Toward Unified Multimodal Feature Quality Assessment via Cross-Modal Feature Alignment

Anonymous authors
Paper under double-blind review

Abstract

For decades, perceptual quality assessment for image, text, and audio relied on three separate toolchains with incompatible metrics and almost no overlap. Yet the common ground is deeper than it appears. If we can project features from all three modalities into the same space, quality assessment becomes a single regression task. To test this, we built a synthetic distortion benchmark covering all three modalities with consistent labels, all traced back to shared visual references. We designed a lightweight pipeline that first identifies the modality and then applies a dedicated regressor, using a simple mechanism to bridge the feature spaces. The result is quality assessment across modalities without heavy custom engineering. Our experiments find a consistent difficulty ordering, clear differences in how models regularize, and a hard limit on how well standard audio features transfer to unseen distortions.

1 Introduction

For decades, image, text, and audio quality assessment evolved independently. Image quality assessment (IQA) matured around LIVE and TID2013 Sheikh et al. (2006); Ponomarenko et al. (2015) with metrics Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) Wang et al. (2004), and Learned Perceptual Image Patch Similarity (LPIPS) Zhang et al. (2018). Text evaluation grew around Bilingual Evaluation Understudy (BLEU) Papineni et al. (2002), Recall-Oriented Understudy for Gisting Evaluation (ROUGE) Lin (2004), and BERTScore Zhang et al. (2020). Audio assessment relied on Perceptual Evaluation of Speech Quality (PESQ) Rix et al. (2001) and deep predictors DNSMOS Reddy et al. (2021) and NISQA Mittag et al. (2021). Each community uses different benchmarks and protocols. No framework compares which modality is hardest to assess or whether one pipeline handles all three. The feature coding community addresses a parallel need for standardized evaluation Gao et al. (2024), with benchmarks for cross-architecture quality comparison.

We discovered a straightforward answer. Once features from all three modalities are projected into the same vector space, quality assessment collapses into a single regression problem. This parallels recent work on cross-architecture feature coding Gao et al. (2025), where distribution alignment enables universal compression without per-architecture tuning. The modality no longer matters because the feature representation does the work. We put this to the test with a simple pipeline. Images and text pass through Contrastive Language-Image Pre-training (CLIP) Vision Transformer (ViT)-B/32 Radford et al. (2021), producing L2-normalized 512-dimensional vectors. Audio goes through Mel-Frequency Cepstral Coefficients (MFCC) extraction followed by a fixed random projection with zero learned parameters, landing in the same 512-D space as CLIP embeddings. A lightweight two-stage architecture first identifies the source modality and then applies a dedicated regressor. Figure 1 sketches the design.

We evaluate on a synthetic benchmark of 50 distortion types spanning all three modalities, derived from shared visual references. Three experiments follow. Eleven regression algorithms reveal a clear difficulty gradient, with SVR-FQA (Feature Quality Assessment) and Extra Trees being strongest on SRCC. Second, an ablation study confirms the pipeline is robust to preprocessing choices. Third, a cross-type generalization test exposes a critical gap: audio features collapse on unseen distortions, showing that MFCC encodes distortion-specific rather than transferable quality patterns.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

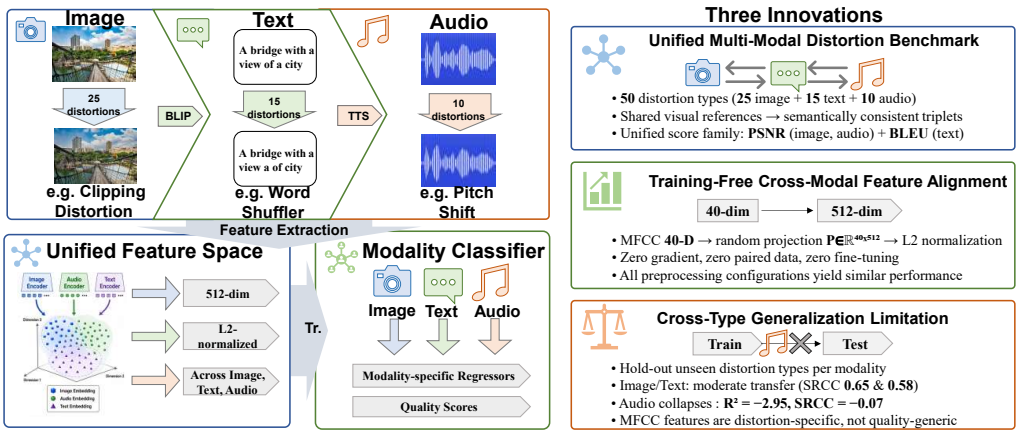

Figure 1: Framework overview. Images and text use CLIP; audio uses MFCCs with a random projection into the shared 512-D space. A modality classifier selects the corresponding regressor.

