

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/187957/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Chen, Baoliang, Huang, Danni, Zhu, Hanwei, Zhu, Lingyu, Zhou, Wei, Wang, Shiqi, Fang, Yuming and Lin, Weisi 2026. From global to granular: revealing IQA model performance via correlation surface. IEEE Transactions on Pattern Analysis and Machine Intelligence 10.1109/tpami.2026.3705184

Publishers page: <https://doi.org/10.1109/tpami.2026.3705184>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



From Global to Granular: Revealing IQA Model Performance via Correlation Surface

Baoliang Chen, Danni Huang, Hanwei Zhu, Lingyu Zhu, Wei Zhou, *Senior Member, IEEE* Shiqi Wang, *Senior Member, IEEE* Yuming Fang, *Fellow, IEEE* and Weisi Lin, *Fellow, IEEE*

Abstract—Evaluation of Image Quality Assessment (IQA) models has long been dominated by global correlation metrics, such as Pearson Linear Correlation Coefficient (PLCC) and Spearman Rank-Order Correlation Coefficient (SRCC). While widely adopted, these metrics reduce performance to a single scalar, failing to capture how ranking consistency varies across the local quality spectrum. For example, two IQA models may achieve identical SRCC values, yet one ranks high-quality images (related to high Mean Opinion Score, MOS) more reliably, while the other better discriminates image pairs with small quality/MOS differences (related to $|\Delta\text{MOS}|$). Such complementary behaviors are invisible under global metrics. Moreover, SRCC and PLCC are sensitive to test-sample quality distributions, yielding unstable comparisons across test sets. To address these limitations, we propose **Granularity-Modulated Correlation (GMC)**, which provides a structured, fine-grained analysis of IQA performance. GMC includes: (1) a **Granularity Modulator** that applies Gaussian-weighted correlations conditioned on absolute MOS values and pairwise MOS differences ($|\Delta\text{MOS}|$) to examine local performance variations, and (2) a **Distribution Regulator** that regularizes correlations to mitigate biases from non-uniform quality distributions. The resulting **correlation surface** maps correlation values as a joint function of MOS and $|\Delta\text{MOS}|$, providing a 3D representation of IQA performance. Experiments on standard benchmarks show that GMC reveals performance characteristics invisible to scalar metrics, offering a more informative and reliable paradigm for analyzing, comparing, and deploying IQA models. Codes are available at <https://github.com/Dniaaa/GMC>.

Index Terms—Image quality assessment, rank correlation indicator, fine-grained evaluation, distribution robustness.

1 INTRODUCTION

DRIVEN by the proliferation of visual media and advances in artificial intelligence, image data has witnessed substantial growth in both volume and complexity. This evolution, in turn, has led to a growing demand for accurate image quality assessment (IQA), which plays a fundamental role in image compression, enhancement, restoration, and generation [1]–[3]. Over the past decades, a variety of objective Full-Reference (FR) and No-Reference (NR) IQA models have been proposed, ranging from perceptual fidelity [4]–[7] to data-driven paradigms [8]–[18]. However, as IQA models become

increasingly sophisticated, a critical question arises:

How should we comprehensively and fairly measure the performance of these IQA models?

Historically, the IQA community has relied on global correlation metrics, such as the Pearson Linear Correlation Coefficient (PLCC), Spearman Rank-Order Correlation Coefficient (SRCC), and Kendall Rank-Order Correlation Coefficient (KRCC), to quantify the alignment between model predictions and Mean Opinion Scores (MOS). However, with the increasing diversification of IQA applications, these global indicators are becoming inadequate for capturing the nuanced demands of real-world scenarios. For instance, in **generative AI**, where the majority of outputs possess high perceptual quality, there is a specific need for models that are highly accurate in the high-MOS regime. Conversely, in **extreme image compression** settings, models are expected to perform fine-grained discrimination among images with marginal differences in perceived quality (small $|\Delta\text{MOS}|$), which is crucial for effective rate-distortion optimization and requires strong fine-grained differentiation capability [19]. In practice, two IQA models might exhibit nearly identical global SRCC scores on a standard benchmark, yet their performance could diverge significantly in these specialized application contexts.

As illustrated in Fig. 1(a), we partition the SPAQ dataset [20] into subsets based on absolute quality score (i.e., MOS) and pair-wise quality difference (i.e., $|\Delta\text{MOS}|$), respectively. Our analysis reveals that CLIP-IQA [21] exhibits superior PLCC performance in high-MOS regions and

- Baoliang Chen is with the School of Computer Science, South China Normal University, and the College of Computing and Data Science, Nanyang Technological University. E-mail: blchen6-c@my.cityu.edu.hk.
- Danni Huang is with the School of Computer Science, South China Normal University. E-mail: dannnyhuang@m.scnu.edu.cn.
- Lingyu Zhu and Shiqi Wang are with the School of Computer Science, City University of Hong Kong. (E-mails: lingyuzhu-c@my.cityu.edu.hk, shiqiwang@cityu.edu.hk.)
- Wei Zhou is with the School of Computer Science and Informatics, Cardiff University, Cardiff, CF24 4AG, United Kingdom. (E-mail: zhouw26@cardiff.ac.uk);
- Yuming Fang is with the School of Computing and Artificial Intelligence, Jiangxi University of Finance and Economics, Nanchang, China, and also with Jiangxi Provincial Key Laboratory of Multimedia Intelligent Processing. (E-mail: fa0001ng@e.ntu.edu.sg).
- Hanwei Zhu, and W. Lin are with the College of Computing and Data Science, Nanyang Technological University. (E-mails: hanwei.zhu@ntu.edu.sg, wslin@ntu.edu.sg.)
- This work was supported in part by the National Natural Science Foundation of China under Grants 62401214, U24A20220, and 62132006, and in part by the Ministry of Education, Singapore, under Tier 1 Grant RG103/24. Corresponding author: Hanwei Zhu.




Fig. 1. **Complementary IQA behaviors along two coupled assessment dimensions on the SPAQ dataset [20].** The overall performance of an IQA model is governed by its *prediction accuracy* relative to absolute quality (MOS) and its *discrimination capability* regarding pairwise differences ($|\Delta\text{MOS}|$). (a) Performance snapshots: Images from the SPAQ dataset are partitioned into four subsets based on the medians of MOS and $|\Delta\text{MOS}|$. PLCC scores on these quadrants reveal the distinct regimes where models excel. For instance, CLIP-IQA [21] is proficient in high-quality and fine-grained discrimination, while NIQE [22] is more robust for severe degradations. (b) Correlation surfaces estimated by the proposed GMC provide a unified and comprehensive view, capturing how correlation performance (e.g., PLCC) evolves seamlessly across the joint MOS and $|\Delta\text{MOS}|$ space.




Fig. 2. **Sensitivity of global correlation metrics (PLCC and SRCC) to shifts in quality score distributions.** Two subsets are sampled from the PIPAL dataset with identical sample sizes but different MOS distributions. The first column shows the MOS distributions of the sampled subsets, while the second column reports the PLCC, SRCC, and the proposed global GMC score (GMC_g) for PSNR and SSIM. Although PSNR and SSIM exhibit reversed performance rankings under PLCC and SRCC under distributional shifts, GMC_g remains stable and preserves consistent model ordering.

for pairs with small $|\Delta\text{MOS}|$, suggesting its proficiency in the fine-grained discrimination of high-quality content. In contrast, NIQE [22] demonstrates higher correlation in low-MOS regimes and for pairs with large $|\Delta\text{MOS}|$, indicating its strength in identifying severe degradations. These

findings underscore that IQA model performance is jointly governed by absolute quality levels and the granularity of quality differences. Such complementary behaviors are obscured by traditional global metrics, which aggregate localized performance into a single scalar and thus provide an incomplete characterization of model performance. Through this empirical analysis, we identify two fundamental limitations of current benchmarking protocols:

- 1) **Dimension Confounding:** Global metrics fail to decouple a model's *Prediction Accuracy* (consistency with absolute MOS) from its *Discrimination Capability* (sensitivity to relative $|\Delta\text{MOS}|$). These two dimensions capture perceptually distinct aspects of human visual assessment, while being mathematically coupled in traditional global metrics.
- 2) **Distributional Bias:** As shown in Fig. 2, global metrics are highly sensitive to the MOS distribution of the test set. Changes in quality span or sampling density can reverse the relative ranking of models, indicating that these metrics confound model capability with dataset-specific distributional bias.

Inspired by these observations, we propose **Granularity-Modulated Correlation (GMC)**, a principled framework for IQA model evaluation. GMC explicitly modulates correlation computation along both MOS and $|\Delta\text{MOS}|$ dimensions, enabling localized behavioral analysis while reducing bias induced by non-uniform quality distributions. Specifically, GMC extends the classical Generalized Correlation Coefficient (GCC) [23] via two complementary mechanisms:

- **Granularity Modulator:** Unlike traditional metrics that assign a uniform weight to every image pair regardless of their absolute quality level (MOS) or pairwise quality difference ($|\Delta\text{MOS}|$), our Granularity Modulator employs explicit Gaussian-based weighting w_k^{ij} conditioned on target query coordinates (Q_k^s, Q_k^d), enabling localized and controllable assessment of model behavior across specific perceptual regimes.
- **Distribution Regulator:** To address the instability shown in Fig. 2, this regulator employs a kernel-smoothed density estimator to suppress the dominance of overrepresented quality regions while compensating for sparsely sampled intervals. This ensures that the evaluation remains fair and stable, effectively decoupling model assessment from specific quality distributions.

