focuses on designing methods to measure image quality in alignment with human perception when a high-quality reference image is unavailable. Most state-of-the-art NR-IQA approaches are opinion-aware, i.e. they require human annotations for training. This dependency limits their scalability and broad applicability. To overcome this limitation, we propose QualiCLIP (Quality-aware CLIP), a CLIP-based self-supervised opinion-unaware approach that does not require human opinions. In particular, we introduce a quality-aware image-text alignment strategy to make CLIP generate quality-aware image representations. Starting from pristine images, we synthetically degrade them with increasing levels of intensity. Then, we train CLIP to rank these degraded images based on their similarity to quality-related antonym text prompts. At the same time, we force CLIP to generate consistent representations for images with similar content and the same level of degradation. Our experiments show that the proposed method improves over existing opinion-unaware approaches across multiple datasets with diverse distortion types. Moreover, despite not requiring human annotations, QualiCLIP achieves excellent performance against supervised opinion-aware methods in cross-dataset experiments, thus demonstrating remarkable generalization capabilities. The code and the model are publicly available at https://github.com/miccunifi/QualiCLIP.

1. Introduction

Image Quality Assessment (IQA) aims to automatically evaluate the quality of images in accordance with human judgments represented by a Mean Opinion Score (MOS). Specifically, No-Reference IQA (NR-IQA) focuses on developing methods that do not require a high-quality reference image and that are consequently more easily applicable in real-world scenarios. NR-IQA plays a critical role in diverse industries and research domains. For exam-
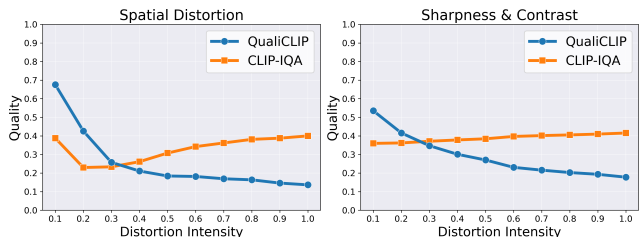

Figure 1. Comparison between the image quality scores predicted by CLIP-IQA [47] and the proposed QualiCLIP for increasing distortion intensities of different types of synthetic degradation. We average the results of 1000 randomly sampled images from the KonIQ-10k [11] dataset. Our method corresponds to a stronger inverse correlation between the predicted quality scores and the severity of the degradation. The distortion intensities are scaled between 0 and 1 for clearer visualization.

ple, given the large number of photos that are captured and shared daily on social media platforms, it is imperative to design approaches that can measure image quality objectively to store and process these images effectively.

Most NR-IQA methods are opinion-aware, *i.e.* they require human-labeled MOS as supervision during the training process [35, 42, 58, 60]. Some approaches, such as HyperIQA [42] or LIQE [58], directly train the model parameters on IQA datasets. Other methods, namely QPT [60] or Re-IQA [35], pre-train an encoder on unlabeled data via self-supervised learning and then either fine-tune the encoder weights or train a linear regressor using MOS. However, annotating IQA datasets is expensive and resource-intensive, as several human ratings are needed for each image for its MOS to be reliable. For example, the FLIVE dataset [53], which contains 40K real-world images, required about 4M ratings, up to 50 for a single image. The need for human annotations significantly hinders the scalability of opinion-aware approaches. In addition, these methods show limited generalization capabilities and thus applicability to scenarios where labeled data is unavailable, as their performance considerably deteriorates on unseen

datasets. To remove the requirement for expensive MOS, several opinion-unaware approaches have been proposed [3, 9, 25]. For instance, CL-MI [3] introduces a two-stage self-supervised approach that employs two different training strategies for synthetically and authentically degraded images. Nevertheless, existing opinion-unaware methods achieve significantly lower performance than opinion-aware approaches in cross-dataset experiments, thus exhibiting limited applicability.

In this context, we propose to leverage recent advancements in Vision-Language Models (VLMs) by presenting a self-supervised opinion-unaware approach based on CLIP [32]. Recently, CLIP-based methods achieved promising performance in NR-IQA [41, 47, 58]. For example, CLIP-IQA [47] proposes to compute the quality score by measuring the similarity between an image and two quality-related antonym prompts without any task-specific training. However, off-the-shelf CLIP models struggle to generate quality-aware image representations [16, 58], as they focus more on high-level semantics than low-level image characteristics, such as noise and blur. To highlight this issue, we randomly sample 1000 images from the KonIQ-10k dataset [11] and synthetically degrade them with several distortions using increasing levels of intensity. Then, we compute the quality score of each image through CLIP-IQA and average the results. We expect the more degraded versions of the images to correspond to lower quality scores. However, Fig. 1 shows that CLIP-IQA exhibits a low correlation between the predicted quality and the degree of distortion. This finding indicates that CLIP is *not* intrinsically quality-aware.

To address this issue, we propose a quality-aware image-text alignment strategy that relies on self-supervised learning to remove the need for human annotations. We start by synthetically degrading pairs of pristine images using increasing levels of intensity. Then, we measure the similarity between each image and *antonym prompts* related to image quality, such as "*Good photo*" and "*Bad photo*". We refer to these prompts as *positive* and *negative* prompts, respectively. Finally, we employ a training strategy based on a margin ranking loss [16, 19] that allows us to achieve two objectives. First, we want CLIP to generate consistent representations for images having similar content and comparable quality, *i.e.* exhibiting the same amount of distortion. Second, the similarity between the positive (negative) prompt and the increasingly degraded versions of the images must correlate inversely (directly) with the intensity of the distortion. Our approach, named QualiCLIP (Quality-aware CLIP), is both self-supervised and opinion-unaware, as we do not rely on any form of supervision – especially MOS – at any step of the training process. Thanks to our training strategy, the image-text alignment in CLIP's embedding space prioritizes low-level image characteristics over high-level semantics. Consequently, QualiCLIP gener-

ates image representations whose similarity to the antonym prompts correlates with the inherent quality of the images, as illustrated in Fig. 1.

The experiments demonstrate that the proposed approach improves over existing opinion-unaware methods across multiple datasets encompassing various degradations. Furthermore, QualiCLIP is the only opinion-unaware approach that consistently obtains remarkable results even when compared against supervised opinion-aware techniques in the cross-dataset setting. The strong and robust performance of our model across different datasets highlights its commendable generalization capabilities and suitability for real-world applications.

We summarize our contributions as follows:
- We propose QualiCLIP, a CLIP-based self-supervised opinion-unaware approach for NR-IQA that does not require any type of supervision, especially MOS;
- We introduce a quality-aware image-text alignment strategy based on ranking increasingly degraded pairs of images according to their similarity to quality-related antonym prompts. After training, QualiCLIP generates quality-aware image representations;
- Our method improves on existing opinion-unaware approaches across multiple datasets and achieves excellent results even when compared to supervised opinion-aware techniques in cross-dataset experiments.

2. Related Work

No-Reference Image Quality Assessment Due to its wide range of applications in real-world scenarios, in recent years research on NR-IQA has gained significant momentum [1, 3, 26, 29, 42]. Traditional methods [28, 29, 55] rely on extracting hand-crafted image features to derive quality scores. Such approaches achieve promising performance on synthetic datasets but struggle on images with authentic distortions. More recently, several methods relying on supervised learning have been introduced [1, 7, 26, 35, 38, 42, 52]. Some approaches directly employ MOS during model training [7, 42, 52]. For example, HyperIQA [42] proposes a self-adaptive hypernetwork that separates content understanding from quality prediction. Another research direction involves pre-training an encoder on unlabeled images via self-supervised learning. Then, the image representations are mapped to quality scores by fine-tuning the encoder weights [60] or training a linear regressor [1, 26, 35] on human annotations. For instance, QPT [60] and Re-IQA [35] use a contrastive loss to train the encoder to separate between images degraded with different types and degrees of distortion. Despite their impressive performance, the scalability and applicability of supervised methods are limited by their need for costly human annotations. This requirement is removed by opinion-unaware approaches [3, 25, 29, 30, 40, 55]. Some of

them, such as NIQE [29], are based on natural scene statistics [29, 55], while others employ self-supervised learning [3, 9, 25, 40]. For example, CL-MI [3] pre-trains an encoder on synthetic data and then fine-tunes it on authentic images via a mutual information-based loss. Nevertheless, existing opinion-unaware approaches fall behind supervised methods in cross-dataset experiments. In contrast, despite not requiring MOS, our method achieves remarkable performance on unseen datasets even when considering opinion-aware techniques.

Vision-Language Models for NR-IQA VLMs, such as CLIP [32], have achieved impressive performance in several low-level vision tasks, including image and video restoration [2, 16, 21] and quality assessment [41, 47–51, 58]. CLIP-IQA [47] studies the capabilities of CLIP in assessing the quality and abstract perception of images without task-specific training. In addition, the authors train a model named CLIP-IQA$^+$ based on learning two antonym prompts using MOS. LIQE [58] fine-tunes CLIP in a supervised way with a multi-task learning approach exploiting scene and distortion-type information. Recently, several methods based on Multimodal Large Language Models (MLLMs) have been proposed [50, 51, 54]. While these approaches achieve impressive results, they require significant computational resources due to the high demands of MLLMs. Among VLM-based methods for NR-IQA, the most similar to our work is GRepQ [41], which trains a low-level and a high-level CLIP-based encoder via self-supervised learning. CLIP is fine-tuned by separating higher and lower-quality groups of images within the same batch with a contrastive loss depending on their predicted quality, obtained by measuring their similarity to antonym prompts. GRepQ predicts the final quality score by combining the features of the two encoders and feeding them as input to a linear regressor, which is trained on IQA datasets using MOS. In contrast, we present a CLIP-only self-supervised approach that removes the need for a low-level encoder. We propose to synthetically degrade pairs of images with increasing levels of intensity and make our model learn to rank them through a ranking loss according to their degree of distortion. The ranking is based directly on the similarity between the text features and each of the antonym prompts, instead of relying on the predicted quality as GRepQ. Also, differently from GRepQ, we do not require any form of supervision at any step of our approach.

Learning to rank Learning to rank images has proven to be an effective technique for image quality and aesthetics assessment [7, 13, 19, 25, 34, 44]. For instance, VILA [13] tackles image aesthetics assessment by training a learnable residual projection on top of CLIP to rank the quality of a single pair of images according to their MOS. Another example is RankIQA [19], which involves synthetically degrading images with varying degrees using dataset-specific

distortions. Then, for each IQA dataset, the authors first pre-train a Siamese network by ranking the images based on their level of degradation and then fine-tune it with the MOS. In our work, we employ a given set of distortions to degrade pairs of crops with increasing levels of intensity. Then, we leverage the information provided by their implicit quality ranking to train a model to rank them according to their similarity to antonym prompts. In this way, our method does not require fine-tuning on ground-truth labels.

3. Proposed Approach

We propose a quality-aware image-text alignment strategy to make CLIP generate quality-aware image representations. First, we synthetically degrade pairs of crops with increasing levels of intensity. Then, we fine-tune CLIP's image encoder by ranking the similarity between two antonym prompts and the progressively distorted image pairs based on their degree of degradation, while guaranteeing consistent representations for each pair of crops. We keep CLIP's text encoder fixed. We use a ResNet50 [10] as the backbone for CLIP. We do not employ any supervision – particularly MOS – at any step of the training process. Due to space limitations, we provide the implementation details in the supplementary material.

3.1. CLIP Preliminaries

CLIP (Contrastive Language-Image Pre-training) [32] is a vision-language model trained on a large-scale dataset to semantically align images and corresponding text captions in a shared embedding space. The authors employ a contrastive loss that maximizes the similarity between paired image-text samples while minimizing the similarity with all the other samples within a batch. CLIP comprises an image encoder ψ_I and a text encoder ψ_T. Given an image I, the image encoder extracts its feature representation $x = \psi_I(I) \in \mathbb{R}^d$, where d is the dimension of CLIP's embedding space. For a given text caption T, each tokenized word is mapped to the token embedding space $\mathcal{W}$ through a word embedding layer E_w. Then, the text encoder ψ_T is used to generate the textual feature representation $y = \psi_T(E_w(T)) \in \mathbb{R}^d$ from the token embeddings. Thanks to its training strategy, CLIP generates similar representations within the common embedding space for images and text expressing the same concepts.

3.2. Synthetic Degradation with Increasing Levels of Intensity

Although authentic distortions cannot be perfectly replicated synthetically, prior studies have shown that synthetic distortions remain effective for training self-supervised NR-IQA models that generalize well to real-world images [1, 26, 35, 60]. Following these works, we synthetically
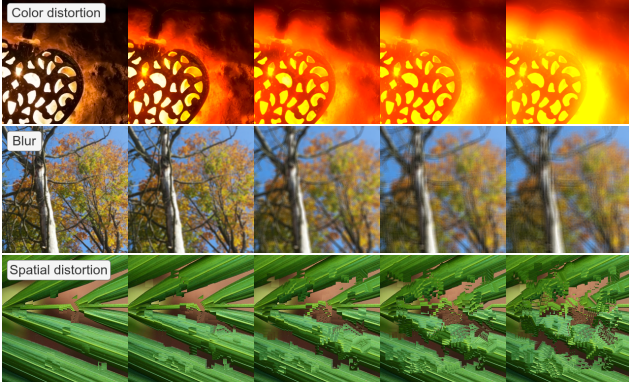
Figure 2. Examples of synthetic degradations for five increasing levels of intensity.

degrade unlabeled pristine images using progressively increasing levels of intensity. In this way, we can train our model in a self-supervised way to rank the different versions of each image according to the severity of their degradation. Following [1], we consider 24 distinct degradation types spanning the 7 distortion groups defined by the KADID-10k dataset [18]: 1) Brightness change; 2) Blur; 3) Spatial distortions; 4) Noise; 5) Color distortions; 6) Compression; 7) Sharpness & contrast. Each distortion has $L=5$ levels of intensity. See the supplementary material for more details on the specific degradation types. Fig. 2 shows some examples of degraded images for varying degrees of intensity.