Our contributions are (1) a unified benchmark for controlled image, text, and audio comparison; (2) a training-free audio-to-CLIP projection without paired data; and (3) SVR-FQA, a modality-specific SVR tuning method that outperforms default SVR and tree ensembles.

2 Related Work

IQA is the most mature sub-field, with full-reference (PSNR, SSIM Wang et al. (2004), LPIPS Zhang et al. (2018)) and no-reference models (BRISQUE Mittal et al. (2012), NIQE Mittal et al. (2013)), including vision-language approaches (CLIP-IQA Wang et al. (2023)). Text evaluation centers on n-gram overlap (BLEU Papineni et al. (2002), ROUGE Lin (2004)) and learned metrics like BERTScore Zhang et al. (2020). Audio quality relies on signal-level (PESQ Rix et al. (2001)) and learned metrics (DNSMOS Reddy et al. (2021), NISQA Mittag et al. (2021)), but these models operate on raw waveforms without producing embeddings aligned to visual-textual spaces.

CLIP Radford et al. (2021) and BLIP (Bootstrapping Language-Image Pre-training) Li et al. (2022) pioneered vision-language alignment; extending to audio typically requires costly contrastive pre-training such as CLAP (Contrastive Language-Audio Pre-training) Wu et al. (2023) and AudioCLIP Guzhov et al. (2022). Our randomized projection offers a training-free alternative. To our knowledge, no existing benchmark covers image, text, and audio quality simultaneously with shared annotations; each domain uses incompatible quality scales.

3 Dataset Construction

3.1 Original Data

We sample 160 reference images from standard vision datasets. For each reference, a BLIP captioning model Li et al. (2022) automatically generates a descriptive English caption, and a text-to-speech (TTS) engine synthesizes a corresponding spoken audio clip. The resulting image, caption, and audio form a semantically consistent cross-modal triplet sharing the same visual origin.

3.2 Distorted Data

Each modality is independently distorted, yielding labeled quality-score pairs across all modalities. The benchmark contains 50 distortion types in total.

Image (25 types): Gaussian blur, Gaussian additive noise, JPEG compression, motion blur, spatial downscaling, color jitter, ISO noise, sharpening, fog, snow, rain, elastic transform, optical/grid distortion, hue-saturation-value shift, posterization, solarization, histogram equalization, multiplica-

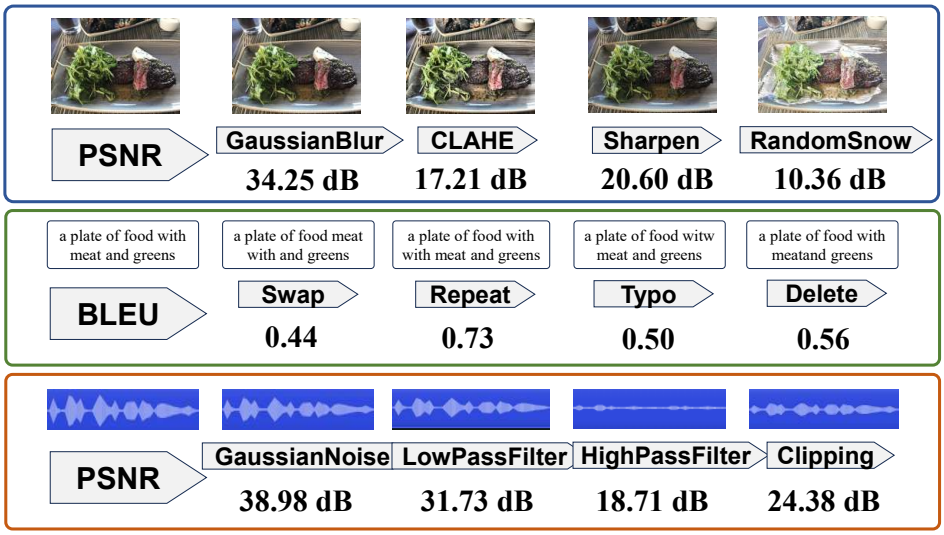

Figure 2: Representative benchmark examples: one reference and four distortions per modality with quality scores.

tive noise, Fancy PCA noise, brightness-contrast adjustment, RGB shift, channel shuffle, CLAHE (Contrast Limited Adaptive Histogram Equalization), and image inversion.

Text (15 types): Adjacent word swap, random deletion, word repetition, letter typo, character deletion/insertion, sentence shuffle, lowercasing, random capitalization, homophone replacement, reversal, punctuation noise/removal, abbreviation substitution, and contraction expansion.

Audio (10 types): Additive white Gaussian noise, time stretching, pitch shifting, low-pass/high-pass filtering, gain scaling, polarity inversion, hard clipping, bit-depth reduction, and tanh saturation.

One distorted variant per distortion type per reference in the benchmark yields a total of 3,966 image samples, 2,400 text samples, and 1,600 audio samples, with some types skipped for unsuitable references. Representative examples are shown in Figure 2.