Beyond providing localized snapshots, our framework synthesizes these measurements into a **unified 3D correlation surface** (see Fig. 1(b)). The surface fitting offers several distinct advantages. First, it provides **enhanced interpretability** by pinpointing exactly where a model fails or excels, allowing researchers to diagnose whether a low global score stems from poor absolute accuracy or a lack of discriminative sensitivity. Second, the localized view offers **task-specific insights**. For example, a model’s high performance in the “high-quality” and “small difference” regime (as seen in CLIP-IQA) directly validates its suitability for generative AI evaluation, whereas global metrics would obscure this specialized capability. Finally, by mapping performance onto a continuous landscape, we estimate the evolution of model behavior while alleviating the severe boundary effects caused by discrete binning. It should be noted that the fidelity of this continuous representation depends on the underlying data density, since the surface is approximated through interpolation or extrapolation in sparsely sampled regions.

More importantly, this 3D representation enables a **robust global synthesis** through surface integration. By deriving a global indicator (denoted as GMC_g) from the entire volume of the surface, the resulting score is inherently distribution-agnostic, yielding a stable and “true” reflection of model capability across diverse application scenarios.

Extensive experiments on standard IQA datasets demonstrate that GMC enables more sophisticated performance analysis. The complementary strengths of different IQA models are systematically revealed, offering reliable evidence for model selection and integration. Moreover, GMC consistently outperforms traditional global metrics in both precision and reliability under varying quality distributions, highlighting its strong potential as a standardized evaluation measure.

2 BACKGROUND

2.1 Image Quality Assessment Models

Objective IQA has evolved from early error-visibility measures to modern data-driven paradigms, paralleling our deepening understanding of the human visual system (HVS). Based on the availability of a reference image, these models are generally categorized into Full-Reference

(FR-IQA) and No-Reference (NR-IQA) frameworks: (1) **FR-IQA**. Early efforts focused on pixel-level fidelity, such as MSE and PSNR, which often fail to align with human perception due to oversimplified assumptions [4], [6], [7]. A major paradigm shift occurred with the structural similarity (SSIM) index [24], which prioritized structural fidelity over error visibility, leading to extensions like MS-SSIM [25] and IW-SSIM [26]. Other classical approaches integrated information theory (VIF [27]) or gradient features (FSIM [28], GMSD [29]). Recently, perceptual comparisons have migrated to deep feature spaces, leveraging pre-trained networks (LPIPS [30], DISTS [5]) and advanced statistical distances (DeepWSD [31], DeepDC [32], DMM [33]) to capture complex HVS nonlinearities. (2) **NR-IQA**. NR-IQA models aim to assess quality without a reference, initially relying on Natural Scene Statistics (NSS) across various domains, including spatial (NIQE [22]), wavelet [34], and DCT [35]. Alternative hand-crafted methods explored the free-energy principle to estimate perceptual uncertainty [36]–[38]. The field has since been dominated by deep learning architectures, ranging from early CNNs (CN-NIQA [39], DBCNN [8]) to advanced Transformer-based and meta-learning models (HyperIQA [40], MetaIQA [41], MUSIQ [14], LIQE [15], TOPIQ [11]). To enhance generalization, strategies like continual and transfer learning have been widely adopted [42], [43]. Most recently, Large Multimodal Models (LMMs) have emerged as powerful zero-shot evaluators. Following foundational benchmarks like Q-Bench [17], specialized LMM-based models such as Q-Align [44], Dog-IQA [45], Compare2Score [46], Q-Debias [47], and VisualQuality-R1 [18] have set new frontiers by aligning large-scale generative knowledge with human quality perception. Despite these advancements, the reliance on monolithic global metrics persists, leaving the local and fine-grained behavioral analysis of these diverse models is still largely unexplored.

2.2 Performance Evaluation for IQA Models

The rapid development of IQA models raises a fundamental question: how to reliably characterize model behavior and select appropriate methods for different application scenarios. Conventional evaluation protocols primarily rely on global correlation-based metrics, including the PLCC, SRCC, and KRCC, which compress model behavior into a single dataset-level statistic. To partially alleviate the uniform treatment of ranking errors, Wu *et al.* proposed Perceptual Weighted Rank Correlation (PWRC), which introduces perceptual weighting and ignores pairs with imperceptible quality differences [48]. However, PWRC remains a global metric and is therefore incapable of characterizing how performance varies across different levels of quality or quality differences. Beyond correlation-based evaluation, several counter-example-driven protocols have been proposed to expose model failure cases through targeted stress testing. The Maximum Differentiation (MAD) competition [49] and its multi-model extension gMAD [50] formulate evaluation as a falsification process by searching for image pairs that maximally contradict a target model, with performance summarized in terms of aggressiveness and resistance. Complementarily, Ma *et al.* [51] introduced

large-scale computational stress tests, including the D-test for pristine-versus-distorted discrimination, the L-test for monotonicity across distortion levels, and the P-test based on billions of automatically generated perceptually discriminable pairs (DIPs) inferred from agreement among multiple full-reference IQA models. These tests reveal extensive failure cases and provide fine-grained behavioral insights that are not captured by conventional benchmarks. From a psychophysical perspective, Eigen-Distortions [52] analyze extremal sensitivity directions of an IQA model via the Fisher Information Matrix and compare model predictions with human detection thresholds, yielding an absolute measure of perceptual alignment. While highly informative, this approach is computationally demanding and not intended for large-scale or routine model comparison.

3 METHOD

3.1 Preliminary: Generalized Correlation Coefficient

Let $x = \{x_i\}_{i=1}^n$ and $y = \{y_i\}_{i=1}^n$ be two sets of real-valued variables associated with n items, such as predicted and ground-truth image quality scores, respectively. The *Generalized Correlation Coefficient* (GCC), as proposed by Kendall and Gibbons [23], is defined as:

$$\Gamma = \frac{\sum_{i,j=1}^n a_{ij} b_{ij}}{\sqrt{\sum_{i,j=1}^n a_{ij}^2} \cdot \sqrt{\sum_{i,j=1}^n b_{ij}^2}}. \quad (1)$$

Here, a_{ij} and b_{ij} are antisymmetric functions that quantify the relational difference between item pairs, i.e., $a_{ij} = -a_{ji}$ and $b_{ij} = -b_{ji}$.

GCC Instantiations: PLCC, SRCC, and KRCC

By choosing appropriate definitions for a_{ij} and b_{ij} , Eqn. (1) recovers several well-known correlation metrics [48]:

- **Pearson Linear Correlation Coefficient (PLCC):**

Let $p_i = x_i$ (i.e., predicted scores), $q_i = y_i$ (i.e., MOS), and define

$$a_{ij} = p_i - p_j, \quad b_{ij} = q_i - q_j. \quad (2)$$

Substituting Eqn. (2) into Eqn. (1), it is straightforward to show that Eqn. (1) reduces to

$$\Gamma = \frac{\sum_i (p_i - \bar{p})(q_i - \bar{q})}{\sqrt{\sum_i (p_i - \bar{p})^2} \cdot \sqrt{\sum_i (q_i - \bar{q})^2}}, \quad (3)$$

where $\bar{p}$ and $\bar{q}$ denote mean quality. This is exactly the standard form of PLCC.

- **Spearman Rank Correlation Coefficient (SRCC):**

Let $p_i = \text{rank}(x_i)$, $q_i = \text{rank}(y_i)$, and again use Eqn. (2). This yields the SRCC, which measures monotonic correlation via rank differences.

- **Kendall Rank Correlation Coefficient (KRCC):** Let $p_i = x_i$ and $q_i = y_i$, and define

$$a_{ij} = \text{sgn}(p_i - p_j), \quad b_{ij} = \text{sgn}(q_i - q_j), \quad (4)$$

where the sign function is defined as

$$\text{sgn}(z) = \begin{cases} 1, & \text{if } z > 0, \\ -1, & \text{if } z < 0, \\ 0, & \text{if } z = 0. \end{cases} \quad (5)$$

Then, Eqn. (1) reduces to the KRCC.

These formulations illustrate that traditional correlation metrics are specific instantiations of the GCC framework, with their behavior governed by the choice of a_{ij} and b_{ij} .

3.2 Limitations of Global Correlation-Based Metrics

Despite their widespread use, PLCC, SRCC, and KRCC exhibit several fundamental limitations when used to evaluate IQA models:

- 1) **Lack of Fine-Grained Evaluability.** All three metrics aggregate over all image pairs equally, implicitly assuming uniform perceptual relevance. Consequently, they fail to resolve how model performance varies across different absolute quality levels (MOS) or pairwise quality differences ($|\Delta\text{MOS}|$). As illustrated in Fig. 1, CLIP-IQA achieves superior performance on high-MOS images, whereas NIQE is more discriminative when image pairs exhibit small $|\Delta\text{MOS}|$ values. Global correlation metrics, by collapsing performance into a single scalar, completely obscure these complementary behaviors.
- 2) **Sensitivity to Quality Distribution.** PLCC, SRCC, and KRCC are highly sensitive to the distribution of quality scores. As shown in Fig. 2, two subsets sampled from the PIPAL dataset with identical sizes but differing MOS distributions produce reversed model rankings under PLCC and SRCC. This sensitivity arises from the uniform treatment of all (i, j) pairs in the correlation computation, which conflates model capability with the specific characteristics of the test set. Ideally, an evaluation metric should remain robust to variations in the underlying MOS distribution, providing stable and consistent comparisons across datasets.
- 3) **Inability to Characterize Performance Landscape.** Global metrics fail to capture the underlying landscape of model behavior, whereas the geometric properties of a continuous correlation surface provide a formal measure of *structural robustness*. For example, a fragmented landscape reveals latent instabilities and hypersensitivity to quality perturbations that are invisible to scalar summaries. Furthermore, the performance gradient across this surface serves as **optimization guidance**. By visualizing transition boundaries and perceptual constraints, such as sensitivity drops at specific bitrates, researchers can identify generalization limits, offering a systematic path for targeted IQA algorithmic refinement.

