

ME-IQA: Memory-Enhanced Image Quality Assessment via Re-Ranking

Kanglong Fan¹[0000–0002–5023–4900], Tianhe Wu¹(✉)[0009–0007–9557–9818],
Wen Wen¹, Jianzhao Liu²,
Le Yang², Yabin Zhang², Yiting Liao², Junlin Li², and Li Zhang²

¹ City University of Hong Kong, Hong Kong SAR, China

² ByteDance Inc., China

tianhewu-c@my.cityu.edu.hk

Abstract. Reasoning-induced vision-language models (VLMs) advance image quality assessment (IQA) with textual reasoning, yet their scalar scores often lack sensitivity and collapse to a few values, so-called *discrete collapse*. We introduce ME-IQA, a plug-and-play, test-time memory-enhanced re-ranking framework. It (i) builds a memory bank and retrieves semantically and perceptually aligned neighbors using reasoning summaries, (ii) reframes the VLM as a probabilistic comparator to obtain pairwise preference probabilities and fuse this ordinal evidence with the initial score under Thurstone’s Case V model, and (iii) performs gated reflection and consolidates memory to improve future decisions. This yields denser, distortion-sensitive predictions and mitigates *discrete collapse*. Experiments across multiple IQA benchmarks show consistent gains over strong reasoning-induced VLM baselines, existing non-reasoning IQA methods, and test-time scaling alternatives.

Keywords: IQA · Memory · Test Time Scaling

1 Introduction

Image quality assessment (IQA), a fundamental problem in computer vision [48, 51], underpins applications from mobile photography [8] and video streaming [5] to image restoration [11]. The rapid progress of vision-language models (VLMs) has catalyzed a shift from conventional score regressors [64, 72, 73] toward VLM-based IQA [57, 68], aiming to leverage perception-reasoning capabilities for more human-aligned judgments.

Reasoning-induced approaches [23, 61] prompt VLMs to generate step-by-step reasoning before emitting a scalar score and often generalize better than direct regression. Yet they frequently suffer from *discrete collapse*: distinct-quality images receive coarse, nearly identical scores. As illustrated in Fig. 1 (a), on KADID [24], VisualQuality-R1 [61] produces scores concentrated at a few discrete levels, even though there are substantial perceptual differences among images assigned to the same level (see Fig. 1 (b)). We attribute this to an objective mismatch: VLMs are pretrained to generate discrete tokens [3, 7], not continuous




Fig. 1: ME-IQA mitigates discrete collapse and improves distortion sensitivity. (a) Scatter plot of ME-IQA versus baseline (VisualQuality-R1) scores on KADID-10K [24]. Points concentrate at a few discrete levels along the x-axis for the baseline, indicating *discrete collapse*, whereas ME-IQA spreads predictions more densely. (b) Four highlighted cases illustrate substantial re-ranking: for a low baseline score around 2.0 (Img A vs. Img B) and a high baseline score around 4.0 (Img C vs. Img D), ME-IQA respectively elevates the higher-quality image and suppresses the lower-quality one relative to the baseline.

perceptual quantities [64]. When forced into numeric prediction, they gravitate to textually salient numbers (*e.g.*, “3.0”, “4.0”, “5.0”), coarsely quantizing perception and dulling sensitivity to fine-grained distortions.

To resolve the collapse, existing work explores two potential remedies. The first extracts token probabilities for special numeric symbols and averages them to form a score [57, 68]. The second aggregates pairwise comparisons under Thurstone’s model [42], as in Depict-QA [69]. While both yield continuous outputs, each has drawbacks. Single-stimulus prompting lacks explicit comparative context and undershoots subtle quality differences. Pure pairwise schemes, though perceptually grounded, scale poorly on large sets (*e.g.*, KonIQ-10K [14]) and are impractical for online testing³.

Anchor-based designs strike a middle ground by comparing a query to a small fixed set (*e.g.*, Compare2Score [80]), tapping the VLM’s comparative strength with manageable cost. However, static anchors under-represent long-tail or novel distortions and degrade under distribution shift. In contrast, human observers judge quality relative to contextually retrieved memories of similar stimuli [13, 38], beyond static reference. A similar principle appears in retrieval-augmented generation (RAG) [20] and external memory for VLMs [15, 28], where instance-specific context improves robustness.

Motivated by dynamic human perceptual memory, we propose **ME-IQA**, a test-time **Memory-Enhanced** re-ranking framework that mitigates *discrete collapse* in reasoning-induced VLMs while avoiding the brittleness of static anchors.

³ A stream of queries arrives sequentially and each should be processed immediately without access to future inputs (“no peeking”).

Concretely, for each query, ME-IQA (i) constructs a hybrid memory bank and retrieves a semantically and perceptually aligned neighborhood using reasoning-derived summaries as retrieval keys; (ii) reframes the VLM as a comparator to elicit pairwise preference probabilities and fuses this ordinal evidence with the initial score by optimizing Thurstone’s Case V [42], yielding denser, perceptually sensitive predictions (see Fig. 1); and (iii) performs gated reflection and consolidates the processed case into memory, strengthening future decisions.

Because ME-IQA operates purely at test time, requires only black-box access to the underlying VLM, and leaves training untouched, it is *plug-and-play*: practitioners can apply it to existing reasoning-induced VLMs without retraining, architectural changes, or extra supervision. Extensive experiments on standard IQA benchmarks show consistent gains over strong VLM baselines and test-time scaling alternatives, narrowing the gap to human MOS distributions while preserving the generalization benefits of reasoning-induced models. In summary, ME-IQA offers an efficient path toward fine-grained, human-aligned IQA, effectively countering *discrete collapse* and improving sensitivity to nuanced distortions.

2 Related Work

2.1 NR-IQA Models

No-reference image quality assessment (NR-IQA) aims to predict visual quality from a single distorted image without access to any pristine reference. Early NR-IQA approaches rely on natural scene statistics and degradation-specific priors [25, 29, 30, 49, 50, 52], modeling regularities of undistorted images to detect departures due to artifacts. With deep learning, convolutional [72] and Transformer-based [64] architectures became dominant, extracting semantics-aware features that better correlate with perceptual quality. Deep NR-IQA has been optimized either by regression loss [2, 17, 18, 27, 40, 64, 68], which treats quality as an absolute scalar, or by ranking loss [4, 9, 26, 42, 46, 47, 71, 72], casting quality as an intrinsically relative quantity. In parallel, memory-like ideas appeared as exemplar storage and local interpolation [21, 44, 58, 59, 63, 66], but were ultimately constrained by hand-crafted or shallow features.

Leveraging large-scale pretraining and multimodal priors, VLM-based NR-IQA has rapidly gained traction [60]. Representative supervised finetuning (SFT) methods include Q-Align [57] and DeQA-Score [68], which directly regress scores. More recently, reasoning-induced models generate textual reasoning before scoring [23, 61]: Q-Insight [23] and VisualQuality-R1 [61] use reinforcement learning (RL) to promote explicit reasoning, and EvoQuality [54] combines self-consistency voting with reinforcement learning to rank. Despite improved generalization, these models remain vulnerable to *discrete collapse*, where scalar predictions concentrate on a few values, dulling sensitivity to fine-grained distortions.