3.3 Quality Labels

Image and audio samples use PSNR (Peak Signal-to-Noise Ratio) as the quality score, each computed on a decibel scale against the original reference. For audio, PSNR is computed over the time-aligned waveform segment and capped at 100 dB. Text samples use BLEU-4 (Bilingual Evaluation Understudy) against the reference caption. All distorted variants share the identical visual scene origin, which is strictly enforced by the generation pipeline. These modality-specific labels preserve the native measurement conventions while keeping the reference identity and semantic content fixed across each triplet.

4 Method

The framework has three stages: feature alignment, classify-then-regress, and cross-type generalization (Figures 1 and 4).

4.1 Unified Feature Extraction and Randomized Projection Alignment

Images and text are processed through CLIP's ViT-B/32 encoders, producing L2-normalized 512-D embedding vectors. For audio, we extract 20 Mel-Frequency Cepstral Coefficients (MFCC) via librosa McFee et al. (2015), computing mean and standard deviation over the time axis to yield a 40-D vector $f_a \in \mathbb{R}^{40}$. A random projection matrix maps audio into the CLIP space:

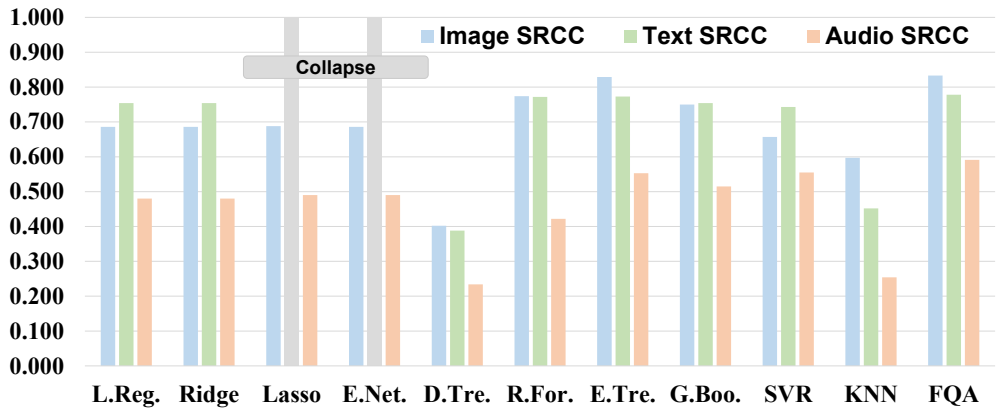

Figure 3: SRCC across regressors and modalities. SVR-FQA leads; audio is consistently hardest, while Lasso and ElasticNet collapse on text.

$$\tilde{f}_a = \frac{f_a P}{\|f_a P\|_2}, \qquad P \in \mathbb{R}^{40\times 512}$$

Here P is drawn from a Gaussian distribution and scaled by $\sqrt{512/40}$ (seed 42). As established by the Johnson–Lindenstrauss lemma Johnson & Lindenstrauss (1984), random Gaussian projections preserve pairwise distances; because we project into a higher-dimensional space (40→512), the resulting distortion is negligible for this benchmark. After L2 normalization, all features lie on the unit sphere, ensuring that the relative structure of the MFCC space is retained in the CLIP-aligned representation. This single matrix multiplication plus normalization requires no training, no paired audio–text data, and no GPU. Whether the resulting 512-D vectors retain sufficient quality information is tested in Section 5.2.

4.2 Classify-then-Regress Architecture

With features from all modalities aligned in 512-D space, the inference pipeline has two stages.

Stage 1: Modality classification. We train a Random Forest (RF) classifier with 100 trees on the combined feature set, using Gini impurity with trees grown to full depth. Despite L2 normalization, the classifier reliably distinguishes modalities because CLIP vision, CLIP text, and projected audio features occupy distinct angular regions of the 512-D sphere.

Stage 2: Per-modality regression. Features are first standardized per modality using scikit-learn's StandardScaler Pedregosa et al. (2011). We compare eleven regressors: Linear Regression, Ridge, Lasso, ElasticNet, Decision Tree, RF, Extra Trees, Gradient Boosting, default SVR, K-Nearest Neighbors (KNN), and SVR-FQA. RF and Extra Trees use 100 trees without a depth limit; Gradient Boosting uses 100 estimators with learning rate 0.1 and maximum depth 3. SVR defaults use an RBF kernel with $C = 1.0$ and $\epsilon = 0.1$, while KNN uses $k = 5$ with uniform Euclidean distance. SVR-FQA uses modality-specific parameters: image ($C = 100, \epsilon = 0.05, \gamma = \text{scale}$), text ($C = 10, \epsilon = 0.05, \gamma = \text{auto}$), and audio ($C = 100, \epsilon = 1.0, \gamma = \text{scale}$).