3.3 Our GMC Measure

Given a test set with n images, predicted scores of the evaluated IQA model $P = \{p_1, p_2, \dots, p_n\}$ and the corresponding ground-truth MOS values $Q = \{q_1, q_2, \dots, q_n\}$, our goal is to estimate the model IQA performance both *locally* (at specific MOS and $|\Delta\text{MOS}|$) and *robustly* (against imbalanced MOS distributions). To achieve this, we model the IQA model performance visualization in a 3D manner (shown as Fig. 3), spanning MOS (x-axis) and ($|\Delta\text{MOS}|$),




Fig. 3. Overview of the proposed 3D GMC performance surface for fine-grained IQA evaluation. The framework consists of a **Granularity Modulator** for analyzing local performance variations conditioned on MOS and $|\Delta\text{MOS}|$, and a **Distribution Regulator** for mitigating biases from non-uniform quality distributions. The resulting correlation surface represents performance as a joint function of MOS and $|\Delta\text{MOS}|$.

with the z-axis representing weighted correlation indicator (e.g., PLCC, SRCC, or KRCC). To motivate our design, consider a *hypothetical* approach where one constructs local subsets via hard thresholding, i.e., selecting pairs satisfying $|q_i - Q_k^s| < \delta$ and $|q_i - q_j| < Q_k^d$ with a predefined threshold δ . In this scheme, a point in the 3D space is represented by (Q_k^s, Q_k^d, V_k) , where Q_k^s denotes the target MOS, Q_k^d represents the $|\Delta\text{MOS}|$ of interest, and V_k reflects model performance over the selected subset. By sampling different representatives, the performance landscape can, in principle, be obtained via plane fitting. However, this approach suffers from three critical limitations: selecting an appropriate δ is nontrivial; binary selection creates discontinuities between neighboring points; and subset construction cannot ensure uniform quality distributions. Consequently, minor variations induce large fluctuations in V_k , rendering plane fitting unreliable. To account for this, build upon the GCC model (Eqn. (1)), we respectively propose a *Granularity Modulator* and *Distribution Regulator* to achieve a smooth and robust performance estimation at each fine-grained point as follows,

$$\Gamma_k = \frac{\sum_{i,j=1}^n w_k^{ij} a_{ij} b_{ij}}{\sqrt{\sum_{i,j=1}^n w_k^{ij} a_{ij}^2} \cdot \sqrt{\sum_{i,j=1}^n w_k^{ij} b_{ij}^2}}, \quad (6)$$

where

$$w_k^{ij} = \underbrace{\frac{P_k^s(i,j)}{\text{Absolute Quality Score}} \times \frac{P_k^d(i,j)}{\text{Pairwise Quality Difference}}}_{\text{Granularity Modulator}} \times \underbrace{\frac{P_k^t(i,j)}{\text{Quality Distribution}}}_{\text{Distribution Regulator}} \quad (7)$$

and $P_k^s(i,j)$ and $P_k^d(i,j)$ are the contribution weights of i -th and j -th images for the correlation measure at the specific quality scale Q_k^s and quality difference Q_k^d in our Granularity Modulator module and $P_k^t(i,j)$ is distribution regularization term in our Distribution Regulator module. When we respectively adopt the PLCC, SRCC and KRCC as the localized performance indicator, the Eqn. (6) can be formed by:

$$\Gamma_k(\text{PLCC}) = \frac{\sum_{i,j=1}^n w_k^{ij} (p_i - p_j)(q_i - q_j)}{\sqrt{\sum_{i,j=1}^n w_k^{ij} (p_i - p_j)^2} \sqrt{\sum_{i,j=1}^n w_k^{ij} (q_i - q_j)^2}}, \quad (8)$$

$$\Gamma_k(\text{SRCC}) = \frac{\sum_{i,j=1}^n w_k^{ij} (r_{p_i} - r_{p_j})(r_{q_i} - r_{q_j})}{\sqrt{\sum_{i,j=1}^n w_k^{ij} (r_{p_i} - r_{p_j})^2} \sqrt{\sum_{i,j=1}^n w_k^{ij} (r_{q_i} - r_{q_j})^2}}, \quad (9)$$

$$\Gamma_k(\text{KRCC}) = \frac{\sum_{i,j=1}^n w_k^{ij} \text{sgn}(p_i - p_j) \text{sgn}(q_i - q_j)}{\sqrt{\sum_{i,j=1}^n w_k^{ij} \text{sgn}(p_i - p_j)^2} \sqrt{\sum_{i,j=1}^n w_k^{ij} \text{sgn}(q_i - q_j)^2}}, \quad (10)$$

where r_{p_i} and r_{q_i} are the rank of p_i and p_j , respectively. The details of the two modules are described in the following.

3.3.1 Granularity Modulator

At a query MOS Q_k^s , we adopt a Gaussian model to estimate the contribution weight of each image pair (I_i and I_j) in the mode performance measure smoothly. In particular,

$$P_k^s(i,j) = \exp\left(-\frac{(Q_k^s - q_i)^2}{2\sigma_i^2} - \frac{(Q_k^s - q_j)^2}{2\sigma_j^2}\right), \quad (11)$$

where σ_i and σ_j are the standard deviations of the subjective quality ratings for images I_i and I_j . Herein, the standard deviations can be obtained from the test dataset if it provides or estimated by a Beta distribution followed by [53]. Herein, the reason we adopt the Gaussian modeling lies in that human ratings (denoted as R_i and R_j for I_i and I_j) usually follows a Gaussian distribution with the MOS, q_i or q_j as the mean value [54], [55], i.e., $R_i \sim \mathcal{N}(q_i, \sigma_i^2)$, $R_j \sim \mathcal{N}(q_j, \sigma_j^2)$. Assuming that the quality scores q_i and q_j are statistically independent, the joint probability of both falling at the target MOS Q_k^s is exactly equal to the product of their individual Gaussian probabilities:

$$P(q_i = Q_k^s \wedge q_j = Q_k^s) = P(q_i = Q_k^s) \cdot P(q_j = Q_k^s). \quad (12)$$

Analogously, let R_i and R_j denote the underlying random variables corresponding to the subjective ratings of images I_i and I_j , respectively. Under the one-shot estimation assumption, the MOS realizations q_i and q_j are treated as estimates of their corresponding means. Consequently,

the signed quality difference $R_i - R_j$ follows a Gaussian distribution:

$$R_i - R_j \sim \mathcal{N}(q_i - q_j, \sigma_i^2 + \sigma_j^2). \quad (13)$$

Since our GMC formulation aims to characterize the discrimination capability of IQA models under different perceptual difference magnitudes, we adopt the absolute quality difference irrespective of direction. Accordingly, the corresponding random variable $|R_i - R_j|$ follows a Folded Normal distribution. Therefore, the probability density that the image pair (I_i, I_j) falls at the target absolute quality difference Q_k^d is evaluated as:

$$P_k^d(i, j) = \frac{1}{\sqrt{2\pi(\sigma_i^2 + \sigma_j^2)}} \left[\exp\left(-\frac{(Q_k^d - (q_i - q_j))^2}{2(\sigma_i^2 + \sigma_j^2)}\right) + \exp\left(-\frac{(Q_k^d + (q_i - q_j))^2}{2(\sigma_i^2 + \sigma_j^2)}\right) \right], \quad Q_k^d \geq 0. \quad (14)$$

Based on the weight modulation results $P_k^s(i, j)$ and $P_k^d(i, j)$, our Granularity Modulator acts as a localized observer, enabling a fine-grained assessment of the model capability on each level of MOS and $|\Delta\text{MOS}|$.

3.3.2 Distribution Regulator

To mitigate the correlation bias caused by the imbalanced MOS distribution, we propose a *Distribution Regulator* in our GMC. Specifically, we define the $P_k^t(i, j)$ in Eqn. (7) by:

$$P_k^t(i, j) = \frac{1}{\mathcal{D}(q_i)} \cdot \frac{1}{\mathcal{D}(q_j)}, \quad (15)$$

where $\mathcal{D}(q_i)$ and $\mathcal{D}(q_j)$ are the estimated density of quality q_i and q_j . Herein, take $\mathcal{D}(q_i)$ as an example, we considered two cases in the density estimation:

- 1) **Standard deviation is available.** As illustrated in Fig. 4, when image-wise standard deviations of quality ratings are available in the test set, we estimate the local density at q_i by accumulating Gaussian kernels centered at all samples:

$$\mathcal{D}(q_i) = \frac{1}{n} \sum_{u=1}^n \exp\left(-\frac{(q_u - q_i)^2}{2\sigma_u^2}\right), \quad (16)$$

where σ_u denotes the sample-specific standard deviation associated with q_u .

- 2) **Standard deviation is unavailable.** In this case, we introduce a symmetric kernel to extract a kernel-smoothed density estimation. Specifically, we first discretize the normalized MOS range (e.g., $[1, 100]$) into $\mathcal{Y}$ equal-width bins and compute the sample frequency of each bin, then the $\mathcal{D}(q_i)$ can be estimated by,

$$\mathcal{D}(q_i) \triangleq \int_{\mathcal{Y}} k(y, y') p(y) dy, \quad (17)$$

where $p(y)$ is the sample frequency of the bin in which q_i is located. $k(\cdot, \cdot)$ is a symmetric and translation-invariant smoothing kernel. Formally, for continuous variables z and z' , the kernel satisfies $k(z, z') = k(z', z)$ and $\nabla_z k(z, z') + \nabla_{z'} k(z', z) = 0$ [56]. By evaluating this kernel at the discrete bin values y and y' , we convolve it with the empirical density distribution to




Fig. 4. Density estimation at a target quality score q_i . Each sample q_u contributes to the density at q_i via a Gaussian kernel with sample-dependent standard deviation.

extract a kernel-smoothed version that accounts for the overlap in information of data samples in nearby quality bins. Without loss of generality, we adopt the Gaussian kernel in our GMC measure.