Fig. 2: Online testing process of ME-IQA. For each online query, the VLM produces a reasoning $\bar{r}_i$ and an initial score $\tilde{s}_i$, which is transformed to s_i via a five-parameter monotonic mapping. The reasoning $\bar{r}_i$ is summarized as a concise quality description $\mathbf{r}_i$ and embedded to retrieve a neighborhood from a memory bank comprising Anchor Memory (AM) and Contrast Memory (CM). Acting as a comparator, the VLM estimates pairwise preference probabilities between the query and retrieved exemplars. The ordinal evidence is fused with the mapped initial score s_i by optimizing Thurstone’s Case V to yield a refined score s_i^* . If $|s_i^* - s_i| > \epsilon$, a reflection step updates the description $\mathbf{r}_i$ to $\mathbf{r}_i^*$, and the case is consolidated to CM for better future decisions.

2.2 Memory for LLMs/VLMs

Beyond purely parametric memory, recent work equips LLMs/VLMs with explicit external memory and agentic modules [76]. Retrieval-augmented generation (RAG) conditions generation on context fetched at inference time [12, 20]; multimodal variants index images, captions, dense embeddings, or structured facts to ground reasoning and improve robustness under distribution shift [16, 65]. In agentic settings, working memory may be represented as plain text [34], latent embeddings [45], or structured graphs [6, 62]. Most existing work emphasizes personalization [74, 79] and long-context management [15, 28, 56]. A complementary line studies self-improving agents that learn from prior trajectories [33, 39, 41, 77, 78].

Our framework specializes memory for perceptual quality assessment through (i) a hybrid memory bank balancing stability (offline anchors) and adaptability (online contrast cases), and (ii) reasoning-aware retrieval keyed by concise quality descriptions. Unlike prior VLM-based NR-IQA methods that rely solely on SFT or end-to-end RL, our method operates purely at test time and targets the *discrete collapse* failure mode by injecting local ordinal constraints and consolidating hard cases for future retrieval.

3 ME-IQA

We adopt an online testing paradigm [53, 55], where a stream of NR-IQA queries $\mathcal{Q} = \{q_t\}_{t=1}^T$ arrives sequentially and should be processed on arrival, meaning no peeking into the future. This mirrors realistic streaming deployments that demand immediate, low-latency scoring and strict causality.

Overview. Given the query image $\mathbf{x}_i$ at time t (omitting t when clear), our system proceeds as in Fig. 2. A VLM first produces a free-form reasoning $\tilde{\mathbf{r}}_i$ and an initial scalar prediction $\tilde{s}_i$. We map $\tilde{s}_i$ to the target quality scale via a monotonic logistic transformation [37], yielding s_i . We then summarize $\tilde{\mathbf{r}}_i$ into a concise, self-contained quality description $\mathbf{r}_i$ and embed it with a text encoder to retrieve a small neighborhood $\mathcal{N}$ from a hybrid memory bank $\mathcal{M} = \mathcal{M}_A \cup \mathcal{M}_C$ composed of (i) Anchor Memory (AM) and (ii) Contrast Memory (CM). Next, the VLM acts as a comparator to estimate pairwise preference probabilities between the query and retrieved exemplars, and the mapped initial score s_i is fused with the ordinal evidence via a Thurstone Case V objective to obtain a refined prediction s_i^* . A reflection mechanism is triggered upon inconsistency: if $|s_i^* - s_i| > \epsilon$, the VLM revises the description to $\mathbf{r}_i^*$ and the system consolidates the case into CM. Prompt templates are provided in the supplementary material. We now detail the memory bank (Sec. 3.1) and quality re-ranking (Sec. 3.2).

3.1 Memory Bank

Design. To enable fine-grained, context-aware refinement at test time, we build a hybrid memory that balances stability and adaptability. AM is constructed offline from labeled anchors with ground-truth (GT) scores and provides a stable scaffold across the full quality range, mitigating drift and rectifying extreme mispredictions. CM is grown online from recently processed queries after re-ranking/reflection, continuously absorbing hard or underrepresented cases (distribution shifts, corner cases, emerging artifacts) to enhance local discrimination.

Schema. Each memory item $\mathbf{m}_i \in \mathcal{M}$ is a structured triple:

$$\mathbf{m}_i = (\mathbf{x}_i, \mathbf{r}_i \text{ or } \mathbf{r}_i^*, s_i^*), \quad (1)$$

where $\mathbf{r}_i/\mathbf{r}_i^*$ summarizes semantic content, clarity, texture detail, and perceptual distortions and their impact on quality; s_i^* stores the GT for AM or the refined estimate for CM.

Retrieval. We employ reasoning-aware retrieval: the VLM first compresses the raw chain-of-thought $\tilde{\mathbf{r}}_i$ into a compact, quality-centric description $\mathbf{r}_i$, suppressing boilerplate and content-heavy detours (*e.g.*, “let’s tackle this step by step”). Encoding $\mathbf{r}_i$ with a text encoder and using cosine similarity, we jointly search $\mathcal{M}_A$ and $\mathcal{M}_C$ under a fixed budget K . We split the budget uniformly:

$$K_A = \left\lfloor \frac{K}{2} \right\rfloor, \quad K_C = K - K_A. \quad (2)$$

In AM, we perform GT-stratified retrieval by partitioning scores into B bins over $[1, 5]$ and taking $k_A = \lfloor K_A/B \rfloor$ nearest neighbors per bin (distributing any

remainder randomly for coverage). In CM, we simply take the top- K_C nearest neighbors. The final neighborhood is $\mathcal{N} = \mathcal{N}_A \cup \mathcal{N}_C$ with $|\mathcal{N}| = K$, used by the re-ranking module.

Anchor Memory consolidation. We pre-build $\mathcal{M}_A$ from a labeled set $\mathcal{D} = \{(\mathbf{x}_i, g_i)\}_{i=1}^D$ (e.g., KONIQ-10K). For each $\mathbf{x}_i$, the VLM emits $(\tilde{\mathbf{r}}_i, \tilde{s}_i)$. We fit the standard five-parameter logistic mapping [37] from $\{\tilde{s}_i\}$ to $\{g_i\}$ and store its parameters $\{\beta_k\}_{k=1}^5$. The mapped score is

$$s_i = \beta_1 \left(\frac{1}{2} - \frac{1}{1 + \exp(\beta_2(\tilde{s}_i - \beta_3))} \right) + \beta_4 \tilde{s}_i + \beta_5, \quad (3)$$

which aligns predictions to the GT scale $[1, 5]$ (higher is better). If $|g_i - s_i| > \epsilon$, we prompt the VLM to reflect given $(\mathbf{x}_i, \tilde{\mathbf{r}}_i, s_i, g_i)$ and produce $\mathbf{r}_i^*$; otherwise, $\tilde{\mathbf{r}}_i$ is directly summarized to $\mathbf{r}_i$. We store $(\mathbf{x}_i, \mathbf{r}_i$ or $\mathbf{r}_i^*, s_i^* = g_i)$. AM is static and persists across sessions, together with the learned mapping in Eq. (3).