Inference procedure. Given a feature vector of unknown modality, (1) the classifier predicts the modality, (2) the per-modality scaler transforms the feature, and (3) the corresponding regressor produces the quality estimate:

$$\hat{k} = \arg\max_{k\in\{I,T,A\}} P(k \mid z), \qquad \hat{q} = \mathcal{R}_{\hat{k}}\big(\psi_{\hat{k}}(z)\big)$$

Here z is the unified feature, ψ_k the modality scaler, and $\mathcal{R}_k$ the corresponding regressor; the classifier uses the full dataset while each regressor sees its own modality.

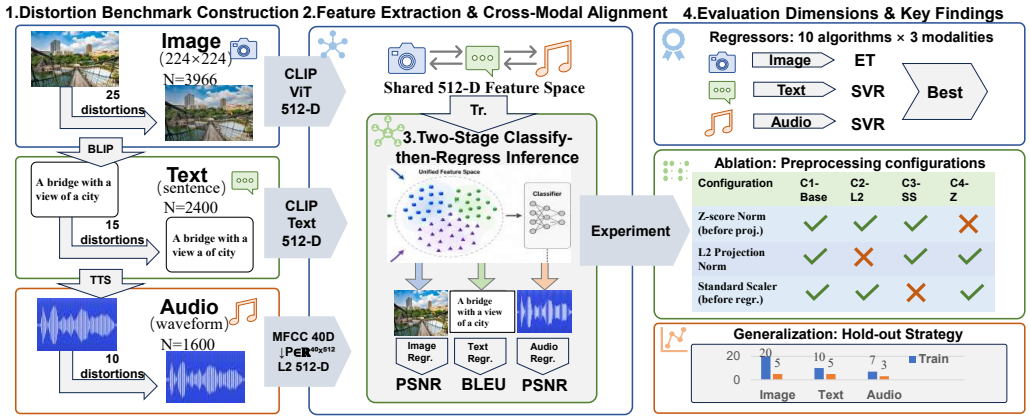

Figure 4: Pipeline from reference images to cross-modal features, classify-then-regress predictions, and the three evaluation settings.

4.3 Cross-Type Generalization

Beyond standard in-distribution evaluation, we test generalization to unseen distortion types by holding out predefined subsets during training: five image, five text, and three audio distortions. The regressor is evaluated exclusively on these unseen types, probing whether it learns generic quality patterns or merely memorizes distortion-specific signatures.

RF and SVR-FQA use the main configuration, training on seen types and testing only on held-out types, thereby providing a direct quantitative measure of cross-distortion transfer.

5 Experiments

5.1 Dataset and Experimental Setup

The benchmark comprises 160 reference images from standard datasets. For each, one distorted variant per type yields 3,966 image pairs (25 types, some skipped for unsuitable references), 2,400 text pairs (15 types), and 1,600 audio pairs (10 types). All experiments use a fixed 80/20 train-test split across the sample pool, with seed 42 for reproducibility.

CLIP feature extraction runs on a single NVIDIA GPU. MFCC extraction uses librosa. Regression runs on CPU via scikit-learn, with each model training in under one minute.

We report seven metrics: Train R^2, Test R^2, Mean Absolute Error (MAE), Root Mean Square Error (RMSE), SRCC, Pearson Linear Correlation Coefficient (PLCC), and Kendall's Tau Correlation Coefficient (KLCC) to provide a multi-faceted evaluation. Test R^2 and SRCC serve as the primary criteria, with MAE and RMSE providing complementary error-scale information.

5.2 Main Results: Regressor Comparison

We begin by comparing all eleven regressors on the full multimodal benchmark with standard 80/20 splits.

Table 1 (also Figure 3) shows that SVR-FQA leads across all three modalities in SRCC: image (SRCC= 0.833), text (SRCC= 0.778), audio (SRCC= 0.591). SRCC values establish a clear, consistent difficulty gradient: image > text > audio, with audio substantially harder than both visual and textual modalities. Relative MAE reduction over the mean predictor (Table 5) independently confirms audio as the hardest modality overall.

On audio specifically, Linear Regression achieves the best R^2 with moderate SRCC (0.480), while SVR-FQA attains the best SRCC (0.591) with competitive R^2 (0.220; Table 3). Linear models capture more variance in the approximately linear MFCC-projected space, but SVR-FQA's nonlinear