Design Rationale: Why Kernel Smoothing Is Preferred in Quality Density Estimation? Though the sample frequency of each bin is widely adopted for the category density estimation in the classification task, we argue that the two limitations lie: (1) The hard assignment of scores to bins, ignoring uncertainty in subjective ratings (e.g., an image with MOS 70.5 might reasonably contribute to both the $[66-70]$ and $[71-75]$ bins); and (2) Certain target values may have no data at all, and the empty bins can cause division-by-zero errors in Eqn. (15), motivating our interpolation manner density smoothing. In Sec. 4.3, our experiments also verify that our kernel-smoothed density estimation strategy results in more robust IQA model performance evaluation.

3.3.3 3D Performance Surface Modeling

Based upon Eqn. (6), we model the IQA model performance in a 3D visualization space (Q_k^s, Q_k^d, V_k) which captures not only the perceptual quality but also integrates the concept of quality discrimination, allowing us to jointly assess model performance in terms of both quality estimation and quality differentiation. To better compare the performance of different IQA models and analyze the performance topology, the performance plane can be fitted when different points (k) are sampled. Herein, the reliability of the plane fitting is higher as the number of sampled points increases. However, exhaustively traversing all possible quality scales and difference across the dataset is inefficient, especially when the dataset volume is large.

Latin Hypercube Sampling. To efficiently explore the parameter space and ensure a representative selection of samples, we employ Latin Hypercube Sampling (LHS) [57], [58]¹ strategy to generate a representative set of K sampling points $\{(Q_k^s, Q_k^d)\}_{k=1}^K$. Specifically, LHS partitions the value range of each variable, i.e., $[Q_{\min}^s, Q_{\max}^s]$ and $[Q_{\min}^d, Q_{\max}^d]$ into K non-overlapping intervals of equal width. For each

1. For efficient implementation, we assume a uniform marginal distribution along each dimension, avoiding complicated Cumulative Distribution Function (CDF) estimation.

sampling index k , a random value is drawn from a unique interval along each dimension, ensuring stratified coverage without redundancy. The sampling coordinates are computed as:

$$Q_k^s = \frac{\pi_x(k) - u_k}{K} \cdot (Q_{\max}^s - Q_{\min}^s) + Q_{\min}^s, \quad (18)$$

$$Q_k^d = \frac{\pi_y(k) - u_k}{K} \cdot (Q_{\max}^d - Q_{\min}^d) + Q_{\min}^d, \quad (19)$$

where π_x and π_y denote independent random permutations of $\{1, 2, \dots, K\}$, and $u_k \sim \mathcal{U}(0, 1)$ is a uniform random variable used to perturb sampling within each interval. This formulation guarantees that each projection of the K samples covers the domain uniformly, yielding improved statistical efficiency and coverage compared to naïve random sampling.

Performance Surface Modeling. Once the weighted correlation score Γ_k is computed for each sampling point, we fit a continuous 3D performance response surface to characterize the model’s behavior across the joint space of quality scale and quality difference. This fitting process can be expressed as:

$$\text{GMC}(Q_k^s, Q_k^d) = \hat{\Gamma}(Q_k^s, Q_k^d) \approx \mathcal{F}\left(\{(Q_k^s, Q_k^d)\}_{k=1}^N\right), \quad (20)$$

where $\mathcal{F}(\cdot)$ denotes a two-dimensional nonparametric function approximation based on Local Linear Kernel Regression. This method performs weighted linear modeling within local neighborhoods, effectively reducing boundary bias.

Global Performance Aggregation. In addition to the localized analysis, a global performance score can also be derived from our GMC (denoted as GMC_g) via a surface integration:

$$\text{GMC}_g = \frac{1}{A} \int_{Q_{\min}^s}^{Q_{\max}^s} \int_{Q_{\min}^d}^{Q_{\max}^d} \hat{\Gamma}(x, y) dx dy, \quad (21)$$

where $A = (Q_{\max}^s - Q_{\min}^s) \times (Q_{\max}^d - Q_{\min}^d)$ represents the normalization area formed by the ranges of Q^s and Q^d . Compared with the traditional global indicators (*e.g.*, PLCC, SRCC or KRCC), our derived indicator GMC_g not only reflects the model’s predictive capability for absolute image quality scales but also captures its sensitivity to quality difference discrimination with the distribution bias mitigated, offering a comprehensive and robust performance assessment of IQA model.

4 EXPERIMENT

4.1 Experimental Settings

We evaluate the proposed GMC measure on both FR and NR IQA models. For FR-IQA, five representative metrics are considered: PSNR, SSIM [24], MS-SSIM [25], LPIPS [30], and DISTS [5]. These FR metrics are evaluated on the full KADID-10k database [59] and the PIPAL dataset [60]. For NR-IQA, we select five recent blind-quality models: NIQE [61], CLIP-IQA (CLIP-IQA and CLIP-IQA+) [62], QualiCLIP [63] and MANIQA [64]. Each NR model is trained on the KonIQ-10k dataset [65] and then tested on two large-scale

“in-the-wild” datasets: LIVE-Challenge (denoted as LIVEC) [66] and SPAQ [67]. In addition, we extend our experiments to two AIGC benchmark datasets, AGIQA-3K [68] and AIG-CIQA2023 [69]. We select three representative recent models, namely CLIP-IQA [21], TIER [70], and DepictQA [55], and conduct cross-dataset testing on two quality dimensions, *i.e.*, *Quality* and *Alignment*. We discretize quality scores into $|\mathcal{Y}| = 100$ bins and employ Latin hypercube sampling with $K = 100$ sample points when computing GMC. We adopt the SRCC, *i.e.*, $\Gamma_k(\text{SRCC})$ in Eqn. (9) for the localized performance evaluation by default.

4.2 Applications of GMC Measure

(1) Insight into FR-IQA and NR-IQA Models. To reveal the local prediction behaviors of existing IQA models, we visualize the joint prediction distribution of each model by our proposed GMC measure, as illustrated in Fig. 5. In addition to qualitative visualization, we enable quantitative comparisons by dividing the MOS range of each dataset into three equal-width quality scales: Low Quality (LQ), Medium Quality (MQ), and High Quality (HQ). Correspondingly, we define three diagnostic intervals: Low Difference (LD), Medium Difference (MD), and High Difference (HD). For each quality scale and diagnostic interval, we compute two fine-grained diagnostic metrics, GMC_s and GMC_d , respectively, formulated as:

$$\text{GMC}_s = \frac{1}{A'} \int_{Q_{\min}^d}^{Q_{\max}^d} \int_{Q_1^s}^{Q_2^s} \hat{\Gamma}(x, y) dx dy, \quad (22)$$

$$\text{GMC}_d = \frac{1}{A''} \int_{Q_{\min}^s}^{Q_{\max}^s} \int_{Q_1^d}^{Q_2^d} \hat{\Gamma}(x, y) dx dy, \quad (23)$$

where $[Q_1^s, Q_2^s]$ and $[Q_1^d, Q_2^d]$ define the boundaries of the selected quality scale and interval, respectively, and A' , A'' are the integration area for normalization. The comparison results of different IQA models are shown in Tab. 1. From the figure and table, several key insights emerge:

- *Different IQA models exhibit distinct behaviors across MOS and $|\Delta\text{MOS}|$.* For instance, on the SPAQ dataset, MANIQA achieves consistently high GMC_s scores in the medium- and high-quality regions, indicating strong robustness under perceptually favorable conditions. In contrast, CLIP-IQA+ exhibits relatively stronger performance in the low-quality and low-difference regimes, demonstrating superior discrimination capability in fine-grained quality assessment scenarios. Similar observations can also be found in AIGC scenarios (see Table 2 and Fig. 6). On the AGIQA-3K dataset, CLIP-IQA shows superior performance in the low-quality and low-difference regions, while TIER gradually outperforms it as the quality level or quality difference increases. More importantly, we observe that PWRC tends to select the same best-performing model as GMC_s in the high-quality regime in most cases. However, it fails to reveal the optimal models in other quality intervals. For instance, on the PIPAL dataset, global metrics consistently identify LPIPS as the best-performing model, whereas GMC_s reveals that its performance in




Fig. 5. Visualization of correlation surfaces generated by different IQA models. The “GMC” response (vertical axis) is shown as a function of the “MOS” (left horizontal axis) and “ $|\Delta\text{MOS}|$ ” (right horizontal axis).

low-quality and low-difference regions is inferior to MS-SSIM.

- Global metrics fail to capture local behaviors.* For example, although MS-SSIM achieves a high overall SRCC on the KADID-10k dataset, its performance (GMC_s) in the low-quality region is notably inferior to that of DISTS. Similarly, while LPIPS exhibits strong performance (GMC_d) in the low-difference regime, it attains a lower overall PLCC than DISTS. These discrepancies indicate that relying solely on global correlation metrics may lead to misleading conclusions in practical evaluation scenarios. The global metrics, including PLCC, SRCC, KRCC, and PWRC, consistently designate a single model as the best-performing one across each entire dataset (e.g., TIER on AGIQA-3K Quality and CLIP-IQA on the remaining three settings). This monolithic ranking collapses heterogeneous local behaviors into a single scalar, thereby obscuring regime-specific strengths essential for scenario-aware deployment.
- Performance degradation under fine-grained quality differences.* Across all four datasets, a consistent performance drop can be observed when moving from high- to low-difference regimes, indicating that existing IQA models struggle to reliably discriminate subtle quality variations. For instance, on KADID-10k, the average GMC_d of representative methods decreases by more than 15% from the high-quality to low-quality region, with similar degradation trends observed on PIPAL, LIVEC, and SPAQ. Although some models (e.g., LPIPS or CLIP-based approaches) maintain competitive global correlation scores, their performance under low-difference conditions remains substantially weaker. These results suggest that strong global correlation does not necessarily imply robust fine-grained perceptual discrimination, and such limitations may be obscured when evaluation relies solely on global metrics. Furthermore, across all alignment experiments, we observe a consistent performance degradation as the quality difference decreases ($\text{LD} < \text{MD} < \text{HD}$), indicating limited sensitivity




Fig. 6. Visualization of correlation surfaces generated by different IQA models. The “GMC” response (vertical axis) is shown as a function of the “MOS” (left horizontal axis) and “ $|\Delta\text{MOS}|$ ” (right horizontal axis).