Contrast Memory consolidation and pruning. During online testing, after re-ranking the current query q_t , we add it to CM. Let $i = D + t$ index the current case. If $|s_i^* - s_i| > \epsilon$, we trigger reflection with $(\mathbf{x}_i, \tilde{\mathbf{r}}_i, s_i, s_i^*)$ to obtain $\mathbf{r}_i^*$. We store $(\mathbf{x}_i, \mathbf{r}_i$ or $\mathbf{r}_i^*, s_i^*)$ with a cached text embedding $\mathbf{e}_i$ and count $c_i=1$. When $|\mathcal{M}_C| > C$, we apply agglomerative-style pruning [31]: repeatedly find the closest pair by cosine distance $d(i, j) = 1 - \mathbf{e}_i^\top \mathbf{e}_j / (\|\mathbf{e}_i\|_2 \|\mathbf{e}_j\|_2)$, keep the higher-count item as the prototype, remove the other item, and add the removed count to the survivor. Pruning stops at $|\mathcal{M}_C| \leq \alpha C$ ($\alpha=0.8$ by default).

3.2 Quality Re-ranking

Pairwise evidence from the VLM. Given the query $\mathbf{x}_i$ and its neighborhood $\mathcal{N}$, the VLM serves as a probabilistic comparator. For each exemplar $j \in \mathcal{N}$, we prompt the model with a binary comparison (query is always “Image A”) and extract the probability of the token “A” from the answer distribution as the soft preference $y_{ij} = P(s_i^* > s_j^*)$ [70, 80].

Thurstone’s Case V with a weak prior. Let s_j^* denote the stored score of neighbor j (fixed constants from memory), and $\Phi(\cdot)$ the standard normal CDF. Under Thurstone’s Case V model [42], the model-implied probability is

$$p_{ij} = \Phi\left(\frac{s_i^* - s_j^*}{\sigma}\right), \quad \sigma = 1. \quad (4)$$

We estimate s_i^* by minimizing the binary cross-entropy (BCE) between p_{ij} and y_{ij} with a quadratic tether to the mapped initial score s_i :

$$\min_{s_i^*} \sum_{j \in \mathcal{N}} \text{BCE}(p_{ij}, y_{ij}) + \lambda (s_i^* - s_i)^2, \quad (5)$$

where λ controls the strength of the prior. The solution nudges s_i^* toward ordinal evidence while remaining conservative when pairwise information is weak or ambiguous.

Closed-form approximation. Although Eq. (5) can be solved by gradient-based optimization, it does not admit a simple closed-form minimizer due to the probit likelihood in the BCE term. For an efficient alternative, we use a common probit linearization that converts each soft preference into a pseudo-observation of the latent score. Specifically, from the optimization objective $y_{ij} \approx \Phi(s_i^* - s_j^*)$ (with $\sigma = 1$), we obtain

$$s_i^* \approx s_j^* + \Phi^{-1}(y_{ij}) \triangleq \mu_{ij}, \quad (6)$$

For numerical stability, before applying Φ^{-1} , we clip y_{ij} into $[\delta, 1 - \delta]$ with a small $\delta > 0$. Replacing the BCE term in Eq. (5) with a squared error yields a ridge-style objective

$$\min_{s_i^*} \sum_{j \in \mathcal{N}} (s_i^* - \mu_{ij})^2 + \lambda (s_i^* - s_i)^2, \quad (7)$$

whose closed-form solution is

$$s_i^* = \frac{\sum_{j \in \mathcal{N}} \mu_{ij} + \lambda s_i}{K + \lambda}. \quad (8)$$

This approximation can be used for fast inference.

4 Experiments

4.1 Experimental Setup

Models under test. We evaluate five reasoning-induced VLMs: Q-Insight [23], VisualQuality-R1 [61], EvoQuality [54], and two proprietary frontier models, Doubao-Seed-1.6-Vision and GPT-5. Q-Insight, VisualQuality-R1, and EvoQuality are built on Qwen-2.5-VL [1] and further fine-tuned via GRPO [36], representing strong open-source baselines for reasoning-based IQA. Unless noted, we use their variants fine-tuned on KONIQ-10K, and we construct the anchor memory for all models uniformly from KONIQ-10K to ensure fairness.

Datasets. We conduct zero-shot evaluation on seven datasets spanning authentic, AI-generated, and synthetic distortions following the same settings in [54]: (i) authentic: SPAQ [8], LIVEW [10]; (ii) AI-generated: AGIQA [22]; (iii) synthetic: CSIQ [19], TID2013 [35], KADID [24], and PIPAL [11].

Implementation details. For the three open-source models, reasoning traces are sufficiently concise; we therefore bypass the summarization step and directly encode the raw reasoning into retrieval embeddings. Quality descriptions are embedded with Qwen3-Embedding-0.6B [75]. Retrieval uses $K = 32$ neighbors, split into $K_A = 16$ from AM and $K_C = 16$ from CM. AM is GT-stratified into $B = 5$ bins with $k_A = \lfloor 16/5 \rfloor = 3$ items per bin and the remaining one item assigned to a random bin; CM contributes the top- K_C nearest neighbors (cosine similarity). The Thurstone fusion prior is set to $\lambda = 0.01$ and the reflection gate to $\epsilon = 0.75$.

Table 1: Performance improvement of ME-IQA over baseline VLM models and comparison with existing non-reasoning IQA methods across seven benchmarks. All models, except the closed-source ones, are trained on KONIQ using the same training protocol. The upper section shows PLCC results, and the lower section shows SRCC results. WAVG denotes the weighted average result, with weights proportional to the number of images in each dataset. The top two results are highlighted in **bold** and underline.

Method	SPAQ	AGIQA	LIVEW	KADID	PIPAL	TID2013	CSIQ	WAVG.
<i>Existing Non-reasoning Methods</i>								
NIQE	0.664	0.533	0.449	0.430	0.120	0.110	0.645	0.349
NIMA	0.838	0.715	0.814	0.701	0.536	0.418	0.854	0.641
MUSIQ	0.868	0.722	0.789	0.668	0.567	0.257	0.860	0.621
CLIP-IQA+	0.866	0.736	0.832	0.762	0.550	0.293	0.900	0.641
MANIQA	0.768	0.723	0.849	0.526	0.603	0.172	0.712	0.582
Q-Align	0.886	0.772	0.853	0.795	0.600	0.282	0.814	0.663
DeQA-Score	0.895	0.809	0.891	0.711	0.608	0.192	<u>0.868</u>	0.652
Compare2Score	0.873	0.756	0.857	0.720	0.596	0.419	<u>0.761</u>	0.686
<i>Q-Insight</i>								
Baseline	0.901	0.823	0.859	0.732	0.604	0.477	0.749	0.714
ME-IQA	<u>0.922</u>	0.840	0.881	0.770	<u>0.635</u>	0.520	0.777	<u>0.744</u>
<i>VisualQuality-R1</i>								
Baseline	0.895	0.827	0.847	0.709	0.567	0.453	0.750	0.698
ME-IQA	0.912	<u>0.850</u>	0.868	0.741	0.599	0.490	0.786	0.726
<i>EvoQuality</i>								
Baseline	0.903	0.827	0.860	0.744	0.602	0.615	0.821	0.748
ME-IQA	0.925	0.847	<u>0.887</u>	<u>0.783</u>	0.642	0.643	0.851	0.777
<i>Doubao-Seed-1.6-vision</i>								
Baseline	0.874	0.821	0.858	0.707	0.517	0.424	0.729	0.678
ME-IQA	0.888	0.851	0.872	0.740	0.561	0.469	0.762	0.711
<i>GPT-5</i>								
Baseline	0.879	0.815	0.851	0.715	0.539	0.399	0.701	0.676
ME-IQA	0.891	0.847	0.860	0.739	0.566	0.447	0.760	0.706
<i>Existing Non-reasoning Methods</i>								
NIQE	0.664	0.533	0.449	0.430	0.120	0.110	0.645	0.349
NIMA	0.856	0.654	0.771	0.693	0.486	0.422	0.844	0.616
MUSIQ	0.863	0.630	0.830	0.650	0.535	0.248	0.815	0.592
CLIP-IQA+	0.864	0.685	0.805	0.756	0.499	0.309	0.890	0.618
MANIQA	0.758	0.636	0.832	0.488	0.551	0.116	0.675	0.534
Q-Align	0.887	0.735	0.860	0.792	0.518	0.284	0.750	0.631
DeQA-Score	0.895	0.729	0.871	0.716	0.565	0.132	<u>0.847</u>	0.614
Compare2Score	0.856	0.758	0.821	0.731	0.588	0.456	0.771	0.688
<i>Q-Insight</i>								
Baseline	0.899	0.766	0.817	0.737	0.528	0.463	0.717	0.683
ME-IQA	0.924	<u>0.796</u>	0.842	<u>0.785</u>	<u>0.566</u>	<u>0.508</u>	0.751	<u>0.719</u>
<i>VisualQuality-R1</i>								
Baseline	0.887	0.760	0.831	0.703	0.509	0.405	0.697	0.661
ME-IQA	0.899	0.791	0.853	0.753	0.551	0.449	0.737	0.696
<i>EvoQuality</i>								
Baseline	0.901	0.782	0.829	0.738	0.543	0.581	0.778	0.716
ME-IQA	<u>0.919</u>	0.809	<u>0.860</u>	0.783	0.588	0.619	0.815	0.751
<i>Doubao-Seed-1.6-vision</i>								
Baseline	0.867	0.748	0.837	0.709	0.483	0.387	0.688	0.648
ME-IQA	0.887	0.775	0.846	0.746	0.533	0.453	0.739	0.687
<i>GPT-5</i>								
Baseline	0.873	0.737	0.823	0.701	0.496	0.373	0.678	0.644
ME-IQA	0.887	0.763	0.832	0.744	0.560	0.424	0.736	0.683