We start from the 140K pristine images of the KADIS-700k dataset [18]. For each image, we extract a pair of random overlapping crops. Then, we randomly sample $D = 1$ distortion groups and a degradation within each group. We apply the D distortions to both crops using $L = 5$ distinct levels of intensity, resulting in L pairs of equally degraded crops, one for each level. Contrary to RankIQA [19], we obtain two images for each degree of distortion, as depicted in the leftmost part of Fig. 3. Given two such pairs of crops, we can infer which has a higher quality based on the corresponding level of degradation. We leverage this information to train our model with a ranking loss.

3.3. Quality-Aware Image-Text Alignment

As Fig. 1 shows, CLIP struggles to generate accurate quality-aware image representations that reflect the severity of the degradation. To address this issue, we propose a quality-aware image-text alignment strategy to fine-tune CLIP's image encoder. The idea of our approach is that given two degraded versions of the same image, a positive prompt referring to high image quality – such as "*Good photo*" – should be more similar to the less degraded version. The opposite consideration applies considering a negative prompt related to low image quality, such as "*Bad photo*". At the same time, two images with overlapping

content and equal degree of degradation should have comparable similarities to such a pair of quality-related antonym prompts. Note that, given two unlabeled images with completely different content, we can not make any assumptions about their relative quality [4], or, in other words, their similarity to the prompts. Our training strategy leverages multiple pairs of increasingly degraded images to achieve two objectives: *O1)*: we want CLIP to generate consistent representations for images with similar content and comparable quality, *i.e.* showing the same amount of distortion; *O2)*: the similarity between the positive (negative) prompt and the distinct versions of the images must correlate inversely (directly) with the corresponding level of degradation.

Let I_i^1 and I_i^2 be the i-th pair of increasingly degraded crops obtained as detailed in Sec. 3.2, where $i \in \{1, \ldots, L\}$ and $L = 5$ is the number of considered distortion levels. For $i, j \in \{1, \ldots, L\}$ with $j > i$, the j-th pair of crops is more degraded than the i-th one. Given each pair of crops, we extract the corresponding features through CLIP's image encoder ψ_I, resulting in $x_i^1 = \psi_I(I_i^1)$ and $x_i^2 = \psi_I(I_i^2)$. Similarly to [47], we remove the positional embedding to relax CLIP's requirement of fixed-size inputs. Let T_p and T_n be a pair of antonym prompts related to image quality, such as "*Good photo*" and "*Bad photo*". We refer to T_p and T_n as *positive* and *negative* prompts, respectively. In practice, we use multiple pairs of antonym prompts, similar to [48, 49]. We use CLIP's text encoder ψ_T to extract the text features associated with the prompts, obtaining $t_p = \psi_T(T_p)$ and $t_n = \psi_T(T_n)$. We normalize both the image and text features to have a unit L_2-norm.

To achieve objective *O1*, we propose to employ a consistency loss term to guarantee that the similarity between the features of the prompts and those of each of the two images composing each degraded pair is comparable. We assume that two overlapping crops extracted from the same image have a comparable quality, analogously to [35, 60]. We rely on a margin ranking loss [7, 16, 19] defined as:

$$\mathcal{L}_{cons} = \sum_{i=1}^{L} \big[\max(0, \big| c(x_i^1, t_p) - c(x_i^2, t_p) \big| - m_{cons}) \\ + \max(0, \big| c(x_i^1, t_n) - c(x_i^2, t_n) \big| - m_{cons}) \big], \quad (1)$$

where $c(\cdot)$ stands for the cosine similarity and the margin m_{cons} is a hyperparameter. Intuitively, m_{cons} must be small enough to force the similarities between the prompts and each of the two crops to be comparable. With CLIP, the cosine similarity between each image-prompt pair is in $[0, 1]$.

Given the i-th level of synthetic degradation, with $i \in \{1, \ldots, L\}$, we assume that the quality of the two distorted crops of the i-th pair is higher than that of the two images composing the $(i + 1)$-th one, analogously to [19, 34]. Thus, we enforce that the similarity between the features of the *positive* prompt and those of two crops is *higher* than
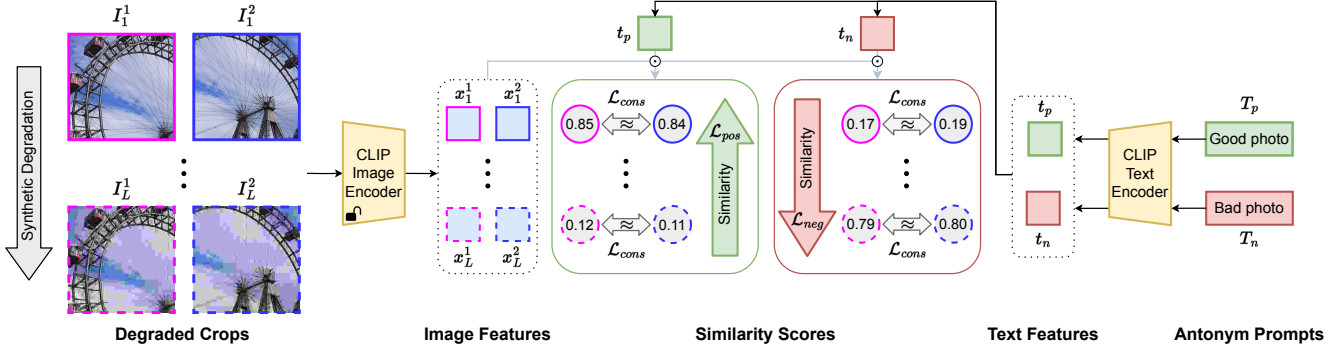
Figure 3. Overview of the proposed quality-aware image-text alignment strategy. Starting from a pair of two random overlapping crops from a pristine image, we synthetically degrade them with L increasing levels of intensity, resulting in L pairs. Then, given two quality-related antonym prompts T_p and T_n, we fine-tune CLIP's image encoder with three margin ranking losses ($\mathcal{L}_{cons}$, $\mathcal{L}_{pos}$, $\mathcal{L}_{neg}$) by considering the similarity between the prompts and the degraded crops. Specifically, we use $\mathcal{L}_{cons}$ to force CLIP to generate consistent representations for the crops belonging to each pair, since they exhibit similar content and the same degree of distortion. At the same time, we make the similarity between the prompt T_p (or T_n) and the increasingly degraded versions of the crops correlate inversely (or directly) with the intensity of the distortion through $\mathcal{L}_{pos}$ (or $\mathcal{L}_{neg}$).

when considering more degraded versions of the two crops. Specifically, we define a margin ranking loss as:

$$\mathcal{L}_{pos} = \sum_{i=1}^{L-1} \sum_{j=i+1}^{L} \sum_{k=1}^{2} \big[\max(0, c(x_j^k, t_p) - c(x_i^1, t_p) + m_{rank}) + \max(0, c(x_j^k, t_p) - c(x_i^2, t_p) + m_{rank}) \big],$$

(2)

where the margin m_{rank} is a hyperparameter. The opposite of the consideration made above applies when we take into account the negative prompt. Therefore, we add a loss term to impose that the similarity between the features of the *negative* prompt and those of two crops is *lower* than when considering more degraded versions of the two crops:

$$\mathcal{L}_{neg} = \sum_{i=1}^{L-1} \sum_{j=i+1}^{L} \sum_{k=1}^{2} \big[\max(0, c(x_i^1, t_n) - c(x_j^k, t_n) + m_{rank}) + \max(0, c(x_i^2, t_n) - c(x_j^k, t_n) + m_{rank}) \big].$$

(3)

Intuitively, m_{rank} must be large enough to make the similarities between the prompts and the increasingly degraded versions of the two crops noticeably different. Using the combination of $\mathcal{L}_{pos}$ and $\mathcal{L}_{neg}$ achieves objective *O2*.

The final training loss is given by:

$$\mathcal{L} = \lambda_{cons}\mathcal{L}_{cons} + \lambda_{pos}\mathcal{L}_{pos} + \lambda_{neg}\mathcal{L}_{neg},$$

(4)

where λ_{cons}, λ_{pos}, and λ_{neg} represent the loss weights. Fig. 3 shows an overview of our training strategy. Given that we do not employ any MOS, our approach is both self-supervised and opinion-unaware. Thanks to the proposed training strategy, CLIP learns to align images and texts based more on low-level characteristics, such as noise and blur, rather than high-level semantics. As a result, the

similarity between the antonym prompts and the image representations obtained by QualiCLIP correlates with the inherent quality of the images, as shown in Fig. 1.

At inference time, given an image I, we extract its features x using CLIP's image encoder. Then, we compute the cosine similarity between x and the features t_p and t_n of the antonym prompts, resulting in s_p and s_n. Finally, we obtain the final quality score $q \in [0, 1]$ as:

$$q = \frac{e^{(s_p/\tau)}}{e^{(s_p/\tau)} + e^{(s_n/\tau)}},$$

(5)

where τ is a temperature hyperparameter. Note that, since we keep CLIP's text encoder weights frozen, we need to compute the text features of the antonym prompts only once, and we can use them both for training and inference. Therefore, at inference time, the computational cost of our method is the same as that of an image-encoder-only model with the same backbone.

Discussion Our quality-aware image-text alignment strategy stems from the inherent limitations of applying common self-supervised NR-IQA training techniques to CLIP. Prior methods, such as QPT [60] and Re-IQA [35], train an encoder using a contrastive loss that maximizes the similarity between the representations of crops from the same degraded image, while minimizing the similarity with the representations of crops coming from different images within the same batch. While this strategy proved its effectiveness for image-encoder-only models, applying it to CLIP's image encoder would introduce a significant mismatch with CLIP's training process [32]. Indeed, CLIP is trained using an inter-modal (*i.e.* image-text) objective, aligning the features – extracted with the image and text encoders – of corresponding images and texts within a common embedding space. Consequently, fine-tuning CLIP's image en-

5

		Authentic Datasets								Image Restoration Datasets				AIGC Datasets			
		KonIQ-10k		CLIVE		FLIVE		SPAQ		CVIU		PIPAL		AGIQA-1K		AGIQA-3K	
Method	OU	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
CONTRIQUE [26]	✗	0.874	0.882	0.806	0.838	0.596	0.648	0.910	0.912	0.912	0.916	0.581	0.611	0.799	0.855	0.817	0.879
Re-IQA [35]	✗	0.883	0.887	0.783	0.797	0.623	0.685	0.909	0.912	0.887	0.895	0.568	0.587	0.783	0.84	0.811	0.874
ARNIQA [1]	✗	0.869	0.883	0.797	0.823	0.595	0.670	0.904	0.909	0.919	0.926	0.634	0.666	0.768	0.849	0.803	0.881
CLIP-IQA$^+$ [47]	✗	0.873	0.890	0.815	0.849	0.602	0.662	0.901	0.904	0.925	0.918	0.552	0.558	0.817	0.855	0.844	0.894
GRepQ [41]	✗	0.882	0.883	0.793	0.813	0.576	0.611	0.902	0.903	0.882	0.857	0.554	0.568	0.740	0.797	0.807	0.858
NIQE [29]	✓	0.527	0.532	0.446	0.469	0.203	0.262	0.700	0.712	0.685	0.696	0.167	0.181	0.623	0.721	0.510	0.526
IL-NIQE [55]	✓	0.511	0.532	0.427	0.481	0.234	0.286	0.713	0.724	0.595	0.668	0.231	0.220	0.645	0.757	0.528	0.544
CL-MI [3]	✓	0.664	0.653	0.494	0.480	0.290	0.322	0.701	0.701	0.433	0.481	0.281	0.282	0.474	0.621	0.591	0.665
MDFS [30]	✓	0.739	0.746	0.427	0.467	0.360	0.407	0.805	0.812	0.702	0.742	0.255	0.262	0.734	**0.818**	**0.697**	0.723
CONTRIQUE-OU [26]	✓	0.651	0.637	0.348	0.364	0.321	0.341	0.677	0.685	0.610	0.715	0.225	0.218	0.644	0.717	0.465	0.489
Re-IQA-OU [35]	✓	0.580	0.568	0.379	0.388	0.350	0.341	0.613	0.616	0.648	0.662	0.152	0.151	0.644	0.764	0.366	0.430
ARNIQA-OU [1]	✓	0.746	0.762	0.468	0.557	0.428	0.468	0.788	0.797	0.673	0.712	0.327	0.331	0.667	0.786	0.614	0.679
CLIP-IQA [47]	✓	0.695	0.730	0.656	0.670	0.344	0.473	0.733	0.734	0.613	0.642	0.274	0.279	0.498	0.651	0.638	0.711
GRepQ-OU [41]	✓	0.776	0.798	0.722	0.766	0.289	0.423	0.802	0.805	0.748	0.785	**0.434**	**0.419**	0.161	0.530	0.613	0.734
QualiCLIP	✓	**0.817**	**0.838**	**0.725**	**0.802**	**0.442**	**0.556**	**0.841**	**0.851**	**0.840**	**0.852**	0.410	0.424	**0.736**	0.805	0.667	**0.735**

Table 1. Quantitative results of the *zero-shot* evaluation setting. OU stands for Opinion-Unaware. Best and second-best scores for OU methods are highlighted in bold and underlined, respectively. The suffix "-OU" indicates approaches modified to be opinion-unaware (see Sec. 4.1). For reference, we report the performance of supervised methods (*i.e.* OU=✗) trained on the training split of each testing dataset.

coder by considering only intra-modal (*i.e.* image-image) similarities without exploiting its alignment with the text encoder contradicts its design [27, 45]. For this reason, we propose to train our model by employing image-text similarities to leverage CLIP's inherent inter-modal alignment. Additionally, using a contrastive loss to maximize (or minimize) the similarity of the antonym prompts with multiple different training samples within the same batch would correspond to making assumptions on the relative quality of unlabeled images with completely different content, which is unfeasible [60]. Instead, by relying on a ranking loss that only considers progressively degraded versions of the same image, we can leverage their inherent quality ranking as supervision to train our model in an effective way.