Image (I)											
Metric	Linear Reg.	Ridge	Lasso	ElasticNet	Dec. Tree	Rand. For.	Extra Trees	Grad. Boost	SVR	SVR-FQA	KNN
MAE↓	5.317	5.298	6.741	6.282	6.506	4.699	**4.358**	5.334	5.615	**3.866**	7.121
RMSE↓	7.233	7.193	8.349	7.863	9.313	6.212	5.975	6.913	7.352	**5.657**	8.957
SRCC↑	0.686	0.689	0.476	0.551	0.587	0.774	0.796	0.685	0.657	**0.833**	0.648
PLCC↑	0.653	0.655	0.488	0.569	0.546	0.757	0.776	0.668	0.647	**0.790**	0.643
KLCC↑	0.500	0.502	0.329	0.386	0.421	0.575	0.598	0.493	0.473	**0.639**	0.466
Text (T)											
Metric	Linear Reg.	Ridge	Lasso	ElasticNet	Dec. Tree	Rand. For.	Extra Trees	Grad. Boost	SVR	SVR-FQA	KNN
MAE↓	0.224	0.222	0.367	0.367	0.216	0.213	0.205	0.224	0.182	**0.153**	0.238
RMSE↓	0.305	0.303	0.413	0.413	0.338	0.294	0.288	0.293	0.272	**0.262**	0.338
SRCC↑	0.567	0.569	—	—	0.490	0.672	0.680	0.614	0.743	**0.778**	0.449
PLCC↑	0.689	0.691	—	—	0.611	0.703	0.718	0.707	0.758	**0.788**	0.604
KLCC↑	0.420	0.423	—	—	0.369	0.531	0.540	0.472	0.588	**0.631**	0.328
Audio (A)											
Metric	Linear Reg.	Ridge	Lasso	ElasticNet	Dec. Tree	Rand. For.	Extra Trees	Grad. Boost	SVR	SVR-FQA	KNN
MAE↓	8.186	8.184	9.014	8.837	9.452	7.255	**6.308**	7.174	7.236	6.439	7.578
RMSE↓	**12.03**	**12.03**	12.72	12.59	18.10	12.46	12.25	12.34	13.59	12.17	13.54
SRCC↑	0.480	0.480	0.362	0.384	0.423	0.474	0.553	0.515	0.555	**0.591**	0.398
PLCC↑	0.487	0.487	0.401	0.425	0.226	0.458	0.487	0.472	0.381	**0.519**	0.353
KLCC↑	0.335	0.335	0.245	0.262	0.362	0.369	0.453	0.394	0.409	**0.456**	0.326

Table 1: Test-set performance of regression algorithms across three modalities. SVR denotes sklearn defaults ($C = 1.0$, $\epsilon = 0.1$, γ = scale); SVR-FQA uses modality-specific parameters from grid search: image ($C = 100$, $\epsilon = 0.05$, γ = scale), text ($C = 10$, $\epsilon = 0.05$, γ = auto), audio ($C = 100$, $\epsilon = 1.0$, γ = scale). Bold denotes best per column within each modality. Lasso and ElasticNet collapse on text (undefined). All results use seed 42; multi-seed stability in Table 2. I, T, A denote Image, Text, Audio.

kernel better preserves monotonic ranking across samples. This is consistent with multi-seed results (Section 5.2.1), where Linear Regression shows lower R^2 variance than SVR-FQA. Tree ensembles outperform linear models substantially on both image and text; KNN is weakest; single Decision Tree overfits (train $R^2 = 1.0$, test $R^2 < 0$).

SVR-FQA: Feature Quality Assessment via modality-specific SVR tuning. We propose SVR-FQA (Feature Quality Assessment), a novel method that finds per-modality optimal SVR hyperparameters. A key finding is that default SVR parameters are far from optimal for feature quality assessment. Standard SVR ($C = 1.0$, $\epsilon = 0.1$, γ = scale) ranks mid-tier on image and audio, and trails SVR-FQA on text (Table 1). SVR-FQA tunes the RBF kernel's capacity (C), insensitivity margin (ϵ), and kernel coefficient (γ) separately for each modality via grid search. The SVR objective with RBF kernel is:

$$\min_{w,b} \frac{1}{2}\|w\|^2 + C\sum_i \max\big(0,\ |y_i - \langle w, \phi(z_i)\rangle - b| - \epsilon\big)$$

where $\phi(z_i)$ is the implicit RBF kernel feature map ($K(z_i, z_j) = \exp(-\gamma\|z_i - z_j\|^2)$), C controls the penalty for margin violations, and ϵ defines the tube width within which errors are not penalized. The optimal configuration differs substantially by modality. Image quality features in CLIP space require high capacity ($C = 100$) and tight margin ($\epsilon = 0.05$), reflecting dense support vector coverage of the PSNR manifold. Text quality with BLEU targets benefits from moderate capacity ($C = 10$) and data-driven kernel scale (γ = auto), indicating that CLIP text embeddings carry quality information at finer granularity than images. Audio MFCC projections demand the widest margin ($\epsilon = 1.0$) with high capacity ($C = 100$), consistent with the noisier MFCC-to-CLIP mapping. Compared to default SVR, SVR-FQA improves SRCC by +0.176 (image), +0.035 (text), and +0.036 (audio). This establishes SVR-FQA not merely as a tuned baseline but as a principled method: modality-specific hyperparameter search uncovers the inductive bias appropriate to each feature representation, yielding a single regression family that dominates across all three modalities without architectural modification.