TABLE 1

Performance comparison of FR- and NR-IQA models across different IQA datasets in terms of PLCC, SRCC, GMC_s, and GMC_d.

KADID-10k [59] (FR)										
Model	GMC _s			GMC _d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
PSNR	0.4509	0.5195	0.6113	0.4047	0.5204	0.6458	0.5557	0.6757	0.4876	21.6874
SSIM [24]	0.3927	0.3286	0.5555	0.3443	0.4245	0.4918	0.5755	0.6188	0.4468	19.2923
MS-SSIM [25]	0.6448	0.6269	0.7357	0.5608	0.6729	0.7652	0.6802	0.8256	0.6350	25.9811
LPIPS [30]	0.6363	0.6453	0.7428	0.5590	0.6779	0.7790	0.7484	0.8224	0.6303	26.2150
DISTS [30]	0.6682	0.6041	0.6867	0.5458	0.6586	0.7505	0.8057	0.8137	0.6254	25.4562
PIPAL [60] (FR)										
Model	GMC _s			GMC _d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
PSNR	0.5364	0.4842	0.3927	0.3769	0.4762	0.5860	0.4033	0.4065	0.2762	5.5463
SSIM [24]	0.5926	0.5853	0.5520	0.4665	0.5763	0.6893	0.4992	0.5041	0.3489	7.2886
MS-SSIM [25]	0.7049	0.6406	0.5637	0.5428	0.6480	0.7525	0.5624	0.5621	0.3975	8.2772
LPIPS [30]	0.6351	0.6531	0.6808	0.5419	0.6556	0.7624	0.5827	0.5854	0.4099	8.7563
DISTS [30]	0.6414	0.6500	0.6733	0.5421	0.6560	0.7629	0.5797	0.5790	0.4069	8.6281
LIVEC [66] (NR)										
Model	GMC _s			GMC _d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
NIQE [61]	0.4768	0.3741	0.2523	0.3035	0.4031	0.4965	0.4791	0.4495	0.3063	6.9598
CLIP-IQA [62]	0.6123	0.5854	0.5897	0.5152	0.6221	0.7005	0.6883	0.6955	0.5065	10.9756
CLIP-IQA+ [62]	0.7561	0.7133	0.7091	0.6606	0.7613	0.8254	0.8312	0.8045	0.6109	12.6591
QualiCLIP [63]	0.7285	0.6540	0.6181	0.6013	0.7075	0.7799	0.7967	0.7553	0.5618	12.0622
MANIQA [64]	0.7640	0.7601	0.7508	0.6879	0.7892	0.8477	0.8401	0.8328	0.6404	13.0035
SPAQ [67] (NR)										
Model	GMC _s			GMC _d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
NIQE [61]	0.6532	0.6296	0.5103	0.4977	0.6186	0.7119	0.6692	0.6928	0.4928	16.5133
CLIP-IQA [62]	0.4541	0.5968	0.6258	0.4548	0.5704	0.6668	0.6432	0.6653	0.4645	16.3178
CLIP-IQA+ [62]	0.7882	0.7917	0.7685	0.7086	0.8137	0.8724	0.8492	0.8516	0.6492	20.7123
QualiCLIP [63]	0.7766	0.7822	0.7613	0.6970	0.8038	0.8649	0.8518	0.8440	0.6386	20.5147
MANIQA [64]	0.7100	0.7969	0.8090	0.6901	0.7983	0.8629	0.8552	0.8509	0.6471	20.8074

in distinguishing subtle consistency variations. By decoupling performance along both the absolute quality and pairwise difference dimensions, GMC exposes these fine-grained limitations that remain invisible to conventional global indicators.

(2) IQA Model Selection for High-Quality Image Retrieval. IQA models are commonly used to identify high-quality images across different quality regimes, such as selecting top-quality samples from generative outputs or

screening out low-quality images in degraded datasets. Standard global correlation metrics may fail to guide optimal model selection in such scenario-specific searches, particularly when the overall ranking performance of competing models is similar. To evaluate the effectiveness of our GMC measure, we construct two types of **query sets** separately for each dataset (KADID-10K and SPAQ) based on their respective MOS distributions: a high-quality set formed by images in the top quartile of MOS, and a low-quality set formed by images in the bottom quartile.

TABLE 2
Performance comparison of AGIQA models across AIGC datasets in terms of PLCC, SRCC, KRCC, PWRC, GMC_s , and GMC_d .

AGIQA-3K (Quality)										
Model	GMC_s			GMC_d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
CLIP-IQA [21]	0.6401	0.5243	0.4080	0.3901	0.5738	0.7068	0.6645	0.6137	0.4278	13.6306
TIER [70]	0.5958	0.5542	0.4610	0.3888	0.5736	0.7084	0.6717	0.6525	0.4594	14.2661
DepictQA [55]	0.6120	0.5201	0.4514	0.3873	0.5726	0.7071	0.6646	0.6184	0.4494	13.7831
AGIQA-3K (Alignment)										
Model	GMC_s			GMC_d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
CLIP-IQA [21]	0.4924	0.3930	0.3488	0.3280	0.4507	0.5503	0.5438	0.4874	0.3365	12.3144
TIER [70]	0.4129	0.3793	0.3511	0.2873	0.4021	0.5028	0.4796	0.4524	0.3077	11.0080
DepictQA [55]	0.4160	0.3824	0.3665	0.2911	0.4090	0.5119	0.4975	0.4525	0.3171	11.0915
AIGCIQA2023 (Quality)										
Model	GMC_s			GMC_d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
CLIP-IQA [21]	0.4659	0.6412	0.3713	0.2924	0.4727	0.6640	0.7177	0.6887	0.4786	12.3565
TIER [70]	0.5452	0.6114	0.3908	0.3189	0.5097	0.6889	0.6794	0.6674	0.4628	11.8953
DepictQA [55]	0.5705	0.5551	0.4174	0.3307	0.5187	0.6827	0.6231	0.6154	0.4371	11.0295
AIGCIQA2023 (Alignment)										
Model	GMC_s			GMC_d			PLCC	SRCC	KRCC	PWRC
	LQ	MQ	HQ	LD	MD	HD				
CLIP-IQA [21]	0.4368	0.4573	0.3100	0.2538	0.4055	0.5297	0.4908	0.5262	0.3673	9.3519
TIER [70]	0.3760	0.4776	0.3152	0.2408	0.3826	0.5118	0.4659	0.4655	0.3171	7.0663
DepictQA [55]	0.3824	0.3105	0.2311	0.1960	0.3200	0.4197	0.3176	0.3159	0.2194	5.2971

For AIGC datasets (AGIQA-3K and AIGCIQA2023), we consider both the *Quality* and *Alignment* dimensions, and construct corresponding high-quality and high-alignment query sets using the top quartile samples in each dimension. For each query set, we compare the top-3 images selected by the model chosen via GMC_s , global SRCC, and PWRC, respectively.

The mean MOS of the selected images is reported in Table 3, and the corresponding selections are visualized in Fig. 7. Results demonstrate that models selected by GMC_s consistently identify higher-quality images in high-quality query sets and more accurately discriminate lower-quality images in low-quality query sets, outperforming those selected based on the best SRCC and PWRC in both scenarios. For AIGC datasets, the results are summarized in Table 4 and Fig. 8. Results demonstrate that, across both traditional and AIGC datasets, models selected by GMC_s consistently outperform those chosen based on SRCC or PWRC. Specifically, GMC_s enables more accurate identification of high-quality images and more reliable discrimination of lower-quality samples in traditional IQA datasets, while in AIGC scenarios, it further facilitates stable retrieval of images with higher perceptual quality or better semantic alignment.

(3) Fine-Grained Quality Optimization. To assess the practical impact of GMC in optimization settings, we consider fine-grained quality optimization, where an IQA model is used as an objective function to guide the iterative refinement of enhanced (distorted) images. In this regime, perceptual differences between reference and optimized images are subtle, and effective optimization depends on the ability of the IQA model to discriminate small yet meaningful quality variations.

Table 1 indicates that MS-SSIM consistently outperforms LPIPS and DISTS in terms of discriminative capability in

TABLE 3
Performance comparison for high-quality image selection on traditional IQA datasets when the best model is selected by SRCC, PWRC, or GMC_s . Blue star (★) denotes the model selected by SRCC, gray star (★) denotes the model selected by PWRC, and red star (★) denotes the model selected by GMC_s .

Query Set	Dataset	Model	Selection	Mean MOS
High-quality	KADID-10k	MS-SSIM LPIPS	★ ★ ★	4.68 4.75
	SPAQ	CLIP-IQA+ MANIQA	★ ★ ★	76.37 86.17
Low-quality	KADID-10k	MS-SSIM DISTS LPIPS	★ ★ ★	1.89 2.05 2.42
	SPAQ	MANIQA CLIP-IQA+	- ★ ★ ★	21.63 28.08

the low quality-difference (LD) regime. To validate this finding, we design an *adversarial optimization* based comparison framework on the PIPAL dataset.