4. Experimental Results

4.1. Quantitative Results

We conduct several experiments to compare the performance of QualiCLIP with existing opinion-unaware and opinion-aware methods. In the supplementary material, we also study the robustness and explainability of our model.

Evaluation protocol We evaluate the performance using Spearman's Rank-order Correlation Coefficient (SRCC) and Pearson's Linear Correlation Coefficient (PLCC), which measure prediction monotonicity and accuracy, respectively. Higher values of SRCC and PLCC correspond to better results. Following [8], we pass the quality predictions through a four-parameter logistic non-linearity before computing PLCC.

We evaluate our method on several IQA datasets, each containing images annotated with human judgments of picture quality in the form of MOS. These datasets feature images with various types of distortions, including authentic degradations, artifacts from image restoration methods, and AI-generated content. Specifically, we consider four authentic datasets: KonIQ-10k [11], CLIVE [6], FLIVE [53], and SPAQ [5]; two image restoration datasets: CVIU [22] and PIPAL [12]; and two AIGC datasets: AGIQA-1K [59] and AGIQA-3K [15]. Additional details on the datasets are provided in the supplementary material, where we also report experiments on images with synthetic distortions. Following [1, 26, 35], we randomly split the datasets into 70% for training, 10% for validation, and 20% for testing. For datasets that include reference images, namely the image restoration and synthetic datasets, we ensure that splits are made based on reference images to prevent content overlap. To mitigate selection bias in the training set, we repeat the training/testing process 10 times and report the median results. Due to its large size, for FLIVE, we follow [1, 26, 35] and use only the official train-validation-test split [53].

We compare our approach to state-of-the-art methods in two settings: *zero-shot* and *cross-dataset*. Our method remains consistent across both settings; the only variation lies in the competing methods. For a fair comparison, we compute the results of the baselines using our evaluation protocol. For each baseline, we employ the official pre-trained model when available, or, otherwise, train the model following the procedure described in the original paper. In the *zero-shot* setting, we compare with existing opinion-unaware methods. In addition, following [3], we consider opinion-aware approaches that can be modified to function without requiring MOS (indicated with the suffix "-OU"). In particular, for GRepQ [41], we follow the zero-shot strategy detailed in the original paper. For methods based on an image encoder and a linear regressor, such as CONTRIQUE [26], we extract the image features via the pre-trained encoder and then employ a NIQE-style framework

| | | Authentic Datasets | | | | | | Image Restoration Datasets | | | | AIGC Datasets | | | |
| | | KonIQ-10k | | CLIVE | | SPAQ | | CVIU | | PIPAL | | AGIQA-1K | | AGIQA-3K | |
Method	OU	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
HyperIQA [42]	✗	0.738	0.735	0.699	0.722	0.649	0.650	0.361	0.400	0.289	0.95	0.403	0.582	0.490	0.570
DBCNN [57]	✗	0.804	0.823	0.660	0.698	0.808	0.817	0.786	0.803	0.356	0.374	0.525	0.726	0.510	0.657
TReS [7]	✗	0.759	0.758	0.696	0.727	0.739	0.735	0.466	0.510	0.326	0.328	0.392	0.585	0.515	0.569
LIQE [58]	✗	0.813	0.801	0.743	0.756	0.713	0.712	0.560	0.625	0.349	0.393	0.450	0.663	0.563	0.633
QCN [39]	✗	0.728	0.789	0.653	0.731	0.815	0.819	0.753	0.757	0.370	0.382	0.458	0.705	0.518	0.595
CONTRIQUE [26]	✗	0.777	0.777	0.734	0.766	0.820	0.827	0.616	0.681	0.357	0.390	0.485	0.663	0.464	0.540
Re-IQA [35]	✗	0.796	0.827	0.690	0.761	0.825	0.830	0.714	0.741	0.306	0.302	0.406	0.675	0.554	0.625
ARNIQA [1]	✗	0.798	0.810	0.699	0.765	0.837	0.846	0.632	0.673	0.351	0.361	0.437	0.682	0.487	0.585
CLIP-IQA+ [47]	✗	0.783	0.798	0.694	0.741	0.755	0.754	0.747	0.770	0.319	0.332	0.502	0.595	0.505	0.518
GRepQ [41]	✗	0.810	0.818	**0.751**	0.772	0.829	0.838	0.799	0.816	**0.448**	**0.460**	0.321	0.610	0.581	0.697
QualiCLIP	✓	**0.817**	**0.838**	0.725	**0.802**	**0.841**	**0.851**	**0.840**	**0.852**	0.410	0.424	**0.736**	**0.805**	**0.667**	**0.735**

Table 2. Quantitative results of the *cross-dataset* evaluation setting. We employ the FLIVE [53] dataset to train the supervised methods. OU stands for Opinion-Unaware. Best and second-best scores are highlighted in bold and underlined, respectively.

to predict the quality scores, similar to [3]. In the *cross-dataset* setting, we evaluate the generalization capabilities of our model by comparing it with supervised opinion-aware methods on testing datasets different from the training one. Due to its large scale, we employ FLIVE as the training dataset for the baselines. Additionally, we report results using CLIVE and PIPAL in the supplementary material. For a fair comparison, we train LIQE [58] using a ResNet50 backbone and restrict our analysis to methods that do not rely on an MLLM, as these models entail substantial computational demands.

Zero-shot setting We report the results for the *zero-shot* setting in Tab. 1. Our approach achieves the best performance on 13 out of 16 evaluation metrics and ranks second on the remaining 3, with SRCC improvements over the best-performing baseline of up to 9.2%, observed on the CVIU dataset. Notably, QualiCLIP sets the new state of the art for opinion-unaware approaches on authentic and image restoration datasets, proving the effectiveness of our training strategy. The improvement over CLIP-IQA highlights that our model generates more accurate quality-aware image representations than off-the-shelf CLIP models. Compared to GRepQ-OU, which is the strongest existing approach in most scenarios, the proposed method obtains better results on all the testing datasets excluding PIPAL. Moreover, while GRepQ-OU combines a low-level encoder with a high-level fine-tuned CLIP-based encoder, QualiCLIP relies solely on CLIP, making it more straightforward and efficient. For reference, Tab. 1 also includes the performance of supervised opinion-aware methods trained on the training split of each testing dataset. We observe that all opinion-unaware methods fall behind supervised opinion-aware approaches when a training set is available, showing that there is still room for improving the performance of opinion-unaware models.

Cross-dataset setting Tab. 2 shows the results for the *cross-*

Ablation	KonIQ-10k	CLIVE	FLIVE	SPAQ
D = 2	0.815	0.718	0.437	0.834
L = 3	0.808	0.719	0.382	0.788
quality-based	0.745	0.708	0.396	0.756
QualiCLIP	**0.817**	**0.725**	**0.442**	**0.841**

Table 3. Ablation study on the training strategy. Best and second-best scores are highlighted in bold and underlined, respectively.

dataset setting. Despite not requiring MOS, QualiCLIP outperforms the baselines on 11 out of 14 evaluation metrics. Specifically, our method achieves excellent performance across datasets with various types of distortions, demonstrating its robustness. This makes our model well-suited for real-world applications where a training set is unavailable. Moreover, comparing Tabs. 1 and 2 reveals two key observations. First, the performance of supervised opinion-aware models significantly decreases when tested on unseen datasets (*e.g.* GRepQ on AIGC datasets), highlighting their limited generalization capabilities. Second, QualiCLIP stands out as the only opinion-unaware approach to consistently obtain remarkable results even against supervised opinion-aware methods.

4.2. Ablation Studies

We conduct ablation studies to analyze the impact of different components of our training strategy, the importance of each loss term, and the contribution of each of the antonym prompts in the quality score computation. For simplicity, we only report the SRCC results on the authentic datasets.

Training strategy We evaluate the performance achieved by modified versions of our approach: 1) $D = 2$: we apply two sequential degradations to each crop in Sec. 3.2 instead of just one; 2) $L = 3$: we consider only three levels of degradation in Secs. 3.2 and 3.3 instead of five; 3) we compute the ranking loss using the predicted quality scores – obtained

$\mathcal{L}_{cons}$	$\mathcal{L}_{pos}$	$\mathcal{L}_{neg}$	KonIQ-10k	CLIVE	FLIVE	SPAQ
✓	✗	✗	0.733	0.668	0.368	0.703
✗	✓	✗	0.764	0.532	0.410	0.769
✗	✗	✓	0.805	0.714	0.420	0.812
✓	✓	✗	0.771	0.550	0.404	0.791
✓	✗	✓	0.800	<u>0.715</u>	0.427	0.820
✗	✓	✓	<u>0.811</u>	0.710	<u>0.434</u>	<u>0.837</u>
✓	✓	✓	**0.817**	**0.725**	**0.442**	**0.841**

Table 4. Ablation study on the loss terms. Best and second-best scores are highlighted in bold and underlined, respectively.

T_p	T_n	KonIQ-10k	CLIVE	FLIVE	SPAQ
✓	✗	0.230	0.063	0.064	0.059
✗	✓	<u>0.661</u>	<u>0.686</u>	<u>0.332</u>	<u>0.818</u>
✓	✓	**0.817**	**0.725**	**0.442**	**0.841**

Table 5. Analysis of the individual prompt contributions in the quality score computation. Best and second-best scores are highlighted in bold and underlined, respectively.

with Eq. (5) – associated to each degraded crop instead of its similarity to the antonym prompts; Tab. 3 shows the results. First, we notice that employing more than one distortion leads to slightly worse performance. We argue that this result stems from the synthetic degradation becoming too severe independently of the level of intensity, making it overly challenging for the model to rank the crops effectively. Moreover, considering only $L = 3$ levels of degradation provides less information to the model during training compared to using five different levels and thus significantly worsens the results. Then, we observe that directly employing the predicted quality scores in the ranking loss instead of the similarity to the prompts achieves poor performance. We attribute this outcome to an increased discrepancy between CLIP's training and fine-tuning process. Indeed, while the predicted quality scores originate from two prompts (see Eq. (5)), the proposed strategy considers multiple pairs of single images and texts, which we argue is more similar to the technique used for training CLIP [32].

Loss terms We study the importance of each loss term in Eq. (4) and report the results in Tab. 4. First, we notice that using only $\mathcal{L}_{cons}$ leads to a significant performance decrease, as $\mathcal{L}_{cons}$ does not exploit the information provided by the intrinsic ranking of the increasingly degraded crops. Nevertheless, $\mathcal{L}_{cons}$ consistently yields a positive impact when combined with any of the other loss terms. Then, we observe that, while $\mathcal{L}_{pos}$ and $\mathcal{L}_{neg}$ differ only for the type of prompt they consider, $\mathcal{L}_{neg}$ proves to be significantly more critical for the training process. Nevertheless, Tab. 4 shows that combining the three loss terms achieves the best results, proving that they are all crucial for training CLIP to generate accurate quality-aware image representations.

Individual prompt contributions The results of the ablation studies on the training loss terms show that $\mathcal{L}_{neg}$ is more critical than $\mathcal{L}_{pos}$ for the training process. We recall that $\mathcal{L}_{pos}$ and $\mathcal{L}_{neg}$ involve the alignment between the images and the positive and negative prompts, respectively. This suggests that the similarity between the image and the negative prompt has a greater influence than that with the positive prompt on the quality score computation (as in Eq. (5)). To support this hypothesis, we study the individual

prompt contributions in obtaining the final quality scores.