5.2.1 Overfitting, Sparsity, and Seed Stability

Tree-based methods exhibit stark overfitting: Extra Trees achieves near-perfect train fits with substantially lower test performance, while SVR-FQA shows moderate train–test gaps (0.056–0.157; Table 3). Linear models outperform RF on audio, suggesting the MFCC-projected quality signal is approximately linear—tree splits amplify noise in a high-dimensional low-sample regime where

Mod.	Algorithm	Test R^2↑		SRCC↑	
I	**SVR-FQA**	**0.623**	**±0.025**	**0.834**	**±0.004**
	Extra Trees	0.591	±0.012	0.797	±0.006
	Rand. For.	0.547	±0.004	0.769	±0.007
T	**SVR-FQA**	**0.534**	**±0.101**	**0.734**	**±0.072**
	Extra Trees	0.471	±0.087	0.649	±0.063
	Rand. For.	0.454	±0.088	0.645	±0.068
A	**SVR-FQA**	**0.299**	**±0.088**	**0.644**	**±0.046**
	Extra Trees	0.172	±0.080	0.597	±0.051
	Rand. For.	0.250	±0.086	0.544	±0.060

Table 2: Multi-seed robustness (3 seeds: 42, 123, 456). Mean ± std of test R^2 and SRCC for top-3 regressors per modality. I, T, A denote Image, Text, and Audio modalities.

variance dominates signal. Lasso and ElasticNet collapse completely on text (test $R^2 \approx 0$, undefined SRCC): the ℓ_1 penalty eliminates all coefficients, indicating text quality information in CLIP space is distributed across many dimensions rather than concentrated in a sparse subset.

Multi-seed analysis (Table 2) further confirms three-seed robustness: image SRCC varies by only ±0.004–0.007, ranking is preserved across seeds, and even the best mean audio R^2 (0.299, SVR-FQA) remains far below image/text levels, confirming audio is genuinely harder to assess than the other two modalities.

To assess statistical reliability, we compute pairwise SRCC differences across the three independent seeds and report the mean ± standard deviation of the observed differences. On image, SVR-FQA leads Extra Trees by $+0.036 \pm 0.010$ and RF by $+0.064 \pm 0.011$, establishing SVR-FQA as the top image regressor across all three seeds. On text, SVR-FQA leads Extra Trees by $+0.085 \pm 0.012$ and RF by $+0.089 \pm 0.015$. On audio, SVR-FQA leads RF by $+0.100 \pm 0.017$ and Extra Trees by $+0.046 \pm 0.042$. All pairwise SRCC differences are strictly positive in every seed, confirming that SVR-FQA's ranking advantage is robust to data split variation. For audio test R^2, Linear Regression and SVR-FQA are not consistently distinguishable (mean difference -0.064 ± 0.079), consistent with complementary strengths: linear models capture more overall variance, while SVR-FQA preserves monotonic ranking better.

5.3 Ablation Study: Preprocessing Pipeline Robustness

We examine whether the audio preprocessing pipeline is robust to component removal with SVR-FQA.

The audio preprocessing chain consists of three steps: per-frequency Z-score normalization over the time axis, L2 normalization after the random projection, and scikit-learn's StandardScaler before regression. This configuration, which we call Baseline, was not selected by grid search or hyperparameter tuning. Each step follows naturally from the pipeline design. Z-score normalization is standard practice for MFCC features to prevent loudness variations from dominating the feature vectors. L2 normalization is required to align audio features with CLIP embeddings, which are L2 normalized by the model. StandardScaler is a conventional pre-regression step that zero-centers and unit-scales each feature dimension. The ablation asks whether any step is redundant.

We evaluate four configurations of the audio preprocessing pipeline: (1) **Baseline:** Z-score + L2 norm + StandardScaler; (2) **w/o Projection Norm:** no L2 after projection; (3) **w/o StandardScaler**; (4) **w/o Audio Z-Score**. Table 3 reports all metrics. Image/text rows are identical for configurations that only affect audio, confirming these ablations target audio exclusively. Removing StandardScaler substantially degrades SVR-FQA on image (SRCC 0.833 → 0.809) and text (SRCC 0.778 → 0.547), as the RBF kernel relies on standardized features for meaningful distance computation.

On audio, the pattern differs from image and text. Removing L2 projection normalization yields the overall best audio performance (SRCC= 0.602, R^2 = 0.350), exceeding the in-distribution result (Baseline and Table 1, SRCC= 0.591). This suggests that raw magnitude in the projected space carries quality-relevant signal that L2 normalization discards. Removing Z-score degrades audio (w/o Z-Score SRCC= 0.553 vs. Baseline 0.591), consistent with the kernel's reliance on well-scaled