Specifically, given a reference image y and an initial distorted image x_0 that exhibits a slight quality degradation relative to y , a target IQA model is optimized to enhance image quality (minimizing its loss), whereas a competing metric is adversarially constrained to prevent improvement. Accordingly, we consider two complementary optimization settings, which are formulated as the following constrained problems:

$$\min_x f_{\text{MS-SSIM}}(x, y), \text{ s.t. } 0 \leq f_{\text{LPIPS}}(x, y) - f_{\text{LPIPS}}(x_0, y) \leq \epsilon, \quad (24)$$




Fig. 7. Comparison of high-quality images retrieved by the best FR/NR IQA models selected via GMC and the SRCC criterion. (a) and (b) show the top-3 images from the high-quality subsets of KADID-10K and SPAQ, respectively; (c) and (d) show the top-3 images from the low-quality subsets of the same datasets. The selected model for each subfigure is indicated in its title.

and

$$\min_x f_{LPIPS}(x, y), \text{ s.t. } 0 \leq f_{MS-SSIM}(x, y) - f_{MS-SSIM}(x_0, y) \leq \epsilon. \quad (25)$$

Here, $f_{MS-SSIM}$ and f_{LPIPS} denote the corresponding loss functions of perceptual quality metrics, and ϵ controls the tolerance threshold that bounds the allowable degradation of the adversarial metric. In practice, we optimize the above objectives using their unconstrained penalized forms, where violations of the adversarial constraints are softly penalized through hinge-based regularization. The threshold ϵ is introduced to prevent the optimization from exploiting trivial shortcuts, where the overall loss could be reduced by excessively deteriorating the adversarial metric without yielding a genuine improvement in the target metric.

As shown in Fig. 9, when optimized against LPIPS or

DISTS as adversarial metrics, improvements in MS-SSIM consistently translate into perceptual quality enhancement. In contrast, when MS-SSIM serves as the defender metric, optimizing LPIPS or DISTS leads to visually unsatisfactory results, even though their metric values improve. These observations align with the GMC analysis, confirming that models exhibiting higher GMC scores under low quality-difference regimes possess superior robustness and reliability for fine-grained quality optimization tasks.

(4) Model Integration for Enhanced IQA Performance. Beyond quality optimization, GMC can also reveal complementary relationships among IQA models, enabling performance improvement through model integration. To demonstrate this, we conduct experiments on both FR- and NR-




Fig. 8. Comparison of images retrieved by the best IQA models selected via SRCC, PWRC, and GMC_s across quality and alignment criteria. (a) and (b) show the top-3 images from the high-quality subsets of AGIQA-3K and AIGCIQA2023, respectively; (c) and (d) show the top-3 images from the high-alignment subsets of the same datasets. The selected model for each subfigure is indicated in its title.

IQA datasets. On the FR-IQA dataset PIPAL, three representative models are selected: a traditional hand-crafted feature model (MS-SSIM) and two deep feature models (LPIPS and DISTS), with LPIPS achieving the best individual performance. Notably, although both SRCC and PWRC suggest that DISTS outperforms MS-SSIM in terms of overall performance, GMC reveals that MS-SSIM provides more complementary information to LPIPS, particularly on low-quality and low-difference subsets, as indicated by the larger GMC values in Table 1. This suggests that integrating LPIPS with MS-SSIM is expected to yield superior performance, despite DISTS exhibiting higher SRCC performance.

The integration is formulated as a weighted combination of the metrics, taking into account their opposite polar-

ity: lower LPIPS and DISTS values indicate better quality, whereas higher MS-SSIM values indicate better quality. Specifically, we define

$$\begin{aligned} V_{Inter1} &= (1 - V_{LPIPS}) + V_{MS-SSIM}, \\ V_{Inter2} &= (1 - V_{DISTS}) + V_{MS-SSIM}, \end{aligned} \quad (26)$$

where V denotes the quality values predicted by each model. Experimental results in Fig. 10 validate that the “LPIPS & MS-SSIM” integration achieves SRCC = 0.6202, compared with “LPIPS & DISTS” (SRCC = 0.6019), resulting in a relative improvement of 3.0%.

A similar trend also holds on the NR-IQA dataset SPAQ: integrating the baseline QualiCLIP model with MANIQA, as revealed by GMC, outperforms the SRCC-selected CLIP-



Fig. 9. Qualitative comparison of adversarial image quality optimization on the PIAPL Dataset. Columns 1 and 2 show the reference and initial distorted images. Columns 3 and 5 present optimization results with LPIPS and DISTS as target metrics, respectively, under MS-SSIM constraints. Columns 4 and 6 show the reverse setting, where MS-SSIM is optimized under LPIPS or DISTS constraints.




Fig. 10. Comparison of IQA model integration results. (a) and (c) show the IQA model complementarity identified by our GMC on the PIPAL and SPAQ datasets, respectively. (b) and (d) report the SRCC and PWRC of individual IQA models and their corresponding integrated counterparts.

TABLE 4

Performance comparison for high-quality and high-alignment image selection on AIGC datasets when the best model is selected by SRCC, PWRC, or GMC_s. Blue star (★) denotes the model selected by SRCC, gray star (★) denotes the model selected by PWRC, and red star (★) denotes the model selected by GMC_s.

Query Set	Dataset	Model	Selection	Mean MOS
Quality	AGIQA-3K	DepictQA	★ ★	3.48
		TIER	★	4.14
	AIGCIQA2023	TIER	★	3.11
Alignment	AGIQA-3K	DepictQA	★ ★	3.16
		CLIP-IQA	★ ★	3.53
	AIGCIQA2023	DepictQA	★	3.65
		CLIP-IQA	★	3.04
		TIER	★	3.20
			★	3.27

IQA+ model, demonstrating that the complementary integration strategy generalizes across both FR and NR scenarios.

We further extend this analysis to the AIGCIQA2023 dataset. The results are presented in Fig. 11. In the quality




Fig. 11. Comparison of AGIQA model integration results. (a) and (c) show the AGIQA model complementarity identified by our GMC on the AIGCIQA2023 datasets at two dimensions: quality and alignment, respectively. (b) and (d) report the SRCC and PWRC of individual IQA models and their corresponding integrated counterparts.

dimension, we consider integrating CLIP-IQA with other candidate models. Although both SRCC and PWRC indicate that TIER outperforms DepictQA globally, GMC analysis reveals a different insight: CLIP-IQA achieves the best performance in the medium-quality region, while DepictQA performs better in low- and high-quality regions, indicating stronger complementarity between CLIP-IQA and DepictQA. Guided by this observation, we perform model integration and find that the combination of CLIP-IQA and DepictQA achieves a higher SRCC than that of CLIP-IQA and TIER.

An analogous trend is observed regarding the Alignment dimension. Although TIER achieves higher global SRCC and PWRC scores than DepictQA, GMC surfaces reveal that DepictQA exhibits stronger complementary behavior with CLIP-IQA across different alignment regimes. Consequently, integrating CLIP-IQA with DepictQA yields better performance than integrating it with TIER.

These results indicate that GMC not only facilitates informed model selection but also provides critical guidance for identifying complementary models, enabling more effective integration across traditional and AIGC IQA scenarios.

(5) Robustness Under Varied Sampling Distributions. In real-world scenarios, image datasets often exhibit imbalanced or non-uniform quality distributions. Under such conditions, a robust evaluation metric is crucial for ensuring fair and reliable comparisons among IQA models. To demonstrate the robustness of GMC_g under varying sampling distributions, we conduct systematic experiments on two widely used benchmarks, PIPAL and SPAQ, involving a diverse set of both FR and NR IQA models.

To simulate realistic distributional shifts, we adopt a Gaussian probability-weighted sampling strategy to construct imbalanced subsets from each dataset. For each dataset, three sampling variants: unimodal, bimodal, and trimodal, are generated. Each variant further produces three sub-variants, yielding a total of nine samples for analysis. As the MOS scales differ across datasets, all sampled scores are normalized to the range $[0, 100]$, partitioned into bins of size 1, and the number of samples in each bin is counted. The resulting sampling distributions are shown in Fig. 12.

We evaluate all considered IQA models on the constructed subsets using both the global SRCC metric and the proposed GMC_g . As shown in Fig. 12, GMC_g consistently exhibits lower variability than SRCC across the nine subsets for every model-dataset pair, demonstrating its enhanced robustness to sampling-induced distributional bias.

4.3 Ablation Study

Study on Kernel-Smoothed Density Estimation. To validate the importance of kernel-smoothed density estimation, we conduct an ablation study comparing the full GMC against a variant that replaces $D(q_i)$ in Eqn. (16) with the reciprocal of the raw sample frequency. For each sampling instance, both surfaces are re-estimated directly from the sampled data, and GMC is computed via global integration over the instance-specific fitted surface. We conduct 50 extreme sampling trials. The hybrid sampling strategy combines: (i) concentrated distributions with small variances; (ii) heuristic bias toward data-sparse regions via histogram-guided Gaussian center placement; and (iii) fully random mixtures with broad variances from a Dirichlet distribution. Results in Table 5 show that the kernel-smoothed variant achieves lower variance, demonstrating superior robustness under distributional perturbations.

Study on Sampling Size in 3D Performance Surface Modeling. To evaluate the impact of sampling density on the stability of our 3D performance surface modeling, we conduct an ablation study on the number of sampling points used in the Latin Hypercube Sampling (LHS) scheme. We vary the number of LHS sampling points (K in Eqn. (19)) from 10 to 1,000 and compute the resulting performance GMC_g for each setting. For comparison, we also report results using naïve random sampling with the same sample sizes. The results are summarized in Fig. 13 and 14.

It is evident that random sampling leads to high variance and unstable evaluations, even when the sample size is sufficiently large. In contrast, LHS achieves stable and reliable estimates with significantly fewer samples. Notably, when the number of LHS samples exceeds 100, the computed metrics converge and remain stable, indicating that LHS not only enhances sampling efficiency but also ensures robustness in surface modeling. These findings validate the use of LHS as an effective and practical strategy for performance surface construction, striking a favorable balance between computational efficiency and statistical reliability.

5 CONCLUSION

In this work, we have revisited the fundamental problem of IQA model evaluation and argued that conventional global




Fig. 12. Robustness comparison between SRCC and GMC_g across different quality distributions. Rows 1 and 3 present the nine sampling sets from PIPAL and SPAQ, respectively, while rows 2 and 4 show the corresponding variations in SRCC and GMC_g values.




Fig. 13. Variation of GMC_g across IQA models versus sampling size under Latin Hypercube Sampling.