We conduct an experiment where we directly use the similarity between the image and each of the antonym prompts as the quality score. This is possible because both the similarities and the quality scores range between 0 and 1. Tab. 5 shows the results. We observe that the similarity between the negative prompt and the image provides significantly more information about its inherent quality than the positive prompt. This result supports our hypothesis and is consistent with the greater importance of $\mathcal{L}_{neg}$ in our training strategy. Nonetheless, Tab. 5 also indicates that both prompts are essential for the quality score computation, as their combination results in the best performance.

We carry out an additional experiment to determine whether the discrepancy in the contributions of the positive and negative prompts arises from our training strategy or is inherent to CLIP itself. Specifically, we follow the experimental setting described above to evaluate the individual contributions of the prompts in the quality score computation of CLIP-IQA [47]. We recall that CLIP-IQA employs an off-the-shelf CLIP model and computes the final quality scores using a strategy similar to Eq. (5). Our experiment reveals that using T_p and T_n individually results in a SRCC of 0.010 and 0.571, respectively, on the KonIQ-10k dataset. This outcome leads us to conclude that the similarity with the negative prompt inherently provides more meaningful information about image quality compared to using the positive prompt. We will investigate this finding more thoroughly in future work.

5. Conclusion

In this work, we propose QualiCLIP, a self-supervised opinion-unaware approach that enhances CLIP's ability to produce accurate quality-aware image representations. In particular, we design a quality-aware image-text alignment strategy that trains CLIP to rank increasingly synthetically degraded images based on their similarity with antonym prompts, while ensuring consistent representations for images with similar content and comparable quality. Compared to existing opinion-unaware methods, QualiCLIP shows significant performance improvements across several datasets. Moreover, it is the only opinion-unaware approach that, in most cases, outperforms opinion-aware methods in

cross-dataset experiments. Thus, we believe that QualiCLIP could serve as a strong baseline for evaluating the generalization capabilities of future NR-IQA approaches.

References

[1] Lorenzo Agnolucci, Leonardo Galteri, Marco Bertini, and Alberto Del Bimbo. ARNIQA: Learning Distortion Manifold for Image Quality Assessment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 189–198, 2024. 2, 3, 4, 6, 7, 13, 14, 16

[2] Lorenzo Agnolucci, Leonardo Galteri, Marco Bertini, and Alberto Del Bimbo. Reference-Based Restoration of Digitized Analog Videotapes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1659–1668, 2024. 3

[3] Nithin C Babu, Vignesh Kannan, and Rajiv Soundararajan. No Reference Opinion Unaware Quality Assessment of Authentically Distorted Images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2459–2468, 2023. 2, 3, 6, 7, 13

[4] Tai-Yin Chiu, Yinan Zhao, and Danna Gurari. Assessing Image Quality Issues for Real-World Problems. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3646–3656, 2020. 4

[5] Yuming Fang, Hanwei Zhu, Yan Zeng, Kede Ma, and Zhou Wang. Perceptual Quality Assessment of Smartphone Photography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3677–3686, 2020. 6, 12

[6] Deepti Ghadiyaram and Alan C Bovik. Massive Online Crowdsourced Study of Subjective and Objective Picture Quality. *IEEE Transactions on Image Processing*, 25(1): 372–387, 2015. 6, 12, 13

[7] S Alireza Golestaneh, Saba Dadsetan, and Kris M Kitani. No-Reference Image Quality Assessment via Transformers, Relative Ranking, and Self-Consistency. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1220–1230, 2022. 2, 3, 4, 7, 12, 13, 16

[8] Video Quality Experts Group. Final Report from the Video Quality Experts Group on the Validation of Objective Models of Video Quality Assessment. 2000. 6

[9] Jie Gu, Gaofeng Meng, Cheng Da, Shiming Xiang, and Chunhong Pan. No-Reference Image Quality Assessment with Reinforcement Recursive List-Wise Ranking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8336–8343, 2019. 2, 3

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 16

[11] 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*, 29:4041–4056, 2020. 1, 2, 6, 12, 15, 16

[12] Gu Jinjin, Cai Haoming, Chen Haoyu, Ye Xiaoxing, Jimmy S Ren, and Dong Chao. PIPAL: a Large-Scale Image Quality Assessment Dataset for Perceptual Image Restoration. In *European Conference on Computer Vision*, pages 633–651. Springer, 2020. 6, 12, 14

[13] Junjie Ke, Keren Ye, Jiahui Yu, Yonghui Wu, Peyman Milanfar, and Feng Yang. VILA: Learning Image Aesthetics from User Comments with Vision-Language Pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10041–10051, 2023. 3

[14] Eric C Larson and Damon M Chandler. Most Apparent Distortion: Full-Reference Image Quality Assessment and the Role of Strategy. *Journal of Electronic Imaging*, 19(1): 011006–011006, 2010. 12

[15] Chunyi Li, Zicheng Zhang, Haoning Wu, Wei Sun, Xiongkuo Min, Xiaohong Liu, Guangtao Zhai, and Weisi Lin. AGIQA-3K: an Open Database for AI-Generated Image Quality Assessment. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(8):6833–6846, 2023. 6, 12

[16] Zhexin Liang, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Iterative Prompt Learning for Unsupervised Backlit Image Enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8094–8103, 2023. 2, 3, 4

[17] Ping-Sung Liao, Tse-Sheng Chen, Pau-Choo Chung, et al. A Fast Algorithm for Multilevel Thresholding. *J. Inf. Sci. Eng.*, 17(5):713–727, 2001. 17

[18] Hanhe Lin, Vlad Hosu, and Dietmar Saupe. KADID-10k: A Large-scale Artificially Distorted IQA Database. In *2019 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–3. IEEE, 2019. 4, 12, 16

[19] Xialei Liu, Joost Van De Weijer, and Andrew D Bagdanov. RankIQA: Learning from Rankings for No-Reference Image Quality Assessment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1040–1049, 2017. 2, 3, 4

[20] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *arXiv preprint arXiv:1711.05101*, 2017. 16

[21] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling Vision-Language Models for Universal Image Restoration. *arXiv preprint arXiv:2310.01018*, 2023. 3

[22] Chao Ma, Chih-Yuan Yang, Xiaokang Yang, and Ming-Hsuan Yang. Learning a No-Reference Quality Metric for Single-Image Super-Resolution. *Computer Vision and Image Understanding*, 158:1–16, 2017. 6, 12

[23] Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang. Waterloo Exploration Database: New Challenges for Image Quality Assessment Models. *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016. 13

[24] Kede Ma, Qingbo Wu, Zhou Wang, Zhengfang Duanmu, Hongwei Yong, Hongliang Li, and Lei Zhang. Group MAD Competition - A New Methodology to Compare Objective Image Quality Models. In *Proceedings of the IEEE Con-*

ference on Computer Vision and Pattern Recognition, pages 1664–1673, 2016. 13, 14

[25] Kede Ma, Wentao Liu, Tongliang Liu, Zhou Wang, and Dacheng Tao. dipIQ: Blind Image Quality Assessment by Learning-to-Rank Discriminable Image Pairs. *IEEE Transactions on Image Processing*, 26(8):3951–3964, 2017. 2, 3

[26] Pavan C Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C Bovik. Image Quality Assessment Using Contrastive Learning. *IEEE Transactions on Image Processing*, 31:4149–4161, 2022. 2, 3, 6, 7, 13, 14, 16

[27] Marco Mistretta, Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Andrew D. Bagdanov. Cross the Gap: Exposing the Intra-modal Misalignment in CLIP via Modality Inversion. In *The Thirteenth International Conference on Learning Representations*, 2025. 6

[28] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-Reference Image Quality Assessment in the Spatial Domain. *IEEE Transactions on Image Processing*, 21 (12):4695–4708, 2012. 2

[29] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a "Completely Blind" Image Quality Analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012. 2, 3, 6, 13

[30] Zhangkai Ni, Yue Liu, Keyan Ding, Wenhan Yang, Hanli Wang, and Shiqi Wang. Opinion-Unaware Blind Image Quality Assessment using Multi-Scale Deep Feature Statistics. *IEEE Transactions on Multimedia*, 2024. 2, 6, 13

[31] Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, et al. Image Database TID2013: Peculiarities, Results and Perspectives. *Signal Processing: Image Communication*, 30:57–77, 2015. 12

[32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual Models from Natural Language Supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 3, 5, 8

[33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 17

[34] Subhadeep Roy, Shankhanil Mitra, Soma Biswas, and Rajiv Soundararajan. Test Time Adaptation for Blind Image Quality Assessment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16742–16751, 2023. 3, 4

[35] Avinab Saha, Sandeep Mishra, and Alan C Bovik. Re-IQA: Unsupervised Learning for Image Quality Assessment in the Wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5846–5855, 2023. 1, 2, 3, 4, 5, 6, 7, 13, 14, 16

[36] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017. 14

[37] Hamid R Sheikh, Muhammad F Sabir, and Alan C Bovik. A Statistical Evaluation of Recent Full Reference Image Quality Assessment Algorithms. *IEEE Transactions on Image Processing*, 15(11):3440–3451, 2006. 12

[38] Jinsong Shi, Pan Gao, and Jie Qin. Transformer-based No-Reference Image Quality Assessment via Supervised Contrastive Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4829–4837, 2024. 2

[39] Nyeong-Ho Shin, Seon-Ho Lee, and Chang-Su Kim. Blind Image Quality Assessment Based on Geometric Order Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12799–12808, 2024. 7, 13, 14, 16

[40] Ankit Shukla, Avinash Upadhyay, Swati Bhugra, and Manoj Sharma. Opinion Unaware Image Quality Assessment via Adversarial Convolutional Variational Autoencoder. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2153–2163, 2024. 2, 3

[41] Suhas Srinath, Shankhanil Mitra, Shika Rao, and Rajiv Soundararajan. Learning Generalizable Perceptual Representations for Data-Efficient No-Reference Image Quality Assessment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 22–31, 2024. 2, 3, 6, 7, 13, 14, 15, 16

[42] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly Assess Image Quality in the Wild Guided by a Self-Adaptive Hyper Network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3667–3676, 2020. 1, 2, 7, 13, 14, 16

[43] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. YFCC100M: The New Data in Multimedia Research. *Communications of the ACM*, 59(2):64–73, 2016. 12

[44] William Thong, Jose Costa Pereira, Sarah Parisot, Ales Leonardis, and Steven McDonagh. Content-Diverse Comparisons Improve IQA. *arXiv preprint arXiv:2211.05215*, 2022. 3

[45] Vishaal Udandarao, Ankush Gupta, and Samuel Albanie. SuS-X: Training-Free Name-Only Transfer of Vision-Language Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2725–2736, 2023. 6

[46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, 9(11), 2008. 15

[47] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring CLIP for Assessing the Look and Feel of Images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2555–2563, 2023. 1, 2, 3, 4, 6, 7, 8, 12, 13, 14, 15, 16, 17

[48] Haoning Wu, Liang Liao, Jingwen Hou, Chaofeng Chen, Erli Zhang, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring Opinion-Unaware Video Quality Assessment with Semantic Affinity Criterion. *arXiv preprint arXiv:2302.13269*, 2023. 4, 16

[49] Haoning Wu, Liang Liao, Annan Wang, Chaofeng Chen, Jingwen Hou, Wenxiu Sun, Qiong Yan, and Weisi Lin. Towards Robust Text-Prompted Semantic Criterion for In-the-Wild Video Quality Assessment. *arXiv preprint arXiv:2304.14672*, 2023. 4, 16

[50] Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Chunyi Li, Wenxiu Sun, Qiong Yan, Guangtao Zhai, et al. Q-Bench: A Benchmark for General-Purpose Foundation Models on Low-level Vision. In *The Twelfth International Conference on Learning Representations*, 2024. 3

[51] Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. Q-Align: Teaching LMMs for Visual Scoring via Discrete Text-Defined Levels. In *Proceedings of the 41st International Conference on Machine Learning*, pages 54015–54029. PMLR, 2024. 3

[52] Kangmin Xu, Liang Liao, Jing Xiao, Chaofeng Chen, Haoning Wu, Qiong Yan, and Weisi Lin. Boosting Image Quality Assessment through Efficient Transformer Adaptation with Local Feature Enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2662–2672, 2024. 2

[53] Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, and Alan Bovik. From Patches to Pictures (PaQ-2-PiQ): Mapping the Perceptual Space of Picture Quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3575–3585, 2020. 1, 6, 7, 12

[54] Zhiyuan You, Zheyuan Li, Jinjin Gu, Zhenfei Yin, Tianfan Xue, and Chao Dong. Depicting Beyond Scores: Advancing Image Quality Assessment Through Multi-Modal Language Models. *arXiv preprint arXiv:2312.08962*, 2023. 3

[55] Lin Zhang, Lei Zhang, and Alan C Bovik. A Feature-Enriched Completely Blind Image Quality Evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015. 2, 3, 6, 13

[56] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 17

[57] Weixia Zhang, Kede Ma, Jia Yan, Dexiang Deng, and Zhou Wang. Blind Image Quality Assessment using a Deep Bilinear Convolutional Neural Network. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(1):36–47, 2018. 7, 13, 14, 16

[58] Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind Image Quality Assessment via Vision-Language Correspondence: A Multitask Learning Perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14071–14081, 2023. 1, 2, 3, 7, 13, 14, 16

[59] Zicheng Zhang, Chunyi Li, Wei Sun, Xiaohong Liu, Xiongkuo Min, and Guangtao Zhai. A Perceptual Quality Assessment Exploration for AIGC Images. In *IEEE International Conference on Multimedia and Expo Workshops*, pages 440–445. IEEE, 2023. 6, 12

[60] Kai Zhao, Kun Yuan, Ming Sun, Mading Li, and Xing Wen. Quality-Aware Pre-Trained Models for Blind Image Quality Assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22302–22313, 2023. 1, 2, 3, 4, 5, 6

Quality-Aware Image-Text Alignment
for Opinion-Unaware Image Quality Assessment

Supplementary Material

Overview

In this document, we present additional details and extended experimental analyses to complement the main paper. The supplementary material is organized as follows:

S1. Datasets: we report the details of the datasets employed in the experiments;

S2. Additional Experimental Results:

  S2.1. Quantitative Results: we provide the zero-shot results on synthetic datasets and the cross-dataset performance using the CLIVE and PI-PAL datasets for training the baselines;

  S2.2. Ablation Studies: we analyze the impact on the performance of the backbone of CLIP's image encoder;

  S2.3. gMAD Competition: we conduct the gMAD competition against GRepQ and CLIP-IQA;

  S2.4. t-SNE Visualization: we visualize the image representations of QualiCLIP with t-SNE;

  S2.5. Supervised QualiCLIP: we extend our approach to leverage human annotations;

  S2.6. Inference Time: we evaluate the inference time of our model;

S3. Implementation Details: we provide the implementation details of the training strategy, the prompts, and the synthetic distortions.

