



HAL
open science

Subjective Media Quality Recovery from Noisy Raw Opinion Scores: A Non-Parametric Perspective

Andrés Altieri, Lohic Fotio Tiotsop, Giuseppe Valenzise

► **To cite this version:**

Andrés Altieri, Lohic Fotio Tiotsop, Giuseppe Valenzise. Subjective Media Quality Recovery from Noisy Raw Opinion Scores: A Non-Parametric Perspective. *IEEE Transactions on Multimedia*, 2024, pp.1-16. 10.1109/TMM.2024.3390113 . hal-04393713v1

HAL Id: hal-04393713

<https://centralesupelec.hal.science/hal-04393713v1>

Submitted on 15 Jan 2024 (v1), last revised 29 Apr 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Subjective Media Quality Recovery from Noisy Raw Opinion Scores: A Non-Parametric Perspective

Andrés Altieri, Lohic Fotio Tiotso, *Member, IEEE*, Giuseppe Valenzise, *Senior Member, IEEE*

Abstract—Recent parametric approaches to recover subjective media quality from noisy opinion scores involve assumptions that are challenged by numerous complex factors influencing viewer behavior in subjective tests. These approaches commonly involve modeling scoring behavior using parametric distributions, with parameters estimated through methods like maximum likelihood or Bayesian techniques. This paper proposes instead a *non-parametric* perspective to address the media quality recovery problem, without making any a priori assumption on the subjects' scoring behavior. We propose and motivate an approach to measure the reliability of each single opinion score of a subject. The proposed measure of reliability is then leveraged to formulate an *Entropy-based Subjective Quality Recovery (ESQR)* algorithm. Simulations and experiments on real datasets show that the proposed ESQR algorithm compares favorably to several state-of-the-art approaches in terms of robustness to noise across various testing conditions. The code to run ESQR and reproduce the results in this paper is made freely available at: <http://media.polito.it/ESQR>.

Index Terms—Media quality assessment, Subjective quality recovery, Opinion scores reliability, Non-parametric model.

I. INTRODUCTION

The opinion scores gathered in subjective tests are known to be affected by noise, caused by factors such as subjects' fatigue and distraction, unforeseen software issues and other uncontrolled elements influencing the emotional state of a subject. Several approaches to recover the subjective media quality from noisy raw ratings have therefore been proposed [1]–[3]. The conventional approach computes the mean of the opinion scores (MOS) gathered for a given stimulus as an estimate of the subjective quality of that stimulus. Despite its popularity, the MOS is known to be particularly sensitive to outlier opinion scores [4]. As a consequence, more sophisticated approaches that exploit parametric statistical models have recently been proposed to cope with the subjective media quality recovery problem [3], [5]–[7].

The diagram in the top part of Figure 1 illustrates the main steps of parametric approaches (e.g., [1], [3], [5], [7], [8]). They essentially implement three steps: i) a parametric statistical model explaining the scoring behavior of each subject is assumed; ii) the model's parameters are estimated;

Andrés Altieri is with Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes (UMR 8506), 91190 Gif-sur-Yvette, France and with CSC-CONICET, Buenos Aires (1425), Argentina. Email: andres.altieri@centralesupelec.fr

Lohic Fotio Tiotso is with the Control and Computer Engineering Department, Corso Duca degli Abruzzi 24, Politecnico di Torino, Torino, Italy. Email: lohic.fotiotiotso@polito.it

Giuseppe Valenzise is with Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes (UMR 8506), 91190 Gif-sur-Yvette, France. Email: giuseppe.valenzise@centralesupelec.fr.

iii) finally the desired subjective quality is derived. The main advantage of parametric approaches is that they explain the scoring behavior of subjects. However, the use of parametric approaches entails several limitations: i) Lack of robustness, as the model's assumptions that allow interpretability and numerical stability in practice do not hold in several application scenarios [5], [7]. ii) High risk of underfitting since despite the huge number of factors that influence the subjects scoring behavior, very few parameters are considered to preserve the model's numerical stability. For instance, authors of parametric approaches (e.g., [1], [3], [5]) very often use a single parameter to capture the reliability of a subject, disregarding the fact that the latter varies with the characteristics of each stimulus that the subject is asked to evaluate. iii) High computational complexity due to the parameter estimation process that involves solving a complex optimization problem [3], [6].

In an attempt to overcome the aforementioned limitations of parametric methods, this paper proposes a novel *non-parametric* approach to recover the subjective quality from noisy opinion scores. The general scheme of the proposed method is illustrated in the bottom part of Figure 1. The proposed non-parametric approach considers two main steps: first, it measures the reliability of each single opinion score without resorting to any assumed probabilistic scoring model. Then, the recovered quality is computed as a weighted sum of the opinion scores, in which the weight of each opinion score is determined by its reliability.