4.2 Results

Improvements over VLM baselines. As summarized in Table 1, ME-IQA consistently and substantially improves all reasoning-induced VLM baselines

Table 2: Comparison with test-time scaling alternatives across seven IQA benchmarks. All models are trained on KONIQ using the same training protocol. The upper section shows PLCC results and the lower section shows SRCC results.

Method	Sec/Img	SPAQ	AGIQA	LIVEW	KADID	PIPAL	TID2013	CSIQ	WAVG.
<i>Q-Insight</i>									
Maj@64	34.97	0.912	0.835	0.870	0.740	0.611	0.487	0.760	0.724
Mean@64	34.97	0.915	0.842	0.875	0.738	0.615	0.490	0.757	0.727
Compare2Score@32	14.02	0.887	0.801	0.842	0.715	0.589	0.460	0.740	0.698
ME-IQA@32	14.34	0.922	0.840	0.881	0.770	0.635	0.520	0.777	0.744
<i>VisualQuality-R1</i>									
Maj@64	34.92	0.903	0.842	0.856	0.722	0.580	0.465	0.762	0.710
Mean@64	34.92	0.906	0.849	0.863	0.720	0.583	0.471	0.772	0.714
Compare2Score@32	13.88	0.881	0.819	0.833	0.689	0.549	0.440	0.737	0.684
ME-IQA@32	14.30	0.912	0.850	0.868	0.741	0.599	0.490	0.786	0.726
<i>EvoQuality</i>									
Maj@64	34.95	0.915	0.839	0.876	0.757	0.613	0.624	0.832	0.759
Mean@64	34.95	0.922	0.847	0.883	0.749	0.610	0.623	0.837	0.761
Compare2Score@32	13.93	0.888	0.815	0.853	0.725	0.590	0.602	0.809	0.734
ME-IQA@32	14.32	0.925	0.847	0.887	0.783	0.642	0.643	0.851	0.777
<i>Q-Insight</i>									
Maj@64	34.97	0.917	0.788	0.832	0.747	0.537	0.487	0.733	0.700
Mean@64	34.97	0.924	0.796	0.838	0.740	0.532	0.485	0.724	0.700
Compare2Score@32	14.02	0.890	0.756	0.799	0.711	0.511	0.449	0.711	0.668
ME-IQA@32	14.34	0.924	0.796	0.842	0.785	0.566	0.508	0.751	0.719
<i>VisualQuality-R1</i>									
Maj@64	34.92	0.905	0.779	0.851	0.725	0.534	0.422	0.715	0.681
Mean@64	34.92	0.909	0.789	0.858	0.719	0.530	0.427	0.719	0.683
Compare2Score@32	13.88	0.876	0.747	0.818	0.687	0.491	0.392	0.683	0.647
ME-IQA@32	14.30	0.899	0.791	0.853	0.753	0.551	0.449	0.737	0.696
<i>EvoQuality</i>									
Maj@64	34.95	0.919	0.799	0.850	0.753	0.560	0.597	0.800	0.734
Mean@64	34.95	0.924	0.811	0.866	0.747	0.557	0.600	0.812	0.738
Compare2Score@32	13.93	0.885	0.770	0.817	0.719	0.531	0.572	0.760	0.703
ME-IQA@32	14.32	0.919	0.809	0.860	0.783	0.588	0.619	0.815	0.751

across diverse benchmarks. Three patterns recur: (i) gains hold across distortion regimes (authentic, AI-generated, and synthetic), suggesting that the hybrid memory serves as a cross-domain scaffold rather than overfitting to any single corpus; (ii) improvements are larger on synthetic datasets (*e.g.*, KADID, PIPAL), indicating that injecting local ordinal evidence sharpens sensitivity to subtle quality differences; and (iii) SRCC gains typically match or exceed PLCC gains, implying that ME-IQA primarily repairs ordinal structure without destabilizing the baseline’s global scale. These observations align with our design: reasoning-aware retrieval provides a fine-grained quality neighborhood, and re-ranking turns the VLM from a coarse scorer into a reliable ranker.

Comparison with existing non-reasoning methods. Table 1 compares ME-IQA with representative non-reasoning NR-IQA models, including: NIQE [30], NIMA [40], MUSIQ [18], CLIP-IQA+ [43], MANIQA [64], Q-Align [57], DeQA-Score [68], and Compare2Score [80] (mPLUG-Owl2 [67] as the backbone). Overall, non-reasoning methods exhibit a clear performance gap under our cross-




Fig. 3: Evidence and mitigation of *discrete collapse* across four IQA benchmarks. Histograms compare predicted score distributions from a no-memory VLM baseline VisualQuality-R1 (blue) and ME-IQA (green) against ground-truth MOS (orange) on (a) SPAQ, (b) AGIQA, (c) KADID, and (d) CSIQ. The baseline exhibits pronounced spikes at a few discrete levels, indicative of *discrete collapse*.

Table 3: Distributional collapse analysis across datasets. All metrics are computed with shared bins per dataset.