S4. Limitations: we discuss the limitations of the proposed approach.

S1. Datasets

To carry out the experiments, we employ several types of datasets, namely authentic, image restoration, AIGC, and synthetic datasets. We rely on four authentic datasets: KonIQ-10k [11], CLIVE [6], FLIVE [53], and SPAQ [5]. KonIQ-10k contains 10K images sampled from the YFCC100M [43] database. CLIVE consists of 1162 images captured with a wide range of mobile devices. FLIVE is the largest existing dataset for NR-IQA and is composed of about 40K real-world images. SPAQ comprises 11K high-resolution photos taken with several smartphones. Following [5], we resize the SPAQ images so that the shorter side is 512 pixels. We use two image restoration datasets: CVIU [22] and PIPAL [12]. CVIU stems from 30 reference images distorted with 9 super-resolution methods, resulting in 1620 images. PIPAL comprises 23200 images degraded with 40 distortion types, including GAN-based super-resolution methods, and originates from 250 refer-ence images. We employ two AIGC datasets: AGIQA-1K [59] and AGIQA-3K [15]. AGIQA-1K includes 1080 im-ages generated with 2 diffusion models. AGIQA-3K con-sists of 2982 images generated via 6 generative models, in-cluding auto-regressive and diffusion ones. We consider four synthetic datasets: LIVE [37], CSIQ [14], TID2013 [31], and KADID-10k [18]. LIVE comprises 779 images degraded with 5 different distortion types at 5 levels of in-tensity, with 29 reference images as the base. CSIQ origi-nates from 30 reference images, each distorted with 6 dis-tinct degradations at 5 intensity levels, resulting in 866 im-ages. TID2013 and KADID-10k comprise 3000 and 10125 images degraded using 24 and 25 types of distortion across 5 different degrees of intensity, originating from 25 and 81 reference images, respectively.

S2. Additional Experimental Results

S2.1. Quantitative Results

Zero-shot setting In Tab. S1, we compare the performance of our model with existing opinion-unaware approaches on synthetic datasets. We observe that our method achieves competitive performance, obtaining the most consistent re-sults among the considered approaches. However, as also observed in Sec. 4.1, supervised opinion-aware meth-ods achieve better performance than opinion-unaware ones when a training set is available.

Cross-dataset setting Tabs. S2 and S3 present the results of the *cross-dataset* setting when employing CLIVE [6] and PIPAL [12], respectively, to train the supervised baselines. Note that we do not report the performance of TReS [7] with PIPAL as the training dataset as there is no public pre-trained model available. The results show that QualiCLIP achieves excellent performance regardless of the dataset used to train the baselines. Moreover, we observe that the opinion-aware approaches generally perform worse when trained on PIPAL compared to CLIVE. We attribute this outcome to the nature of the degradation types included in PIPAL, which are different from those in the other testing datasets. This result highlights the sensitivity of the super-vised opinion-aware methods to the training data, and fur-ther confirms their limited generalization capabilities.

S2.2. Ablation Studies

Backbone of CLIP's image encoder Following [47], we evaluate the impact of the backbone architecture of CLIP's image encoder on performance. Specifically, we examine

Method	OU	LIVE SRCC	LIVE PLCC	CSIQ SRCC	CSIQ PLCC	TID2013 SRCC	TID2013 PLCC	KADID-10k SRCC	KADID-10k PLCC
CONTRIQUE [26]	✗	0.962	0.967	0.946	0.960	0.821	0.843	0.927	0.928
Re-IQA [35]	✗	0.954	0.956	0.946	0.954	0.873	0.895	0.867	0.865
ARNIQA [1]	✗	0.965	0.970	0.961	0.969	0.874	0.896	0.905	0.908
CLIP-IQA$^+$ [47]	✗	0.950	0.946	0.862	0.882	0.804	0.834	0.818	0.822
GRepQ [41]	✗	0.926	0.931	0.844	0.871	0.668	0.681	0.808	0.797
NIQE [29]	✓	<u>0.903</u>	<u>0.906</u>	0.642	0.727	0.308	0.419	0.380	0.442
IL-NIQE [55]	✓	0.878	0.884	**0.848**	**0.883**	0.504	0.636	0.557	0.599
CL-MI [3]	✓	0.749	0.731	0.619	0.615	0.249	0.317	0.518	0.529
MDFS [30]	✓	**0.926**	**0.930**	0.797	0.829	<u>0.554</u>	<u>0.657</u>	0.608	0.634
CONTRIQUE-OU [26]	✓	0.841	0.840	0.684	0.700	0.321	0.344	0.552	0.564
Re-IQA-OU [35]	✓	0.781	0.779	0.713	0.725	0.284	0.339	0.516	0.534
ARNIQA-OU [1]	✓	0.853	0.852	<u>0.832</u>	0.810	0.467	0.527	<u>0.632</u>	<u>0.638</u>
CLIP-IQA [47]	✓	0.626	0.651	0.748	0.805	0.525	0.632	0.465	0.473
GRepQ-OU [41]	✓	0.714	0.698	0.740	0.742	0.423	0.565	0.416	0.463
QualiCLIP	✓	0.898	0.885	0.804	<u>0.842</u>	**0.651**	**0.732**	**0.654**	**0.664**

Table S1. Quantitative results of the *zero-shot* evaluation setting using synthetic datasets. OU stands for Opinion-Unaware. Best and second-best scores for OU methods are highlighted in bold and underlined, respectively. The suffix "-OU" indicates approaches modified to be opinion-unaware (see Sec. 4.1). For reference, we report the performance of supervised methods (*i.e.* OU=✗) trained on the training split of each testing dataset.

| | | Authentic Datasets | | | | | | Image Restoration Datasets | | | | AIGC Datasets | | | |
| | | KonIQ-10k | | FLIVE | | SPAQ | | CVIU | | PIPAL | | AGIQA-1K | | AGIQA-3K | |
Method	OU	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
HyperIQA [42]	✗	0.734	0.773	0.328	0.483	0.771	0.796	0.632	0.642	0.361	0.384	0.183	0.351	0.594	0.686
DBCNN [57]	✗	0.703	0.755	0.299	0.435	0.806	0.816	0.728	0.748	0.332	0.341	<u>0.445</u>	<u>0.693</u>	0.448	0.558
TReS [7]	✗	0.727	0.756	0.359	0.494	0.858	0.863	0.713	0.739	<u>0.397</u>	**0.453**	0.364	0.643	0.549	0.666
LIQE [58]	✗	<u>0.810</u>	0.787	0.277	0.271	0.830	0.840	0.704	0.717	**0.446**	<u>0.445</u>	0.417	0.609	<u>0.653</u>	<u>0.732</u>
QCN [39]	✗	0.762	<u>0.835</u>	0.428	0.537	<u>0.860</u>	<u>0.867</u>	0.765	0.770	0.393	0.413	0.337	0.677	0.584	0.729
CONTRIQUE [26]	✗	0.736	0.749	0.399	0.527	0.834	0.842	0.583	0.601	0.310	0.370	0.256	0.393	0.503	0.597
Re-IQA [35]	✗	0.600	0.618	0.252	0.341	0.820	0.829	0.576	0.685	0.360	0.386	0.330	0.548	0.431	0.516
ARNIQA [1]	✗	0.745	0.775	<u>0.434</u>	0.526	**0.861**	**0.871**	0.610	0.643	0.369	0.382	0.320	0.629	0.526	0.649
CLIP-IQA$^+$ [47]	✗	0.769	0.808	0.410	<u>0.547</u>	0.846	0.854	<u>0.837</u>	<u>0.851</u>	0.369	0.385	0.395	0.584	0.637	0.712
GRepQ [41]	✗	0.742	0.750	0.391	0.464	0.855	0.865	0.701	0.717	0.368	0.385	0.346	0.589	0.616	0.730
QualiCLIP	✓	**0.817**	**0.838**	**0.442**	**0.556**	0.841	0.851	**0.840**	**0.852**	0.412	0.422	**0.736**	**0.805**	**0.667**	**0.735**

Table S2. Quantitative results of the *cross-dataset* evaluation setting. We employ the CLIVE [6] dataset to train the supervised methods. OU stands for Opinion-Unaware. Best and second-best scores are highlighted in bold and underlined, respectively.

the ResNet50 and ViT-B/32 backbones. In Tab. S4 we compare the results of QualiCLIP with CLIP-IQA [47], which leverages an off-the-shelf CLIP model. As also observed by Wang et al. [47], relying on the ViT-B/32 significantly hinders the performance over the ResNet50. This outcome stems from the stronger inductive bias of convolutional networks compared to transformers, which are more sensitive to the removal of the positional embedding. Nevertheless, for both backbones, we observe that QualiCLIP outperforms CLIP-IQA.

S2.3. gMAD Competition

To assess the robustness of our model, we carry out the group maximum differentiation (gMAD) competition [24]. In particular, we compare QualiCLIP against GRepQ and CLIP-IQA using the Waterloo Exploration Database [23] dataset, which comprises 95K synthetically degraded images without MOS annotations. In this evaluation, one model is fixed to function as a defender, and its quality predictions are grouped into two distinct levels. The other model assumes the role of the attacker, tasked with identifying image pairs within each level that exhibit the greatest

| | | Authentic Datasets | | | | | | | | IR Dataset | | AIGC Datasets | | | |
| | | KonIQ-10k | | CLIVE | | FLIVE | | SPAQ | | CVIU | | AGIQA-1K | | AGIQA-3K | |
Method	OU	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
HyperIQA [42]	✗	0.570	0.564	0.560	0.642	0.246	0.310	0.752	0.762	<u>0.821</u>	<u>0.835</u>	<u>0.493</u>	0.486	0.497	0.535
DBCNN [57]	✗	<u>0.718</u>	<u>0.744</u>	0.570	0.638	0.267	0.378	0.779	0.795	0.803	0.814	0.427	0.651	0.516	0.601
LIQE [58]	✗	0.681	0.696	<u>0.657</u>	<u>0.691</u>	0.138	0.178	0.692	0.690	0.797	0.810	0.413	0.465	0.560	0.621
QCN [39]	✗	0.584	0.663	0.463	0.559	<u>0.292</u>	<u>0.408</u>	<u>0.785</u>	<u>0.797</u>	0.766	0.767	0.484	<u>0.679</u>	0.563	0.650
CONTRIQUE [26]	✗	0.445	0.460	0.342	0.408	0.228	0.308	0.596	0.600	0.719	0.723	0.311	0.535	0.344	0.392
Re-IQA [35]	✗	0.351	0.376	0.297	0.358	0.206	0.273	0.633	0.641	0.670	0.692	0.291	0.410	0.317	0.375
ARNIQA [1]	✗	0.246	0.325	0.389	0.486	0.224	0.252	0.614	0.644	0.757	0.768	0.173	0.564	0.557	0.629
CLIP-IQA⁺ [47]	✗	0.568	0.586	0.551	0.597	0.231	0.313	0.758	0.766	0.748	0.764	0.442	0.560	**0.679**	<u>0.720</u>
GRepQ [41]	✗	0.679	0.682	0.536	0.583	0.285	0.345	0.754	0.765	0.792	0.820	0.174	0.591	0.573	0.679
QualiCLIP	✓	**0.817**	**0.838**	**0.725**	**0.802**	**0.442**	**0.556**	**0.841**	**0.851**	**0.840**	**0.852**	**0.736**	**0.805**	<u>0.667</u>	**0.735**

Table S3. Quantitative results of the *cross-dataset* evaluation setting. We employ the PIPAL [12] dataset to train the supervised methods. OU stands for Opinion-Unaware. IR signifies Image Restoration. Best and second-best scores are highlighted in bold and underlined, respectively.