Modality	Configuration	T.R^2↑	Te.R^2↑	MAE↓	RMSE↓	SRCC↑	PLCC↑	KLCC↑
Image	Baseline	0.780	0.623	3.866	5.657	0.833	0.790	0.639
	w/o Proj.Norm	0.780	0.623	3.866	5.657	0.833	0.790	0.639
	w/o StandardScaler	0.702	0.579	4.091	5.978	0.809	0.763	0.614
	w/o Audio Z-Score	0.780	0.623	3.866	5.657	0.833	0.790	0.639
Text	Baseline	0.654	0.598	0.153	0.262	0.778	0.788	0.631
	w/o Proj.Norm	0.654	0.598	0.153	0.262	0.778	0.788	0.631
	w/o StandardScaler	0.379	0.375	0.257	0.327	0.547	0.631	0.409
	w/o Audio Z-Score	0.654	0.598	0.153	0.262	0.778	0.788	0.631
Audio	Baseline	0.296	0.220	6.439	12.17	0.591	0.519	0.456
	w/o Proj.Norm	0.411	0.350	5.931	11.11	0.602	0.613	0.467
	w/o StandardScaler	0.420	0.297	6.018	11.55	0.556	0.576	0.441
	w/o Audio Z-Score	0.425	0.295	6.077	11.57	0.553	0.573	0.439

Table 3: SVR-FQA ablation results for four audio preprocessing configurations. All metrics are reported per modality–configuration pair. StandardScaler substantially affects SVR-FQA on image and text (the RBF kernel relies on standardized features). For audio, removing L2 projection normalization yields the best results (SRCC= 0.602), and removing Z-score normalization degrades performance (Baseline SRCC= 0.591 vs. w/o Z-Score SRCC= 0.553), confirming that per-frequency variance normalization benefits the kernel's distance computation.

features. The StandardScaler step has minimal effect on audio, consistent with features already being roughly scale-normalized by the projection. We retain StandardScaler for image and text.

5.4 Cross-Type Generalization

We formally test generalization to held-out distortion types: five from image (Fog, Snow, Rain, Posterize, Solarize), five from text (lowercase, homophone, reverse, abbreviation, contraction), and three from audio (Gain, PolarityInversion, TanhDistortion). Table 4 compares RF and SVR-FQA exclusively on held-out distortion types the models have never seen during training.

SVR-FQA outperforms RF across all modalities. On image, SVR-FQA achieves SRCC= 0.750 and Te.R^2 = 0.460 (vs. RF SRCC= 0.652, Te.R^2 = 0.286), a 15% SRCC improvement. On text, SVR-FQA attains Te.R^2 = 0.654, surpassing its in-distribution performance (Te.R^2 = 0.598; Table 3). On audio, RF fails markedly (Te.R^2 = −2.954, SRCC= −0.073), whereas SVR-FQA degrades less: Te.R^2 = −0.252 and SRCC remains positive at 0.144, indicating that SVR regularization prevents memorizing narrow distortion signatures.

5.5 Simple Baselines: What Does "Good" Mean?

To contextualize the regressor performance, we establish trivial lower bounds using constant predictors.

We evaluate two trivial baselines: Mean Predictor ($\bar{y}_{\text{train}}$) and Median Predictor, to establish performance lower bounds (Table 5). SRCC values establish a difficulty gradient (image > text > audio; Section 5.2), with audio substantially harder than both image and text. Relative MAE reduction over the mean predictor reflects this and is more defensible than raw R^2 comparison across different target distributions. Cross-Type audio with RF (R^2 = −2.954, negative SRCC) confirms that tree ensembles learn anti-informative patterns on unseen distortions, while SVR-FQA avoids this collapse (Te.R^2 = −0.252, SRCC= 0.144; Table 4), demonstrating that kernel regularization provides meaningful protection against overfitting to distortion-specific signatures.

5.6 Limitations

Several limitations should be noted.

(1) The benchmark is entirely synthetic; generalizability to real-world distortions (camera noise, codec cascades, and transmission errors) remains unevaluated.

Mod.	Method	Te.R^2↑	SRCC↑	PLCC↑	MAE↓	RMSE↓
I	Rand. For.	0.286	0.652	0.720	5.167	6.139
	SVR-FQA	**0.460**	**0.750**	**0.839**	**4.233**	**5.338**
T	Rand. For.	0.202	0.583	0.758	0.344	0.378
	SVR-FQA	**0.654**	**0.659**	**0.865**	**0.187**	**0.249**
A	Rand. For.	-2.954	-0.073	-0.150	13.66	18.93
	SVR-FQA	**-0.252**	**0.144**	**0.104**	**8.129**	**10.65**

Table 4: Cross-distortion-type generalization: Rand. For. vs. SVR-FQA. Bold denotes the better result per metric–modality pair. SVR-FQA dominates across all settings. On audio, Rand. For. collapses catastrophically (Te.$R^2 = -2.954$, SRCC$= -0.073$), while SVR-FQA degrades gracefully with positive SRCC (0.144). I, T, A denote Image, Text, Audio.