Fig. 14. Variation of GMC_g across IQA models versus sampling size under random sampling.

TABLE 5
Effect of kernel smoothing on GMC_g variance across IQA models.

	Kernel	Model	Standard Deviation	Average
PIPAL	✗	PSNR	0.1438	0.1664
		SSIM	0.1717	
		MS-SSIM	0.1932	
		LPIPS	0.1610	
		DISTS	0.1621	
	✓	PSNR	0.1476	0.1651
		SSIM	0.1678	
		MS-SSIM	0.1924	
		LPIPS	0.1601	
		DISTS	0.1574	
SPAQ	✗	NIQE	0.1445	0.1778
		CLIP-IQA	0.1966	
		CLIP-IQA+	0.1665	
		QualiCLIP	0.1672	
		MANIQA	0.2141	
	✓	NIQE	0.1484	0.1737
		CLIP-IQA	0.1777	
		CLIP-IQA+	0.1615	
		QualiCLIP	0.1671	
		MANIQA	0.2138	
AIGCQA2023	✗	CLIP-IQA	0.1928	0.1783
		TIER	0.1665	
		DepictQA	0.1758	
	✓	CLIP-IQA	0.1912	0.1626
		TIER	0.1551	
		DepictQA	0.1416	

correlation metrics are inherently limited in their ability to reveal fine-grained model behavior and are highly sensitive to quality distribution biases. To address these challenges, we proposed a GMC framework that extends classical correlation analysis with granularity-aware modulation and distribution-aware regularization. By constructing a continuous correlation surface over quality score and quality difference, GMC enables a structured and interpretable characterization of IQA performance.

The GMC analysis reveals that a critical performance decoupling where strong global correlation frequently fails to translate into fine-grained discrimination—particularly in low-difference regimes. This finding suggests that future IQA model design must move beyond broad statistical alignment and prioritize localized sensitivity to satisfy the stringent requirements of fine-grained perceptual tasks. Furthermore, GMC advocates for a shift from searching for a “universally optimal” metric to identifying “scenario-appropriate” models. By visualizing the performance landscape, we provide a principled basis for model deployment and optimization tailored to specific perceptual constraints. Finally, by exposing how IQA models inherit and manifest biases from their training distributions, GMC could serve as a diagnostic lens to guide the acquisition of balanced datasets and the development of distribution-robust architectures. In essence, GMC transforms evaluation from a passive post-hoc measurement into an active, diagnostic bridge that links model behavior, dataset composition, and

real-world deployment requirements.

REFERENCES

- [1] Z. Wang and A. C. Bovik, “Modern image quality assessment,” in *Synthesis Lectures on Image, Video, and Multimedia Processing*, 2006, pp. 1–156.
- [2] G. Zhai and X. Min, “Perceptual image quality assessment: a survey,” *Science China Information Sciences*, vol. 63, pp. 1–52, 2020.
- [3] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, “Comparison of full-reference image quality models for optimization of image processing systems,” *International Journal of Computer Vision*, vol. 129, no. 4, pp. 1258–1281, 2021.
- [4] J. Mannos and D. Sakrison, “The effects of a visual fidelity criterion of the encoding of images,” *IEEE Transactions on Information Theory*, vol. 20, no. 4, pp. 525–536, 1974.
- [5] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, “Image quality assessment: Unifying structure and texture similarity,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 5, pp. 2567–2581, 2020.
- [6] Z. Wang and A. C. Bovik, “Mean squared error: Love it or leave it? A new look at signal fidelity measures,” *IEEE Signal Processing Magazine*, vol. 26, no. 1, pp. 98–117, 2009.
- [7] B. Girod, “What’s wrong with mean-squared error?” in *Digital Images and Human Vision*, 1993, pp. 207–220.
- [8] W. Zhang, K. Ma, J. Yan, D. Deng, and Z. Wang, “Blind image quality assessment using a deep bilinear convolutional neural network,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 1, pp. 36–47, 2018.
- [9] J. Kim, A.-D. Nguyen, and S. Lee, “Deep CNN-based blind image quality predictor,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 1, pp. 11–24, 2018.
- [10] S. Bosse, D. Maniry, K.-R. Müller, T. Wiegand, and W. Samek, “Deep neural networks for no-reference and full-reference image quality assessment,” *IEEE Transactions on Image Processing*, vol. 27, no. 1, pp. 206–219, 2017.
- [11] C. Chen, J. Mo, J. Hou, H. Wu, L. Liao, X. Min, and W. Lin, “TOPIQ: A top-down approach from semantics to distortions for image quality assessment,” *IEEE Transactions on Image Processing*, 2024.
- [12] J. You and J. Korhonen, “Transformer for image quality assessment,” in *IEEE International Conference on Image Processing*, 2021, pp. 1389–1393.
- [13] M. Cheon, S.-J. Yoon, B. Kang, and J. Lee, “Perceptual image quality assessment with transformers,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 433–442.
- [14] J. Ke, Q. Wang, Y. Wang, P. Milanfar, and F. Yang, “MUSIQ: Multi-scale image quality transformer,” in *IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5148–5157.
- [15] W. Zhang, G. Zhai, Y. Wei, X. Yang, and K. Ma, “Blind image quality assessment via vision-language correspondence: A multitask learning perspective,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 14 071–14 081.
- [16] K. Xu, L. Liao, J. Xiao, C. Chen, H. Wu, Q. Yan, and W. Lin, “Boosting image quality assessment through efficient transformer adaptation with local feature enhancement,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2662–2672.
- [17] H. Wu, Z. Zhang, E. Zhang, C. Chen, L. Liao, A. Wang, C. Li, W. Sun, Q. Yan, G. Zhai *et al.*, “Q-Bench: A benchmark for general-purpose foundation models on low-level vision,” in *International Conference on Learning Representations*, 2024.
- [18] T. Wu, J. Zou, J. Liang, L. Zhang, and K. Ma, “VisualQuality-R1: Reasoning-induced image quality assessment via reinforcement learning to rank,” *Advances in Neural Information Processing Systems*, vol. 38, pp. 88 167–88 190, 2026.
- [19] X. Zhang, W. Lin, and Q. Huang, “Fine-grained image quality assessment: A revisit and further thinking,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 5, pp. 2746–2759, 2021.
- [20] Y. Fang, H. Zhu, Y. Zeng, K. Ma, and Z. Wang, “Perceptual quality assessment of smartphone photography,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3677–3686.
- [21] J. Wang, K. C. Chan, and C. C. Loy, “Exploring CLIP for assessing the look and feel of images,” in *AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 2555–2563.