	Method	KonIQ-10k	CLIVE	FLIVE	SPAQ
RN50	CLIP-IQA	0.695	0.656	0.344	0.733
	QualiCLIP	**0.817**	**0.725**	**0.442**	**0.841**
B/32	CLIP-IQA	0.390	0.328	0.211	0.608
	QualiCLIP	**0.626**	**0.334**	**0.217**	**0.782**

Table S4. Ablation study on the CLIP backbone. RN50 and B/32 stand for ResNet50 and ViT-B/32, respectively. Best SRCC scores for each backbone are highlighted in bold.

quality difference. For a model to demonstrate robustness, the selected image pairs should show comparable quality when acting as the defender while exhibiting a notable quality disparity when assuming the role of the attacker. We observe that when we fix QualiCLIP at a low-quality level (Fig. S1a top), GRepQ fails to find picture pairs with an obvious quality difference. When considering a high-quality level (Fig. S1a bottom), the image pair identified by GRepQ shows a slight quality gap. However, when assuming the role of the attacker (Fig. S1b), QualiCLIP successfully exposes the failures of GRepQ, as it pinpoints image pairs displaying a significant quality disparity. Fig. S2 shows that the same considerations can be drawn when analyzing the results of the gMAD competition between QualiCLIP and CLIP-IQA. Hence, our approach demonstrates greater robustness than GRepQ and CLIP-IQA.

S2.4. gradCAM visualization

We evaluate the explainability of our model and CLIP-IQA via a gradCAM [36] visualization. gradCAM is a visualization technique aimed at understanding which regions of an input image are most influential for a model's decision by studying the gradients of a given layer. We employ gradCAM to produce a heatmap of the regions of the image that activate the most for each of the antonym prompts. We em-
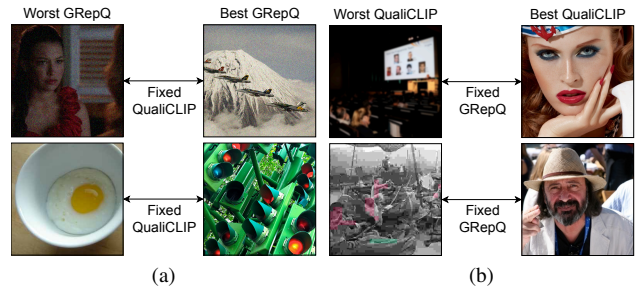

Figure S1. gMAD competition results between QualiCLIP and GRepQ [41]. (a): Fixed QualiCLIP at a low- (top) and high-quality (bottom) level, respectively. (b): Fixed GRepQ at a low- (top) and high-quality (bottom) level, respectively.
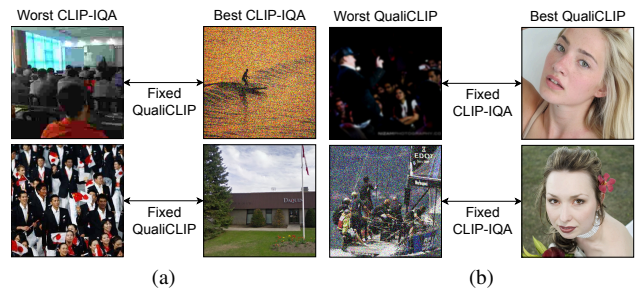

Figure S2. gMAD [24] competition results between QualiCLIP and CLIP-IQA [47]. (a): Fixed QualiCLIP at a low- (top) and high-quality (bottom) level, respectively. (b): Fixed CLIP-IQA at a low- (top) and high-quality (bottom) level, respectively.

ploy "*Good photo*" and "*Bad photo*" as the positive and negative prompts, respectively. Following [36], we consider the last convolutional layer of the ResNet50 backbone. Fig. S3a shows the result for the positive prompt. We observe that, compared to CLIP-IQA, our model leads to a better alignment with high-quality areas of the image, such as the head of the horse. Similarly, Fig. S3b illustrates that QualiCLIP
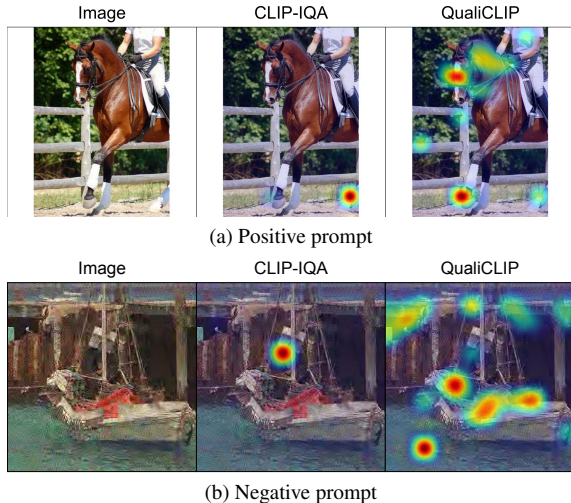
(a) Positive prompt

(b) Negative prompt

Figure S3. gradCAM visualization of the most important regions of the input image for each of the antonym prompts.

focuses on the most degraded parts of the images when considering the negative prompt, in contrast with CLIP-IQA. The improved alignment between the antonym prompts and the corresponding regions of the images makes QualiCLIP more easily explainable than CLIP-IQA.

S2.5. t-SNE Visualization

We compare the image representations generated by QualiCLIP and CLIP-IQA via a t-SNE [46] visualization. Following [41], we consider images from the CLIVE dataset with very high or very low quality. In particular, we take into account images with a labeled MOS greater than 75 and lower than 25, respectively. Fig. S4 shows the results. We observe that the representations of high- and low-quality images obtained by the proposed approach (Fig. S4b) correspond to more easily separable clusters compared to those of CLIP-IQA (Fig. S4a), which are more intertwined. This result confirms that QualiCLIP generates more accurate quality-aware representations.

S2.6. Supervised QualiCLIP

Although our approach is designed to remove the requirement for human annotations, it can easily be extended to leverage ground-truth labels. Similar to CLIP-IQA$^+$ [47], we exploit the MOS to fine-tune the antonym prompts with an MSE loss via standard backpropagation, while keeping the network weights fixed. We refer to this supervised opinion-aware variant of our approach as QualiCLIP$^+$. We train QualiCLIP$^+$ on the training split of each of the testing datasets and report the performance in Tab. S5. For a fair comparison, we only consider supervised baselines that do not fine-tune the network weights but instead train a smaller set of parameters, such as a linear regressor or the antonym prompts. We observe that QualiCLIP$^+$ achieves competi-
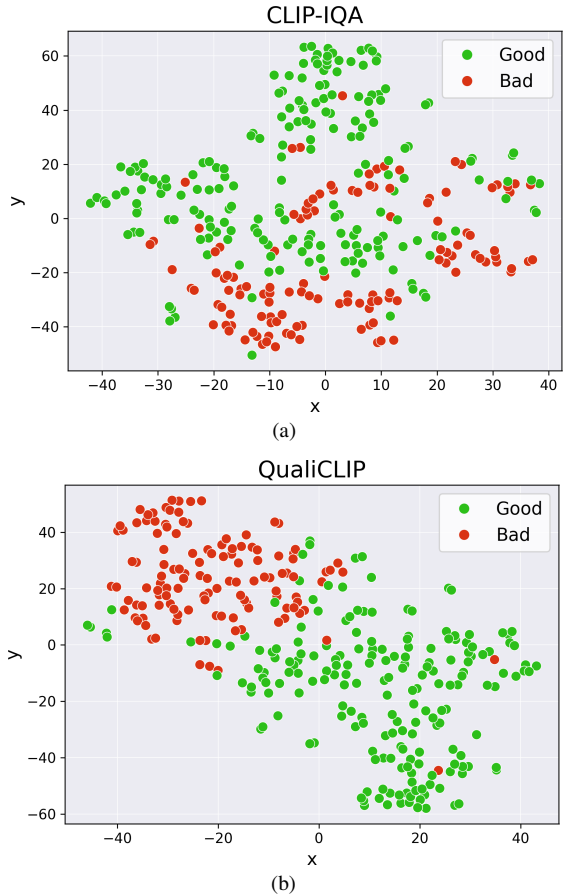

(a)

(b)

Figure S4. Comparison of the t-SNE visualizations related to the image representations of the CLIVE dataset generated by CLIP-IQA (a) and QualiCLIP (b), respectively. Good and bad points refer to images with a MOS greater than 75 and lower than 25, respectively.

tive performance also in this evaluation setting, outperforming the baselines on most metrics. This outcome shows that the proposed method can also be applied in scenarios where human annotations are available.

S2.7. Inference Time

As detailed in Sec. 3.3, we do not finetune CLIP's text encoder. Consequently, we need to compute the text features of the antonym prompts only once, and we can employ them both for training and inference. Therefore, at inference time, QualiCLIP has the same computational cost as an image-encoder-only model with a ResNet50 backbone.

To validate this, we compare the average inference time of our model with that of the supervised baselines on the KonIQ-10k dataset [11], which comprises 10K 1024 × 768px images. We conduct the experiments on an NVIDIA RTX 2080Ti GPU. We report the results in Tab. S6. As expected, QualiCLIP has a similar inference time to models based solely on ResNet50, such as CLIP-IQA$^+$, LIQE,

Method	OU	Authentic Datasets								Image Restoration Datasets				AIGC Datasets			
		KonIQ-10k		CLIVE		FLIVE		SPAQ		CVIU		PIPAL		AGIQA-1K		AGIQA-3K	
		SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
CONTRIQUE [26]	✗	0.874	0.882	0.806	0.838	0.596	0.648	0.910	0.912	0.912	0.916	0.581	<u>0.611</u>	0.799	<u>0.855</u>	0.817	0.879
Re-IQA [35]	✗	<u>0.883</u>	0.887	0.783	0.797	**0.623**	**0.685**	<u>0.909</u>	<u>0.912</u>	0.887	0.895	0.568	0.587	0.783	0.840	0.811	0.874
ARNIQA [1]	✗	0.869	0.883	0.797	0.823	0.595	<u>0.670</u>	0.904	0.909	0.919	<u>0.926</u>	**0.634**	**0.666**	0.768	0.849	0.803	0.881
CLIP-IQA$^+$ [47]	✗	0.873	<u>0.890</u>	<u>0.815</u>	<u>0.849</u>	0.602	0.662	0.901	0.904	<u>0.925</u>	0.918	0.552	0.558	<u>0.817</u>	0.854	<u>0.844</u>	<u>0.894</u>
GRepQ [41]	✗	0.882	0.883	0.793	0.813	0.576	0.611	0.902	0.903	0.882	0.857	0.554	0.568	0.740	0.797	0.807	0.858
QualiCLIP$^+$	✗	**0.889**	**0.898**	**0.821**	**0.867**	<u>0.618</u>	0.665	**0.911**	**0.914**	**0.936**	**0.936**	<u>0.586</u>	0.590	**0.826**	**0.871**	**0.860**	**0.903**

Table S5. Comparison between QualiCLIP$^+$ and supervised approaches. OU stands for Opinion-Unaware. Best and second-best scores methods are highlighted in bold and underlined, respectively.

Method	Inference Time (ms)
HyperIQA [42]	16 ± 3
DBCNN [57]	25 ± 3
TReS [7]	103 ± 5
LIQE [58]	13 ± 2
QCN [39]	101 ± 6
CONTRIQUE [26]	15 ± 3
Re-IQA [35]	42 ± 4
ARNIQA [1]	15 ± 3
CLIP-IQA+ [47]	13 ± 2
GRepQ [41]	27 ± 3
QualiCLIP	13 ± 2

Table S6. Comparison of the average inference time between QualiCLIP and supervised methods on the KonIQ-10k [11] dataset.

CONTRIQUE, and ARNIQA. Note that CONTRIQUE and ARNIQA employ the full- and half-scale versions of the input image to compute the final quality score, and thus are more computationally expensive. Moreover, our model is faster than methods based on two encoders, namely GRepQ and Re-IQA. This outcome demonstrates the efficiency and applicability of QualiCLIP in real-world scenarios.

S3. Implementation Details

Training details We rely on a ResNet50 [10] as the backbone for CLIP's image encoder. Similar to [47], we remove the positional embedding from the encoder to make our model capable of taking images of any resolution as input. The dimension d of CLIP's embedding space is 1024. Differently from [41], we do not train a projector head on top of CLIP's image encoder. We keep CLIP's text encoder frozen. We train our model for 10 epochs. We employ an AdamW [20] optimizer with a weight decay and a learning rate of $1e-2$ and $1e-9$, respectively. During training, we use a patch size of 224 and a batch size of 16. We set the margins m_{cons} in Eq. 1 and m_{rank} in Eq. 2 and 3 to $2.5e-3$

and $6.75e-2$, respectively. The loss weights λ_{cons}, λ_{pos} and λ_{neg} in Eq. 4 are all equal to 1. We set the temperature hyperparameter τ in Eq. 5 to 2. At inference time, our model takes the whole image as input.

Prompts Following [48, 49], we employ multiple pairs of antonym prompts during training and inference. In particular, we use: 1) "Good/Bad photo"; 2) "Good/Bad picture"; 3) "High-resolution/Low-resolution image"; 4) "High-quality/Low-quality image"; 5) "Sharp/Blurry image"; 6) "Sharp/Blurry edges"; 7) "Noise-free/Noisy image". We average the similarities between the images and the pairs of prompts.