Model	SPAQ			AGIQA			KADID			CSIQ		
	JS↓	Entropy↑	Eff.Bins↑	JS↓	Entropy↑	Eff.Bins↑	JS↓	Entropy↑	Eff.Bins↑	JS↓	Entropy↑	Eff.Bins↑
<i>Q-Inight</i>												
Baseline	0.49	2.62	9.10	0.47	2.94	15.50	0.41	2.98	16.80	0.46	2.72	12.40
ME-IQA	0.10	4.23	58.40	0.06	4.36	79.10	0.06	4.31	73.60	0.10	4.33	70.10
<i>VisualQuality-R1</i>												
Baseline	0.46	2.72	10.00	0.46	3.01	16.14	0.39	3.08	17.73	0.44	2.83	13.09
ME-IQA	0.08	4.36	61.71	0.05	4.49	84.20	0.06	4.44	77.50	0.09	4.46	73.95

domain evaluation protocol, with strong models such as CLIP-IQA+ and Q-Align far behind. Compare2Score is the strongest non-reasoning competitor, benefiting from comparative prompting, yet it still falls short of ME-IQA. These results indicate that (i) reasoning-induced VLMs provide richer perceptual evidence than non-reasoning scorers, and (ii) ME-IQA effectively converts such evidence into stable, dense scalar predictions at test time.

Test-time scaling alternatives. Table 2 contrasts ME-IQA with common test-time scaling strategies, majority voting or mean aggregation over multiple samples (Maj@64, Mean@64), and Compare2Score@32 [80], under comparable compute budgets, that is 32 pairwise comparisons for ME-IQA/Compare2Score vs. 64 samples for Maj/Mean. We observe that: (i) ME-IQA matches or surpasses these baselines on nearly all datasets and backbones, yielding the highest WAVG SRCC/PLCC; (ii) relative to Compare2Score, our adaptive neighborhood retrieval and hybrid memory reduce anchor sensitivity, improving generalization in the presence of shift; (iii) gaps widen on difficult, fine-grained regimes (KADID, PIPAL, CSIQ), where rank-based refinement is most impactful; and (iv) while Mean/Maj remain competitive on authentic sets (SPAQ, LIVEW), ME-IQA still edges them out, offering a better efficiency-effectiveness trade-off at the same FLOP budget. We also report end-to-end latency (sec/img) on 1,000 images at 1K resolution using an NVIDIA A100 GPU with batch size 128. As shown in Table 2, ME-IQA@32 has a cost comparable to Compare2Score@32,

Table 4: Order robustness on four IQA benchmarks with 5 permutations.

Dataset	PLCC	SRCC
SPAQ	0.913±0.004	0.897±0.005
AGIQA	0.849±0.007	0.789±0.007
KADID	0.737±0.007	0.754±0.008
CSIQ	0.783±0.006	0.734±0.006

Table 5: Effects of Neighborhood Size (K) on WAVG performance across seven IQA benchmarks. Relative cost is normalized to $K = 32$ (*i.e.*, Rel. Cost = $K/32$).

K	Rel. Cost	PLCC	SRCC
0	0	0.698	0.661
8	0.25	0.705	0.673
16	0.50	0.714	0.689
32†	1.00	0.726	0.696
64	2.00	0.732	0.703
128	4.00	0.734	0.707
256	8.00	0.735	0.706

yet yields a clear accuracy gain. Moreover, ME-IQA is $2.4\times$ faster and more accurate than Mean@64/Maj@64.

Discrete collapse analysis. As shown in Fig. 3, the baseline VisualQuality-R1 exhibits an obvious *discrete collapse*: the score distributions concentrate at a few discrete levels. In comparison, ME-IQA redistributes probability mass from the spikes into the inter-level valleys, leading to denser, MOS-like profiles. To quantify this phenomenon, we report distribution statistics computed under the same binning scheme for MOS and model outputs (Table 3). First, the Jensen-Shannon divergence [32] (JS) between the prediction and MOS histograms decreases (lower is better). Second, the entropy of the prediction histogram increases, and so does its entropy-derived effective number of bins (higher is better), jointly indicating broader support and closer alignment to MOS. We attribute these distributional gains to reasoning-aware neighborhood retrieval that better aligns with distortion semantics and the VLM-as-comparator that supplies informative pairwise preferences, which converts local ordinal evidence into stable, fine-grained scalar predictions. Together, these components mitigate discrete collapse while improving human-consistent calibration of the score distribution.

4.3 Ablations

Unless otherwise noted, ablations use VisualQuality-R1 as the backbone, our default retrieval budget of $K=32$ with $K_A=16/K_C=16$, and the standard prompts and mapping described in Sec. 3. We report PLCC/SRCC and WAVG (where applicable) across datasets.

Order robustness. To gauge sensitivity to stream order in the online setting, we evaluate ME-IQA under five random permutations of the same test split on four benchmarks (SPAQ, AGIQA, KADID, CSIQ). For each dataset we report mean

Table 6: Efficiency comparison. DS- L denotes direct scoring with average thinking length L . Off. is batched offline latency, while On. and Mem. denote single-query online latency and peak memory.

Method	Off. ↓	On. ↓	Mem. ↓	PLCC	SRCC
DS- $L=100$	0.51	3.12	17	0.698	0.661
DS- $L=300$	1.03	8.07	22	0.701	0.663
DS- $L=500$	1.64	13.23	28	0.689	0.649
ME-IQA@8	3.07	6.32	27	0.705	0.673
ME-IQA@32	14.30	16.24	60	0.726	0.696

Table 7: Scoring mechanism comparison. PA is pairwise ordering accuracy on 1000 same-content KADID pairs with MOS gap $0.1 < \Delta < 0.25$.

Method	PLCC	SRCC	PA ↑
DS-100	0.698	0.661	57%
TokenProb-Avg.	0.682	0.647	63%
RandPair-Thurstone@32	0.691	0.664	67%
ME-IQA@32	0.726	0.696	75%

$\pm$ std over permutations (Table 4). Fluctuations are small and markedly below the gains that ME-IQA delivers over its baselines, indicating robust performance under varying arrival orders. We attribute this stability to (i) GT-stratified retrieval from AM, which supplies a consistent global scaffold across the quality range, and (ii) the weak quadratic prior that regularizes refinement around the initial scale.

Neighborhood size and efficiency. We vary the neighborhood size K and normalize relative cost to $K=32$ (Table 5). Accuracy improves from $K=0$ and saturates after $K=128$, with $K=32$ capturing most gains. Table 6 further shows that longer direct reasoning is not an effective substitute: DS- $L=300$ gives tiny gains, DS- $L=500$ is slower and less accurate than DS- $L=100$, while ME-IQA@8 improves over all direct baselines with comparable online latency to DS- $L=300$ and ME-IQA@32 gives the best accuracy.

Scoring mechanism. Table 7 separates ME-IQA from two alternative densification strategies. TokenProb-Avg raises PA but lowers correlations, and RandPair-Thurstone@32 uses the same comparison budget but remains weaker. ME-IQA reaches 75% PA and the best correlations, indicating that retrieved local ordinal evidence is more useful than token-probability averaging or random pairwise aggregation.

Re-ranking vs. re-scoring. We compare two inference schemes under the same retrieval budget on KADID (Table 11). Re-scoring performs kernel regression over retrieved neighbor scores using similarity weights and therefore only interpolates within the existing calibrated score scale. This strategy improves over the direct baseline but remains substantially behind re-ranking. In contrast, ME-IQA elicits pairwise preference probabilities and fits the query score through Thurstone fusion, which introduces explicit local ordinal constraints. The gap shows that correcting *discrete collapse* requires more than smoothing neighbor scores:

Table 8: Compact hyperparameter sensitivity. Extra denotes the reflection rate for ϵ .