- [22] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "completely blind" image quality analyzer," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 209–212, 2012.
- [23] M. G. Kendall and J. D. Gibbons, *Rank Correlation Methods*, 5th ed. Oxford University Press, 1990.
- [24] 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, 2004.
- [25] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *Asilomar Conference on Signals, Systems & Computers*, 2003, pp. 1398–1402.
- [26] Z. Wang and Q. Li, "Information content weighting for perceptual image quality assessment," *IEEE Transactions on Image Processing*, vol. 20, no. 5, pp. 1185–1198, 2010.
- [27] H. R. Sheikh and A. C. Bovik, "Image information and visual quality," *IEEE Transactions on Image Processing*, vol. 15, no. 2, pp. 430–444, 2006.
- [28] 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, 2011.
- [29] W. Xue, L. Zhang, X. Mou, and A. C. Bovik, "Gradient magnitude similarity deviation: A highly efficient perceptual image quality index," *IEEE Transactions on Image Processing*, vol. 23, no. 2, pp. 684–695, 2013.
- [30] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [31] X. Liao, B. Chen, H. Zhu, S. Wang, M. Zhou, and S. Kwong, "Deep-WSD: Projecting degradations in perceptual space to wasserstein distance in deep feature space," in *ACM International Conference on Multimedia*, 2022, pp. 970–978.
- [32] H. Zhu, B. Chen, L. Zhu, S. Wang, and W. Lin, "DeepDC: Deep distance correlation as a perceptual image quality evaluator," *IEEE Transactions on Image Processing*, vol. 34, pp. 7859–7873, 2025.
- [33] B. Chen, H. Zhu, L. Zhu, S. Wang, J. Pan, and S. Wang, "Debiased mapping for full-reference image quality assessment," *IEEE Transactions on Multimedia*, 2025.
- [34] A. K. Moorthy and A. C. Bovik, "A two-step framework for constructing blind image quality indices," *IEEE Signal Processing Letters*, vol. 17, no. 5, pp. 513–516, 2010.
- [35] M. A. Saad, A. C. Bovik, and C. Charrier, "Blind image quality assessment: A natural scene statistics approach in the DCT domain," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3339–3352, 2012.
- [36] G. Zhai, X. Wu, X. Yang, W. Lin, and W. Zhang, "A psychovisual quality metric in free-energy principle," *IEEE Transactions on Image Processing*, vol. 21, no. 1, pp. 41–52, 2011.
- [37] K. Gu, G. Zhai, X. Yang, W. Zhang, and L. Liang, "No-reference image quality assessment metric by combining free energy theory and structural degradation model," in *IEEE International Conference on Multimedia and Expo*, 2013, pp. 1–6.
- [38] B. Chen, L. Zhu, C. Kong, H. Zhu, S. Wang, and Z. Li, "No-reference image quality assessment by hallucinating pristine features," *IEEE Transactions on Image Processing*, vol. 31, pp. 6139–6151, 2022.
- [39] L. Kang, P. Ye, Y. Li, and D. Doermann, "Convolutional neural networks for no-reference image quality assessment," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1733–1740.
- [40] S. Su, Q. Yan, Y. Zhu, C. Zhang, X. Ge, J. Sun, and Y. 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*, 2020, pp. 3667–3676.
- [41] H. Zhu, L. Li, J. Wu, W. Dong, and G. Shi, "MetaIQA: Deep meta-learning for no-reference image quality assessment," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 143–14 152.
- [42] W. Zhang, K. Ma, G. Zhai, and X. Yang, "Task-specific normalization for continual learning of blind image quality models," *IEEE Transactions on Image Processing*, vol. 33, pp. 1898–1910, 2024.
- [43] B. Chen, L. Zhu, G. Li, F. Lu, H. Fan, and S. Wang, "Learning generalized spatial-temporal deep feature representation for no-reference video quality assessment," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 4, pp. 1903–1916, 2021.
- [44] H. Wu, Z. Zhang, W. Zhang, C. Chen, L. Liao, C. Li, Y. Gao, A. Wang, E. Zhang, W. Sun *et al.*, "Q-Align: Teaching LMMs for visual scoring via discrete text-defined levels," *arXiv preprint arXiv:2312.17090*, 2023.
- [45] K. Liu, Z. Zhang, W. Li, R. Pei, F. Song, X. Liu, L. Kong, and Y. Zhang, "Dog-IQA: Standard-guided zero-shot mllm for mix-grained image quality assessment," *arXiv preprint arXiv:2410.02505*, 2024.
- [46] H. Zhu, H. Wu, Y. Li, Z. Zhang, B. Chen, L. Zhu, Y. Fang, G. Zhai, W. Lin, and S. Wang, "Adaptive image quality assessment via teaching large multimodal model to compare," *arXiv preprint arXiv:2405.19298*, 2024.
- [47] B. Chen, S. Pan, D. Wu, L. Xie, X. Sui, L. Zhu, and H. Zhu, "Mitigating perception bias: A training-free approach to enhance lmm for image quality assessment," *AAAI Conference on Artificial Intelligence*, 2026.
- [48] Q. Wu, H. Li, F. Meng, and K. N. Ngan, "A perceptually weighted rank correlation indicator for objective image quality assessment," *IEEE Transactions on Image Processing*, vol. 27, no. 5, pp. 2499–2513, 2018.
- [49] Z. Wang and E. P. Simoncelli, "Maximum differentiation (MAD) competition: A methodology for comparing computational models of perceptual quantities," *Journal of Vision*, vol. 8, no. 12, pp. 8–8, 2008.
- [50] 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, 2018.
- [51] K. Ma, Z. Duanmu, Q. Wu, Z. Wang, H. Yong, H. Li, and L. Zhang, "Waterloo exploration database: New challenges for image quality assessment models," *IEEE Transactions on Image Processing*, vol. 26, no. 2, pp. 1004–1016, 2016.
- [52] A. Berardino, V. Laparra, J. Ballé, and E. Simoncelli, "Eigen-distortions of hierarchical representations," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [53] S. Ferrari and F. Cribari-Neto, "Beta regression for modelling rates and proportions," *Journal of Applied Statistics*, vol. 31, no. 7, pp. 799–815, 2004.
- [54] W. Zhang, K. Ma, G. Zhai, and X. Yang, "Uncertainty-aware blind image quality assessment in the laboratory and wild," *IEEE Transactions on Image Processing*, vol. 30, pp. 3474–3486, 2021.
- [55] Z. You, J. Gu, Z. Li, X. Cai, K. Zhu, T. Xue, and C. Dong, "Descriptive image quality assessment in the wild," *arXiv preprint arXiv:2405.18842*, 2024.
- [56] Y. Yang, K. Zha, Y. Chen, H. Wang, and D. Katabi, "Delving into deep imbalanced regression," in *International Conference on Machine Learning*. PMLR, 2021, pp. 11 842–11 851.
- [57] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 42, no. 1, pp. 55–61, 2000.
- [58] T. J. Santner, B. J. Williams, W. I. Notz, and B. J. Williams, *The design and analysis of computer experiments*. Springer, 2003, vol. 1.
- [59] H. Lin, V. Hosu, and D. Saupe, "KADID-10k: A large-scale artificially distorted IQA database," in *International Conference on Quality of Multimedia Experience*, 2019.
- [60] G. Jinjin, C. Haoming, C. Haoyu, Y. Xiaoxing, J. S. Ren, and D. Chao, "PIPAL: a large-scale image quality assessment dataset for perceptual image restoration," in *European Conf. on Computer Vision*. Springer, 2020, pp. 633–651.
- [61] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "completely blind" image quality analyzer," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 209–212, 2013.
- [62] J. Wang, K. C. Chan, and C. C. Loy, "Exploring CLIP for assessing the look and feel of images," in *AAAI Conference on Artificial Intelligence*, 2023, pp. 1357–1365.
- [63] L. Agnolucci, L. Galteri, and M. Bertini, "Quality-aware image-text alignment for opinion-unaware image quality assessment," *arXiv preprint arXiv:2403.11176*, 2024.
- [64] S. Yang, T. Wu, S. Shi, S. Lao, Y. Gong, M. Cao, J. Wang, and Y. Yang, "MANIQA: Multi-dimension attention network for no-reference image quality assessment," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1191–1200.
- [65] 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, pp. 4041–4056, 2020.

- [66] 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, 2015.
- [67] Y. Fang, H. Zhu, Y. Zeng, K. Ma, and Z. Wang, "Perceptual quality assessment of smartphone photography," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3677–3686.
- [68] C. Li, Z. Zhang, H. Wu, W. Sun, X. Min, X. Liu, G. Zhai, and W. Lin, "AGIQ-3K: An open database for AI-generated image quality assessment," *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.
- [69] J. Wang, H. Duan, J. Liu, S. Chen, X. Min, and G. Zhai, "AIG-CIQA2023: A large-scale image quality assessment database for AI-generated images: from the perspectives of quality, authenticity and correspondence," in *CAAI International Conference on Artificial Intelligence*. Springer, 2023, pp. 46–57.
- [70] J. Yuan, X. Cao, J. Che, Q. Wang, S. Liang, W. Ren, J. Lin, and X. Cao, "TIER: Text-image encoder-based regression for AIGC image quality assessment," *arXiv preprint arXiv:2401.03854*, 2024.



Baoliang Chen (Member, IEEE) received the Ph.D. degree in Computer Science from the City University of Hong Kong, Hong Kong, in 2022. From 2024 to 2026, he was an Associate Professor with South China Normal University. Since 2026, he has been a Postdoctoral Research Fellow with Nanyang Technological University, Singapore. His research interests include image/video quality assessment and transfer learning.




Danni Huang is currently pursuing the B.E. degree with the Department of Computer Science, South China Normal University, China. Her research interests include image quality assessment and image quality enhancement.



Hanwei Zhu (Member, IEEE) received Ph.D. degree in computer science from the City University of Hong Kong, Hong Kong, in 2025. He is currently a research scientist with the Alibaba-NTU Global e-Sustainability CorpLab (ANGEL) at Nanyang Technological University. His research interests include perceptual image processing, computational vision, and computational photography.



Lingyu Zhu (Member, IEEE) received his Ph.D. degree in computer science from the City University of Hong Kong, Hong Kong SAR, China, in 2024. He is currently a postdoctoral researcher with the Department of Computer Science at City University of Hong Kong. His research interests include image/video compression, image/video enhancement, image/video quality assessment, and deep learning.




Wei Zhou (Senior Member, IEEE) received the Ph.D. degree from the University of Science and Technology of China in 2021, jointly with the University of Waterloo, Canada. They were a Visiting Professor with Dalian University of Technology, a Visiting Scholar with the National Institute of Informatics, Japan, a Research Assistant with Intel, and a Research Intern with Microsoft Research and Alibaba Cloud. They were a Post-Doctoral Fellow with the University of Waterloo. They are currently an Assistant Professor with Cardiff University, U.K. Their research interests include multimedia computing, perceptual image processing, and computational vision.



Shiqi Wang (Senior Member, IEEE) received the Ph.D. degree in computer application technology from Peking University in 2014. He is currently an Professor with the Department of Computer Science, City University of Hong Kong. He has proposed more than 70 technical proposals to ISO/MPEG, ITU-T, and AVS standards. He authored or coauthored more than 300 refereed journal articles/conference papers, including more than 100 IEEE Transactions. His research interests include video compression, image/video quality assessment, video coding for machine, and semantic communication. He served or serves as an Associate Editor for IEEE TIP, TCSVT, TMM, TCyber, Access, and APSIPA Transactions on Signal and Information Processing.



Yuming Fang (Fellow, IEEE) received the B.E. degree from Sichuan University, Chengdu, China, the M.S. degree from the Beijing University of Technology, Beijing, China, and the Ph.D. degree from Nanyang Technological University, Singapore. He is currently a Professor with the School of Computing and Artificial Intelligence, Jiangxi University of Finance and Economics, Nanchang, China. His research interests include visual attention modeling, visual quality assessment, image retargeting, computer vision, and 3D image/video processing. He serves as an Associate Editor for IEEE TRANSACTIONS ON MULTIMEDIA. He is an Editorial Board of Signal Processing: Image Communication.



Weisi Lin (Fellow, IEEE) received the Ph.D. degree from the King's College, University of London, U.K. He is currently a Professor with the College of Computing and Data Science, Nanyang Technological University. His areas of expertise include image processing, perceptual signal modeling, video compression, and multimedia communication, in which he has published over 200 journal articles, over 230 conference papers, filed seven patents, and authored two books. He has been an invited/panelist/keynote/tutorial speaker at over 20 international conferences. He is a fellow of IET and an Honorary Fellow of Singapore Institute of Engineering Technologists. He has been the Technical Program Chair of IEEE ICME 2013, PCM 2012, QoMEX 2014, and IEEE VCIP 2017. He has been an Associate Editor of IEEE TRANSACTIONS ON IMAGE PROCESSING, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE TRANSACTIONS ON MULTIMEDIA, and IEEE SIGNAL PROCESSING LETTERS. He was a Distinguished Lecturer of Asia-Pacific Signal and Information Processing Association (APSIPA) from 2012 to 2013 and the IEEE Circuits and Systems Society from 2016 to 2017.