Synthetic distortions As detailed in Sec. 3.2, during training we synthetically degrade pristine images with increasing intensity levels. Specifically, similar to [1], we consider 24 distinct distortion types divided into the 7 degradation groups defined by the KADID-10k [18] dataset. Each degradation has 5 degrees of progressively higher intensity. We report an example for all the intensity levels of each type of degradation in Figs. S5 to S11. Each distortion is described as follows:

1. Brightness change:
   - *Brighten*: applies a sequence of color space transformations, curve adjustments, and blending operations to increase the brightness of the image;
   - *Darken*: similar to the brighten operation, but reduces the brightness instead of increasing it;
   - *Mean shift*: adjusts the average intensity of image pixels by adding a constant value to all pixel values. Then, it constrains the resulting values to stay within the original image range;
2. Blur:
   - *Gaussian blur*: applies a Gaussian kernel filter to each image pixel;
   - *Lens blur*: applies a circular kernel filter to each image pixel;
   - *Motion blur*: applies a linear motion blur kernel to each image pixel, simulating the effect of either a moving camera or a moving object in the scene. This results in the image appearing blurred in the direction of the motion;

3. Spatial distortions:
   - *Jitter*: randomly displaces image data by applying small offsets to warp each pixel;
   - *Non-eccentricity patch*: randomly selects patches from the image and places them in random neighboring positions;
   - *Pixelate*: employs a combination of downscaling and upscaling operations using nearest-neighbor interpolation;
   - *Quantization*: quantizes the image into N uniform levels. The quantization thresholds are dynamically computed using Multi-Otsu's method [17];
   - *Color block*: randomly superimposes uniformly colored square patches onto the image;
4. Noise:
   - *White noise*: adds Gaussian white noise to the image;
   - *White noise in color component*: transforms the image to the YCbCr color space and then adds Gaussian white noise to each channel;
   - *Impulse noise*: adds salt and pepper noise to the image;
   - *Multiplicative noise*: adds speckle noise to the image;
5. Color distortions:
   - *Color diffusion*: transforms the image to the LAB color space and then applies Gaussian blur to each channel;
   - *Color shift*: randomly shifts the green channel and then blends it into the original image, masking it with the normalized gradient magnitude of the original image;
   - *Color saturation 1*: transforms the image to the HSV color space and then scales the saturation channel by a factor;
   - *Color saturation 2*: transforms the image to the LAB color space and then scales each color channel by a factor;
6. Compression:
   - *JPEG2000*: applies the standard JPEG2000 compression to the image;
   - *JPEG*: applies the standard JPEG compression to the image;
7. Sharpness & contrast:
   - *High sharpen*: applies unsharp masking to sharpen the image in the LAB color space;
   - *Nonlinear contrast change*: applies a nonlinear tone mapping operation to adjust the contrast of the image;
   - *Linear contrast change*: applies a linear tone mapping operation to adjust the contrast of the image;

S4. Limitations

The proposed approach leverages the intrinsic quality ranking of progressively degraded images as supervision to train a model in a self-supervised way. This involves defining a method to synthetically generate an inherent ranking from unlabeled images, which, in our work, is represented by the application of synthetic degradations. While we find this strategy beneficial for assessing technical image quality, it is not directly applicable to abstract perception (*e.g.* happiness or naturalness) [47] or aesthetic quality assessment. Indeed, generating an inherent image ranking for such abstract or aesthetic quantities without employing annotations requires more sophisticated strategies compared to the straightforward yet effective application of low-level degradations. Future work could focus on developing such strategies, for example by exploiting the capabilities of text-to-image generation models [33, 56] to directly synthesize intrinsically ranked images.
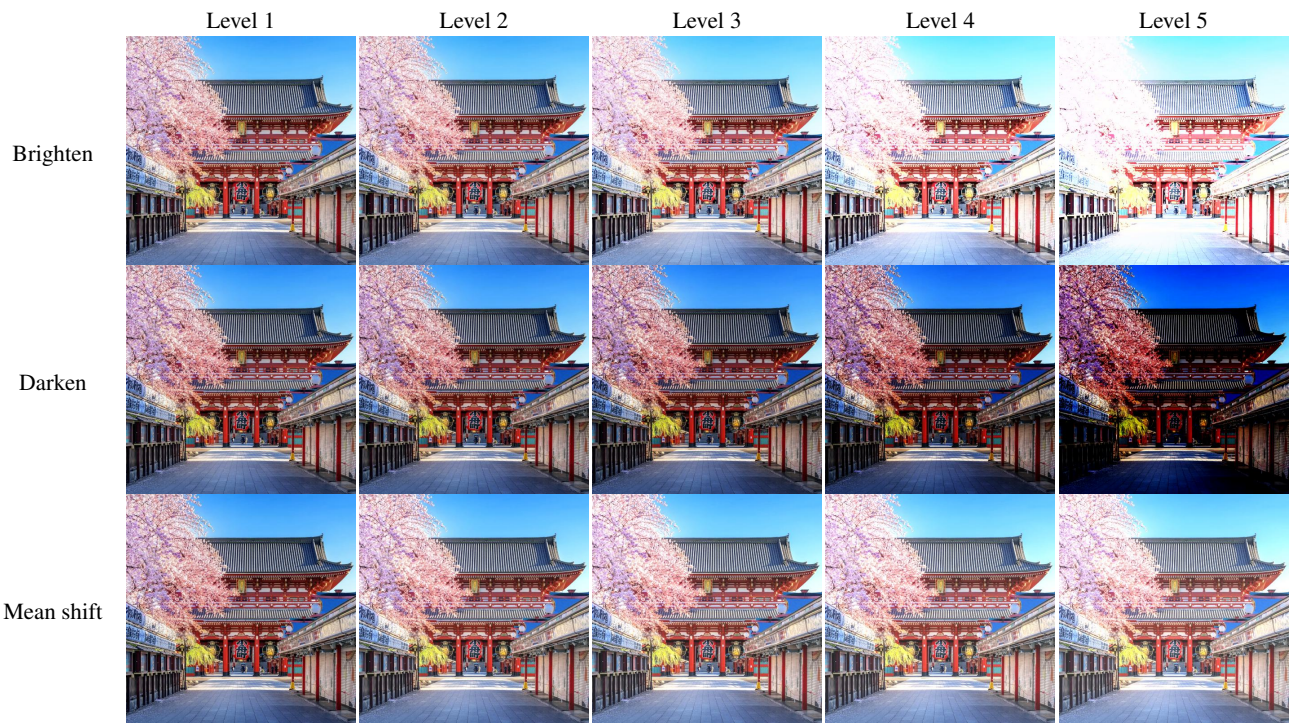
Figure S5. Visualization of the distortion types belonging to the *Brightness change* group for increasing intensity levels.
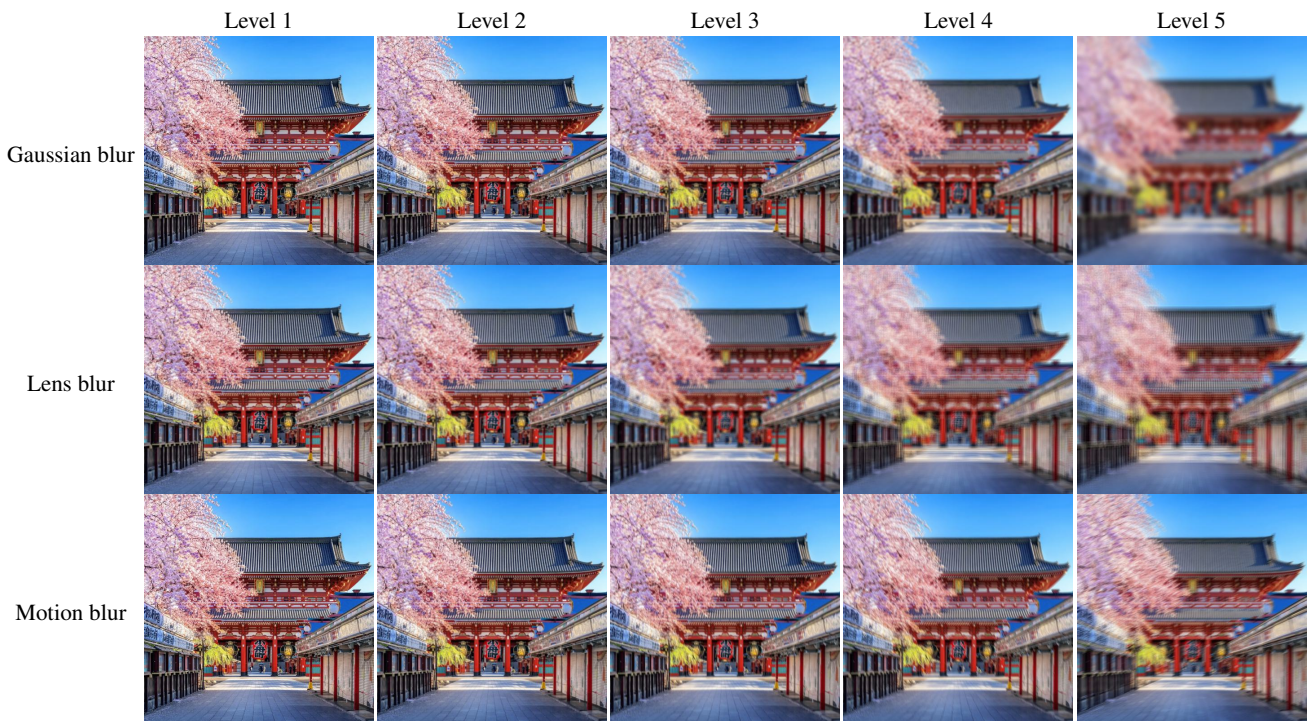