Factor	Setting	PLCC	SRCC	Extra ↓
$K_A:K_C$	32:0	0.714	0.682	—
	24:8	0.719	0.685	—
	16:16	0.726	0.696	—
	8:24	0.724	0.692	—
	0:32	0.717	0.685	—
ϵ	0.25	0.727	0.696	64%
	0.50	0.725	0.698	35%
	0.75	0.726	0.696	31%
	1.00	0.721	0.694	19%
	1.25	0.717	0.691	17%

Table 9: Ablations on Retrieval Embeddings. We compare (i) Random Retrieval as a baseline, (ii) Image as Embeddings, and (iii) Reasoning as Embeddings. We further include (iv) Reasoning as Embeddings* with an additional consistency reward during the RL finetuning. All runs use $K = 32$ and report WAVG over seven datasets (backbone: VisualQuality-R1).

Retrieval Strategy	PLCC	SRCC
Random Retrieval	0.703	0.670
Image as Embeddings	0.706	0.677
Reasoning as Embeddings	0.726	0.696
Reasoning as Embeddings*	0.737	0.710

pairwise comparisons provide richer evidence about fine-grained relative quality and yield stronger rank alignment.

Hybrid memory and sensitivity. Table 8 shows that a balanced AM/CM split performs best, while nearby ratios remain close. The reflection gate mainly controls update frequency: performance is stable for $\epsilon=0.25\text{--}0.75$, while the reflection rate drops from 64% to 31%. We also vary the AM source among KONIQ, KADID, and SPAQ and test on held-out out-of-distribution datasets; the largest and average PLCC spreads are only 0.019 and 0.010, indicating stable generalization and no obvious source bias.

Causality and cold start. The hybrid design also addresses the cold-start issue of online IQA. At the beginning of a stream, CM is empty, so AM provides broad quality coverage and anchors the score scale before any target-domain case has been observed. As queries arrive, only past processed samples can enter CM, and the reflection gate filters updates before they are stored. Thus ME-IQA never uses future test samples or target MOS, preserving the causal online protocol. The five-parameter mapping used for AM is only a scale-alignment step between VLM scores, anchor MOS, and Thurstone fusion, rather than supervision from the target test domain. In practice, AM supplies stable global calibration while CM gradually specializes retrieval to the observed stream.

Retrieval keys: reasoning vs. visuals. Under a fixed budget, we compare retrieval keys for neighborhood construction (Table 9). Random Retrieval offers only modest gains; Image-as-Embeddings improves slightly, suggesting limited alignment between raw visual proximity and perceived quality. In contrast,

Table 10: Comparison under training/testing regimes in Compare2Score [80].

Method	LIVE	CSIQ	KADID	BID	CLIVE	KONIQ
<i>mPLUG-Owl2</i>						
Baseline	0.902	0.899	0.903	0.887	0.861	0.910
Compare2Score	0.970	0.932	0.935	0.947	0.917	0.945
ME-IQA	0.973	0.936	0.943	0.939	0.925	0.947
<i>VisualQuality-R1</i>						
Baseline	0.956	0.910	0.909	0.915	0.903	0.921
Compare2Score	0.953	0.914	0.900	0.912	0.890	0.922
ME-IQA	0.975	0.929	0.936	0.943	0.915	0.955

Table 11: Re-ranking vs. re-scoring under the same retrieval budget ($K = 32$). Metrics are computed on KADID with VisualQuality-R1.

Method	PLCC	SRCC
Re-scoring	0.707	0.704
Re-ranking	0.741	0.753

Reasoning-as-Embeddings substantially improves retrieval fidelity and downstream re-ranking, supporting our premise that reasoning-grounded representations capture distortion semantics and ordinal relations. Adding a consistency reward during RL fine-tuning to stabilize the reasoning space yields the strongest results (Reasoning-as-Embeddings*; see supplementary), indicating that intra-task consistency further improves retrieval geometry.

Resolution robustness. ME-IQA does not rely on a fixed resizing rule. We process images at native resolution whenever feasible and otherwise follow the backbone VLM’s default resizing/tokenization policy. Table 12 shows that performance on SPAQ is stable from 0.5K to 4K resolution, with PLCC/SRCC varying only mildly. This suggests that the memory-enhanced comparison stage is not tied to a particular pixel scale; the retrieved reasoning summaries and pairwise preferences remain effective as long as the backbone preserves the relevant quality cues.

Robustness to training/testing regimes. Table 10 tests ME-IQA under the Compare2Score [80] protocol with a non-reasoning mPLUG-Owl2 backbone and a stronger VisualQuality-R1 backbone trained on the same six-dataset mixture. ME-IQA consistently improves both baselines and remains competitive with, or superior to, Compare2Score; the smaller mPLUG margin likely comes from using less quality-aligned vision-encoder retrieval keys.

Prior weight in Thurstone fusion. We sweep the quadratic prior weight λ in the Case V objective on KADID (Table 13). Removing the prior degrades correlations, suggesting that purely pairwise evidence can be noisy and prone to over-correction. A very weak prior stabilizes optimization and consistently improves performance. Results peak around $\lambda \in [1e-2, 5e-2]$, which anchors refinement to the initial scale while preserving informative ordinal signals. Larger λ (e.g., 1e-1) slightly harms performance, implying under-adjustment when the prior dominates.

Table 12: Resolution ablation on SPAQ. ME-IQA remains stable across input resolutions.

Resolution	PLCC	SRCC
0.5K	0.912	0.899
1K	0.917	0.906
2K	0.915	0.908
4K	0.921	0.906

Table 13: Effects of prior hyperparameter λ in Thurstone fusion. All runs use $K = 32$ and report results on the KADID dataset (backbone: VisualQuality-R1).

Metric	λ				
	0	1e-3	1e-2†	5e-2	1e-1
PLCC	0.724	0.733	0.741	0.740	0.737
SRCC	0.727	0.739	0.753	0.755	0.749

Failure cases. By aggregating neighbors and using the initial score as a weak prior, ME-IQA rarely overturns a correct direct order: on the 1000 KADID pairs in Table 7, this happens only 5 times. Remaining failures often involve subtle impulse noise or JPEG compression artifacts, suggesting a shared perceptual limitation of the underlying VLM.

5 Conclusion

We presented ME-IQA, a plug-and-play test-time memory-enhanced re-ranking framework that mitigates *discrete collapse* in reasoning-induced VLMs. By retrieving aligned exemplars, eliciting pairwise preferences, and applying Thurstone fusion, ME-IQA produces denser calibrated scores without retraining. Across seven benchmarks and multiple backbones, it improves over reasoning and non-reasoning IQA baselines and test-time scaling alternatives. Ablations show that gains come from image-specific ordinal evidence rather than longer reasoning, token-probability averaging, random comparisons, or neighbor-score interpolation, while bounded AM/CM keeps online deployment causal and practical.

Acknowledgements

We thank Professor Kede Ma for his helpful suggestions and discussions. Yabin Zhang is the project leader of this work.