Mod.	Method	Te.R^2↑	MAE↓	RMSE↓	SRCC↑
I	Mean Pred.	0.000	7.496	9.210	—
	Median Pred.	-0.006	7.486	9.233	—
	SVR-FQA	**0.623**	**3.866**	**5.657**	**0.833**
T	Mean Pred.	0.000	0.367	0.413	—
	Median Pred.	-0.072	0.345	0.428	—
	SVR-FQA	**0.598**	**0.153**	**0.262**	**0.778**
A	Mean Pred.	0.000	10.54	13.78	—
	Median Pred.	-0.089	9.802	14.38	—
	SVR-FQA	**0.220**	**6.439**	**12.17**	**0.591**
	RF Cross-Type	-2.954	13.66	18.93	-0.073

Table 5: Simple baselines vs. best regressors per modality. Mean Predictor achieves $R^2 \approx 0$ by definition. RF Cross-Type audio (Te.$R^2 = -2.954$) is included for contrast, highlighting the severity of tree ensemble collapse on unseen distortions. I, T, A denote Image, Text, Audio.

(2) The audio pipeline uses a coarse MFCC summary that discards temporal structure, and the fixed random projection cannot recover distortion-invariant representations that a learned projection might capture.

(3) With 160 reference images and 512-D features, the high-dimensional low-sample regime amplifies tree-based overfitting and limits statistical power for per-distortion subgroup analyses.

(4) Evaluation uses objective metrics (PSNR/BLEU) rather than human perceptual judgments; bridging to subjective quality requires cross-modal human annotations, which remain scarce.

The benchmark and code will be publicly released.

6 Conclusion

We introduced a unified multimodal quality assessment framework with three contributions: a 50-type synthetic benchmark, training-free randomized projection for audio-to-CLIP alignment, and SVR-FQA, a modality-specific SVR tuning method that achieves the best ranking performance among all compared regressors across all three modalities and evaluation metrics. SVR-FQA uses per-modality grid-searched hyperparameters (image: $C = 100, \epsilon = 0.05, \gamma =$ scale; text: $C = 10, \epsilon = 0.05, \gamma =$ auto; audio: $C = 100, \epsilon = 1.0, \gamma =$ scale) that substantially and consistently outperform default SVR and tree ensembles. It leads all modalities in SRCC (image 0.833, text 0.778, audio 0.591), establishing a clear and consistent difficulty gradient with audio as the hardest modality. Multi-seed analysis confirms statistical robustness across seeds. SVR-FQA also proves more robust in cross-type generalization experiments, outperforming RF on all modalities and avoiding catastrophic collapse on unseen audio distortions (SVR-FQA Te.$R^2 = -0.252$, SRCC$= 0.144$ vs. RF Te.$R^2 = -2.954$, SRCC$= -0.073$). This demonstrates that kernel regularization provides inherent, training-free protection against memorizing distortion-specific signatures. Future work should pursue learned audio-visual projections, temporal audio features, and evaluation on naturally distorted, real-world media.

7 Ethics Statement

This work uses only synthetic data generated from publicly available images, with no human subjects, personally identifiable information, or real-world user data involved. We acknowledge the risk of quality assessment systems being misused for mass surveillance or content manipulation; however, our framework is purely evaluative and its transparency (open code, reproducible benchmarks) mitigates such risks by enabling public scrutiny. We see no direct path from this research to harmful applications, and the synthetic benchmark does not encode demographic or cultural biases beyond those present in the underlying vision-language models used for caption generation.

AI use statement

No generative AI was used to produce the scientific content, experiments, analyses, or conclusions. AI assistance was limited to LaTeX formatting and final layout inspection; the authors verified and approved all content.

References

Y. Gao et al. Toward standardized benchmarks for cross-model feature coding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024.

Y. Gao et al. Cross-architecture feature coding via distribution alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. AudioCLIP: Extending CLIP to image, text and audio. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 976–980, 2022.

William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.

Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 12888–12900, 2022.

Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, 2004.

Brian McFee, Colin Raffel, Dawen Liang, Daniel P. W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in Python. In *Proceedings of the 14th Python in Science Conference (SciPy)*, pp. 18–25, 2015.

Gabriel Mittag, Babak Naderi, Assmaa Chehadi, and Sebastian Möller. NISQA: A deep CNN-self-attention model for multidimensional speech quality prediction with crowdsourced datasets. In *Proceedings of Interspeech*, pp. 2127–2131, 2021.

Anish Mittal, Anush K. Moorthy, and Alan C. Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012.

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 311–318, 2002.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, and C.-C. Jay Kuo. Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication*, 30:57–77, 2015.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 8748–8763, 2021.

Chandan K. A. Reddy, Vishak Gopal, and Ross Cutler. DNSMOS: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6493–6497, 2021.

Antony W. Rix, John G. Beerends, Michael P. Hollier, and Andries P. Hekstra. Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 2, pp. 749–752, 2001.

Hamid R. Sheikh, Muhammad F. Sabir, and Alan C. Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing*, 15(11):3440–3451, 2006.

Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. BERTScore: Evaluating text generation with BERT. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.