Figure S6. Visualization of the distortion types belonging to the *Blur* group for increasing intensity levels.
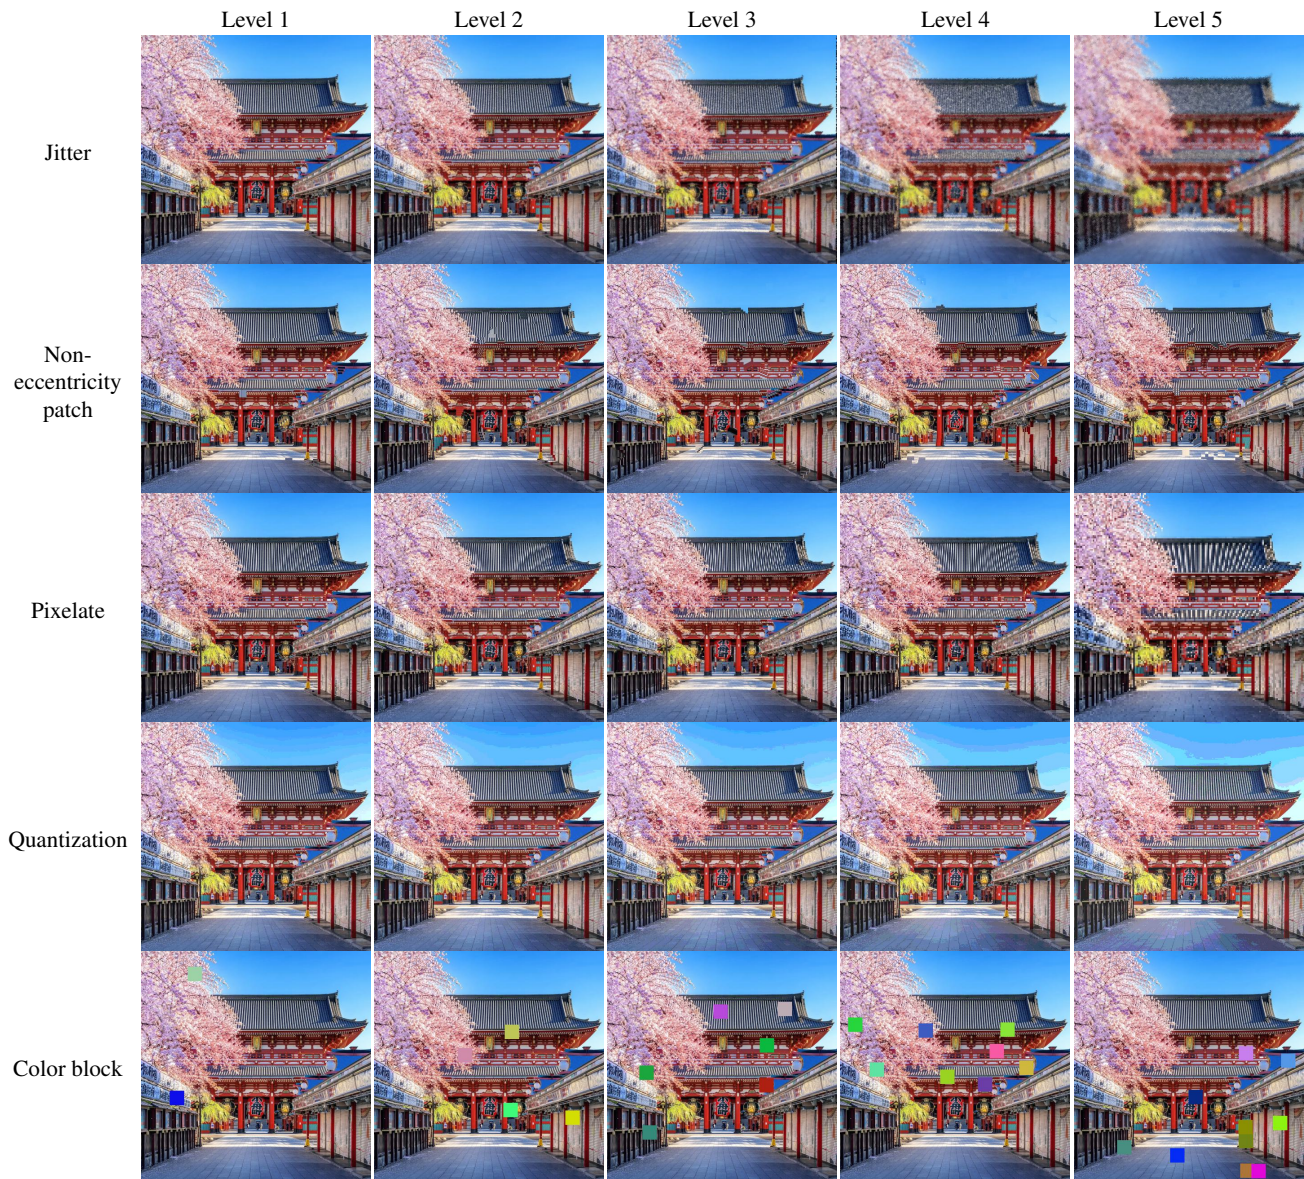
Figure S7. Visualization of the distortion types belonging to the *Spatial distortions* group for increasing intensity levels.
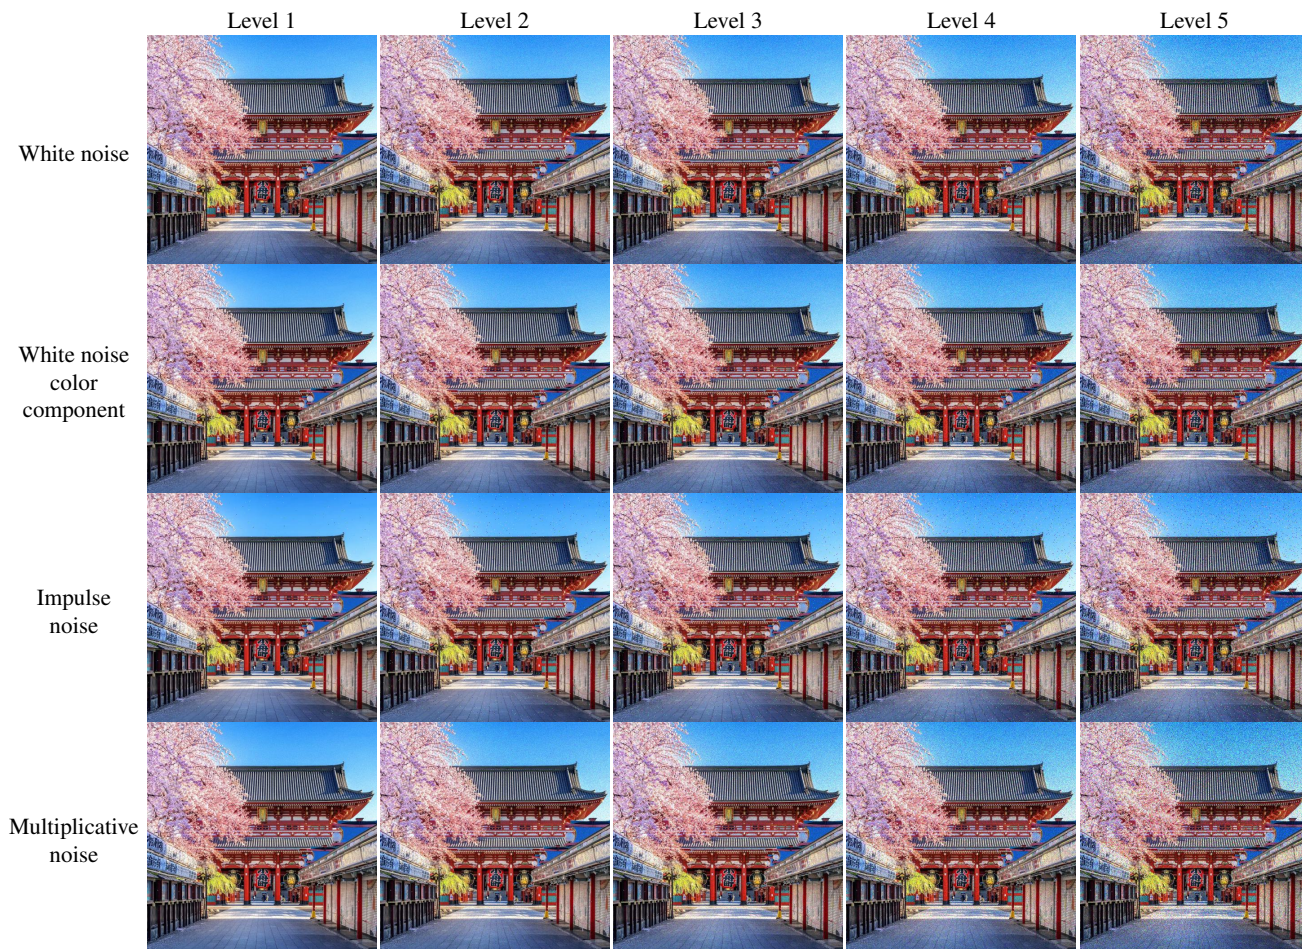
Figure S8. Visualization of the distortion types belonging to the *Noise* group for increasing intensity levels.
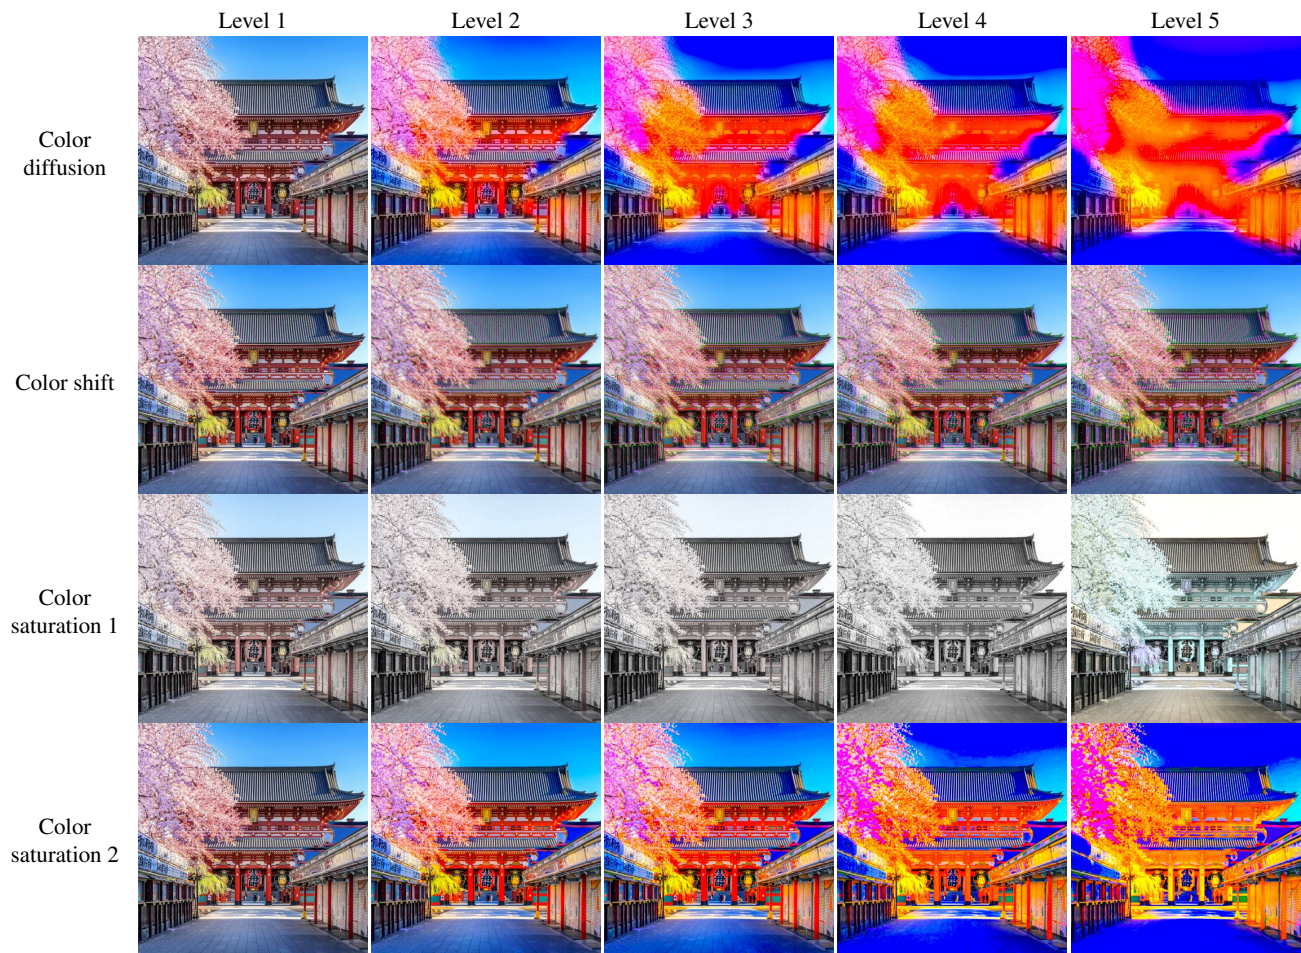
Figure S9. Visualization of the distortion types belonging to the *Color distortions* group for increasing intensity levels.

Figure S10. Visualization of the distortion types belonging to the *Compression* group for increasing intensity levels.
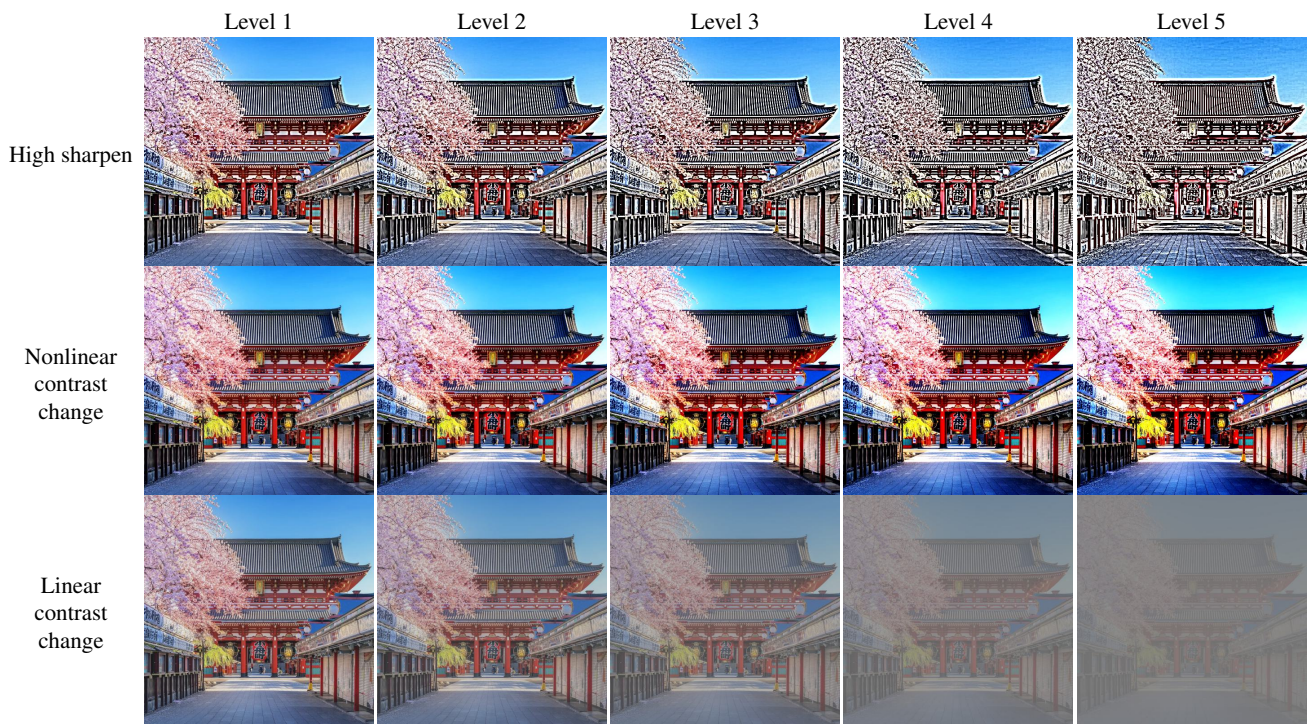
Figure S11. Visualization of the distortion types belonging to the *Sharpness & contrast* group for increasing intensity levels.