References

1. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2.5-VL technical report. arXiv preprint arXiv:2502.13923 (2025)

2. Bosse, S., Maniry, D., Müller, K.R., Wiegand, T., Samek, W.: Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Transactions on Image Processing* **27**(1), 206–219 (2017)
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: *Advances in Neural Information Processing Systems*. pp. 1877–1901 (2020)
4. Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., Hullender, G.: Learning to rank using gradient descent. In: *International Conference on Machine Learning*. pp. 89–96 (2005)
5. Chen, P., Li, L., Huang, Y., Tan, F., Chen, W.: Qoe evaluation for live broadcasting video. In: *IEEE International Conference on Image Processing*. pp. 454–458 (2019)
6. Chhikara, P., Khant, D., Aryan, S., Singh, T., Yadav, D.: Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413* (2025)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 4171–4186 (2019)
8. Fang, Y., Zhu, H., Zeng, Y., Ma, K., Wang, Z.: Perceptual quality assessment of smartphone photography. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3677–3686 (2020)
9. Gao, F., Tao, D., Gao, X., Li, X.: Learning to rank for blind image quality assessment. *IEEE Transactions on Neural Networks and Learning Systems* **26**(10), 2275–2290 (2015)
10. Ghadiyaram, D., Bovik, A.C.: Massive online crowdsourced study of subjective and objective picture quality. *IEEE Transactions on Image Processing* **25**(1), 372–387 (2016)
11. Gu, J., Cai, H., Chen, H., Ye, X., Ren, J.S., Dong, C.: PIPAL: A large-scale image quality assessment dataset for perceptual image restoration. In: *European Conference on Computer Vision*. pp. 633–651 (2020)
12. Guu, K., Lee, K., Tung, Z., Pasupat, P., Chang, M.: Retrieval augmented language model pre-training. In: *International Conference on Machine Learning*. pp. 3929–3938 (2020)
13. Helson, H.: *Adaptation-level Theory: An Experimental and Systematic Approach to Behavior*. Harper & Row (1964)
14. Hosu, V., Lin, H., Sziranyi, T., Saupe, D.: KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing* **29**, 4041–4056 (2020)
15. Hu, Y., Wang, Y., McAuley, J.: Evaluating memory in llm agents via incremental multi-turn interactions. *arXiv preprint arXiv:2507.05257* (2025)
16. Hu, Z., Iscen, A., Sun, C., Wang, Z., Chang, K.W., Sun, Y., Schmid, C., Ross, D.A., Fathi, A.: Reveal: Retrieval-augmented visual-language pre-training with multi-source multimodal knowledge memory. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23369–23379 (2023)
17. Kang, L., Ye, P., Li, Y., Doermann, D.: Convolutional neural networks for no-reference image quality assessment. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1733–1740 (2014)

18. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: MUSIQ: Multi-scale image quality Transformer. In: IEEE/CVF International Conference on Computer Vision. pp. 5148–5157 (2021)
19. Larson, E.C., Chandler, D.M.: Most apparent distortion: Full-reference image quality assessment and the role of strategy. *Journal of Electronic Imaging* **19**(1), 1–21 (2010)
20. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* **33**, 9459–9474 (2020)
21. Li, C., Bovik, A.C., Wu, X.: Blind image quality assessment using a general regression neural network. *IEEE Transactions on Neural Networks* **22**(5), 793–799 (2011)
22. Li, C., Zhang, Z., Wu, H., Sun, W., Min, X., Liu, X., Zhai, G., Lin, W.: 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)
23. Li, W., Zhang, X., Zhao, S., Zhang, Y., Li, J., Zhang, L., Zhang, J.: Q-Insight: Understanding image quality via visual reinforcement learning. *arXiv preprint arXiv:2503.22679* (2025)
24. Lin, H., Hosu, V., Saupe, D.: KADID-10K: A large-scale artificially distorted IQA database. In: IEEE International Conference on Quality of Multimedia Experience. pp. 1–3 (2019)
25. Liu, H., Klomp, N., Heynderickx, I.: A no-reference metric for perceived ringing artifacts in images. *IEEE Transactions on Circuits and Systems for Video Technology* **20**(4), 529–539 (2009)
26. Ma, K., Liu, W., Liu, T., Wang, Z., Tao, D.: dipIQ: Blind image quality assessment by learning-to-rank discriminable image pairs. *IEEE Transactions on Image Processing* **26**(8), 3951–3964 (2017)
27. Ma, K., Liu, W., Zhang, K., Duanmu, Z., Wang, Z., Zuo, W.: End-to-end blind image quality assessment using deep neural networks. *IEEE Transactions on Image Processing* **27**(3), 1202–1213 (2017)
28. Maharana, A., Lee, D.H., Tulyakov, S., Bansal, M., Barbieri, F., Fang, Y.: Evaluating very long-term conversational memory of llm agents. In: Annual Meeting of the Association for Computational Linguistics. pp. 13851–13870 (2024)
29. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing* **21**(12), 4695–4708 (2012)
30. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters* **20**(3), 209–212 (2012)
31. Murtagh, F., Contreras, P.: Algorithms for hierarchical clustering: An overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **2**(1), 86–97 (2012)
32. Nielsen, F.: On a variational definition for the jensen-shannon symmetrization of distances based on the information radius. *Entropy* **23**(4), 464 (2021)
33. Ouyang, S., Yan, J., Hsu, I., Chen, Y., Jiang, K., Wang, Z., Han, R., Le, L.T., Daruki, S., Tang, X., et al.: Reasoningbank: Scaling agent self-evolving with reasoning memory. *arXiv preprint arXiv:2509.25140* (2025)
34. Packer, C., Wooders, S., Lin, K., Fang, V., Patil, S.G., Stoica, I., Gonzalez, J.E.: MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560* (2023)

35. Ponomarenko, N., Jin, L., Ieremeiev, O., Lukin, V., Egiazarian, K., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., Kuo, C.C.J.: Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication* **30**, 57–77 (2015)
36. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., Guo, D.: DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024)
37. Sheikh, H.R., Sabir, M.F., Bovik, A.C.: A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing* **15**(11), 3440–3451 (2006)
38. Stewart, N., Chater, N., Brown, G.D.: Decision by sampling. *Cognitive psychology* **53**(1), 1–26 (2006)
39. SU, H., Sun, R., Yoon, J., Yin, P., Yu, T., Arik, S.O.: Learn-by-interact: A data-centric framework for self-adaptive agents in realistic environments. In: *International Conference on Learning Representations*. pp. 1–13 (2025)
40. Talebi, H., Milanfar, P.: NIMA: Neural image assessment. *IEEE Transactions on Image Processing* **27**(8), 3998–4011 (2018)
41. Tang, X., Hu, T., Ye, M., Shao, Y., Yin, X., Ouyang, S., Zhou, W., Lu, P., Zhang, Z., Zhao, Y., Cohan, A., Gerstein, M.: Chemagent: Self-updating memories in large language models improves chemical reasoning. In: *International Conference on Learning Representations*. pp. 1–15 (2025)
42. Thurstone, L.L.: A law of comparative judgment. *Psychological Review* **34**, 273–286 (1927)
43. Wang, J., Chan, K.C., Loy, C.C.: Exploring clip for assessing the look and feel of images. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. pp. 2555–2563 (2023)
44. Wang, X., Xiong, J., Gao, H., Lin, W.: Regression-free blind image quality assessment with content-distortion consistency. *arXiv preprint arXiv:2307.09279* (2023)
45. Wang, Y., Krotov, D., Hu, Y., Gao, Y., Zhou, W., McAuley, J., Gutfreund, D., Feris, R., He, Z.: M+: Extending memoryLLM with scalable long-term memory. In: *International Conference on Machine Learning*. pp. 1–12 (2025)
46. Wang, Z., Ma, K.: Active fine-tuning from gMAD examples improves blind image quality assessment. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 4577–4590 (2021)
47. Wang, Z., Wang, H., Chen, T., Wang, Z., Ma, K.: Troubleshooting blind image quality models in the wild. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16256–16265 (2021)
48. Wang, Z., Bovik, A.C.: *Modern Image Quality Assessment*. Morgan & Claypool Publishers (2006)
49. Wang, Z., Bovik, A.C.: Reduced-and no-reference image quality assessment. *IEEE Signal Processing Magazine* **28**(6), 29–40 (2011)
50. Wang, Z., Bovik, A.C., Evan, B.L.: Blind measurement of blocking artifacts in images. In: *IEEE International Conference on Image Processing*. pp. 981–984 (2000)
51. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
52. Wang, Z., Simoncelli, E.P.: Local phase coherence and the perception of blur. In: *Advances in Neural Information Processing Systems*. pp. 1435–1442 (2003)
53. Wang, Z.Z., Gandhi, A., Neubig, G., Fried, D.: Inducing programmatic skills for agentic tasks. *arXiv preprint arXiv:2504.06821* (2025)

54. Wen, W., Zhi, T., Fan, K., Li, Y., Peng, X., Zhang, Y., Liao, Y., Li, J., Zhang, L.: Self-evolving Vision-Language Models for image quality assessment via voting and ranking. In: International Conference on Learning Representations. pp. 1–14 (2026)
55. Wu, C., Tam, Z.R., Lin, C., Chen, Y., Lee, H.: Streambench: Towards benchmarking continuous improvement of language agents. In: Advances in Neural Information Processing Systems. pp. 107039–107063 (2024)
56. Wu, D., Wang, H., Yu, W., Zhang, Y., Chang, K.W., Yu, D.: LongMemEval: Benchmarking chat assistants on long-term interactive memory. In: International Conference on Learning Representations. pp. 1–16 (2025)
57. Wu, H., Zhang, Z., Zhang, W., Chen, C., Liao, L., Li, C., Gao, Y., Wang, A., Zhang, E., Sun, W., Yan, Q., Min, X., Zhai, G., Lin, W.: Q-Align: Teaching LMMs for visual scoring via discrete text-defined levels. In: International Conference on Machine Learning. pp. 54015–54029 (2024)
58. Wu, Q., Li, H., Meng, F., Ngan, K.N., Luo, B., Huang, C., Zeng, B.: Blind image quality assessment based on multichannel feature fusion and label transfer. *IEEE Transactions on Circuits and Systems for Video Technology* **26**(3), 425–440 (2015)
59. Wu, Q., Li, H., Ngan, K.N., Ma, K.: Blind image quality assessment using local consistency aware retriever and uncertainty aware evaluator. *IEEE Transactions on Circuits and Systems for Video Technology* **28**(9), 2078–2089 (2017)
60. Wu, T., Ma, K., Liang, J., Yang, Y., Zhang, L.: A comprehensive study of multi-modal large language models for image quality assessment. In: European Conference on Computer Vision. pp. 143–160 (2024)
61. Wu, T., Zou, J., Liang, J., Zhang, L., Ma, K.: VisualQuality-R1: Reasoning-induced image quality assessment via reinforcement learning to rank. *arXiv preprint arXiv:2505.14460* (2025)
62. Xu, W., Mei, K., Gao, H., Tan, J., Liang, Z., Zhang, Y.: A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110* (2025)
63. Xue, W., Zhang, L., Mou, X.: Learning without human scores for blind image quality assessment. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 995–1002 (2013)
64. Yang, S., Wu, T., Shi, S., Lao, S., Gong, Y., Cao, M., Wang, J., Yang, Y.: MANIQA: Multi-dimension attention network for no-reference image quality assessment. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 1191–1200 (2022)
65. Yasunaga, M., Aghajanyan, A., Shi, W., James, R., Leskovec, J., Liang, P., Lewis, M., Zettlemoyer, L., Yih, W.t.: Retrieval-augmented multimodal language modeling. In: International Conference on Machine Learning (2022)
66. Ye, P., Doermann, D.: No-reference image quality assessment using visual codebooks. *IEEE Transactions on Image Processing* **21**(7), 3129–3138 (2012)
67. Ye, Q., Xu, H., Ye, J., Yan, M., Hu, A., Liu, H., Qian, Q., Zhang, J., Huang, F.: mPLUG-Owl2: Revolutionizing multi-modal large language model with modality collaboration. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 13040–13051 (2024)
68. You, Z., Cai, X., Gu, J., Xue, T., Dong, C.: Teaching large language models to regress accurate image quality scores using score distribution. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14483–14494 (2025)
69. You, Z., Li, Z., Gu, J., Yin, Z., Xue, T., Dong, C.: Depicting beyond scores: Advancing image quality assessment through multi-modal language models. In: European Conference on Computer Vision. pp. 259–276 (2024)

70. Zhang, L., Hosseini, A., Bansal, H., Kazemi, M., Kumar, A., Agarwal, R.: Generative verifiers: Reward modeling as next-token prediction. arXiv preprint arXiv:2408.15240 (2024)
71. Zhang, W., Li, D., Ma, C., Zhai, G., Yang, X., Ma, K.: Continual learning for blind image quality assessment. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(3), 2864–2878 (2022)
72. Zhang, W., Ma, K., Zhai, G., Yang, X.: Uncertainty-aware blind image quality assessment in the laboratory and wild. *IEEE Transactions on Image Processing* **30**, 3474–3486 (2021)
73. Zhang, W., Zhai, G., Wei, Y., Yang, X., Ma, K.: Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14071–14081 (2023)
74. Zhang, X.F., Beauchamp, N., Wang, L.: Prime: Large language model personalization with cognitive dual-memory and personalized thought process. In: *Conference on Empirical Methods in Natural Language Processing*. pp. 33695–33724 (2025)
75. Zhang, Y., Li, M., Long, D., Zhang, X., Lin, H., Yang, B., Xie, P., Yang, A., Liu, D., Lin, J., Huang, F., Zhou, J.: Qwen3 embedding: Advancing text embedding and reranking through foundation models. arXiv preprint arXiv:2506.05176 (2025)
76. Zhang, Z., Dai, Q., Bo, X., Ma, C., Li, R., Chen, X., Zhu, J., Dong, Z., Wen, J.R.: A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems* **43**(6), 1–47 (2025)
77. Zhao, A., Huang, D., Xu, Q., Lin, M., Liu, Y., Huang, G.: Expel: LLM agents are experiential learners. In: *AAAI Conference on Artificial Intelligence*. pp. 19632–19642 (2024)
78. Zheng, L., Wang, R., Wang, X., An, B.: Synapse: Trajectory-as-exemplar prompting with memory for computer control. In: *International Conference on Learning Representations*. pp. 1–13 (2024)
79. Zhong, W., Guo, L., Gao, Q., Ye, H., Wang, Y.: MemoryBank: Enhancing large language models with long-term memory. In: *AAAI Conference on Artificial Intelligence*. pp. 19724–19731 (2024)
80. Zhu, H., Wu, H., Li, Y., Zhang, Z., Chen, B., Zhu, L., Fang, Y., Zhai, G., Lin, W., Wang, S.: Adaptive image quality assessment via teaching large multimodal model to compare. In: *Advances in Neural Information Processing Systems*. pp. 32611–32629 (2024)