The proposed approach to measure the reliability of an opinion score finds its theoretical explanation in information theory. In fact, it is inversely proportional to a measure of how surprising that opinion score is for the quality of the stimulus under evaluation. The proposed method addresses the main shortcomings of parametric approaches since it: i) does not consider any scoring model involving simple yet restrictive assumptions on the subject's behavior; ii) does not require to solve an optimization problem for parameters estimation, thus a significant computational burden is avoided; iii) does not suffer underfitting or overfitting issues since it is a non-parametric approach.

This paper's contribution can be summarized as follows:

- 1) we propose a non-parametric approach to measure the reliability of each opinion score of each subject;
- 2) based on this reliability score, we introduce a novel algorithm, called Entropy-based Subjective Quality Recovery (ESQR), to robustly recover the quality of stimuli.

Computational experiments demonstrate that ESQR offers a subjective quality estimate with reduced uncertainty compared to current techniques. Additionally, ESQR exhibits lower sen-

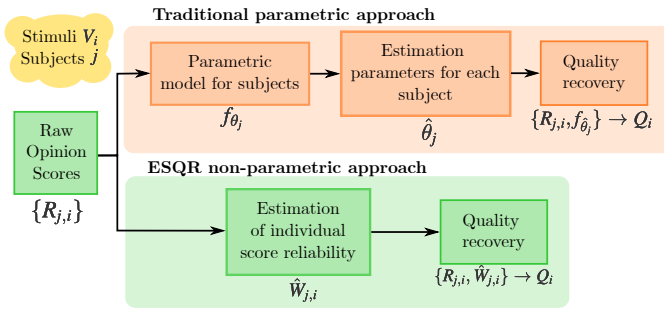


Fig. 1. Parametric approaches assume that the observed opinion scores come from a probabilistic model f_{θ} , that depends on a finite set of parameters θ . The parameters θ are then estimated, and the subjective quality Q is recovered from the assumed probabilistic model using the estimated parameters. The proposed non-parametric approach (ESQR) instead measures how reliable is each single opinion score without resorting to a predefined model. The recovered quality Q is then computed as a weighted sum of the opinion scores in which the weight of each opinion score is determined by its reliability.

sitivity to noise when compared to five state-of-the-art quality recovery approaches. Furthermore, the findings suggest that the performance of ESQR is relatively robust across various application scenarios. This resilience could be attributed to the fact that, unlike parametric approaches, our method circumvents assumptions about subjects' scoring behavior that may be invalid in certain application contexts.

The rest of this paper is organized as it follows: Section II reviews and analyzes the prior art. Our non-parametric approach to analyze subject's behavior in subjective tests is described in Section III, while the proposed quality recovery algorithm is discussed in Section IV. Computational experiments and the related results are presented in Section V, and the conclusions are drawn in Section VI.

II. RELATED WORK

In recent years, the research on subjective quality recovery has gained increasing importance, specially since the COVID-19 pandemic democratized the collection of opinion scores in non-highly controlled environments, i.e., crowdsourcing subjective tests [9], [10]. Subjective tests can be conducted using a rating quality scale [11] or by performing pair comparisons [12]. Examples of approaches to deal with noise in pair comparison-based subjective tests can be found, e.g., in [13]–[16]. Pairwise preferences and rating scores can also be fused into a common quality scale, thus reducing the noise of individual experiments [17], [18]. In this paper, we focus on the scenario where a rating quality scale is used to elicit and collect opinions.

The mean of the opinion scores (MOS) gathered for a given stimulus has long been considered as a good estimate of the subjective quality of that stimulus. Unfortunately, the MOS is very sensitive to outlier ratings since it attributes the same importance to reliable and unreliable subjects. To address this limitation, several relevant ITU recommendations have appeared. The ITU-R BT.500 [4] recommends to identify outlier subjects and exclude their opinion scores from the dataset before computing the MOS. The ITU-T P.910 [19] introduces the so-called Absolute Category Rating (ACR) scale and recommends the use of the quality recovery algorithm proposed in ITU-R BT.500 [4] together with Confidence

Intervals (CIs) to deal with noisy raw opinion scores collected using the ACR scale. Finally, the ITU-T P.913 [20] suggests to first perform a bias removal step on the data and then potentially use the algorithm proposed in ITU-R BT.500 [4].

All these ITU recommendations support the use of an algorithm that performs subject exclusion. However, removing all the opinion scores of a subject from the dataset probably causes an unnecessary loss of relevant information. In fact, as underlined in [3], [5], considering that the excluded subjects gave inaccurate opinion scores for all the stimuli is overly conservative. Therefore, researchers have recently proposed advanced statistical approaches to circumvent subject exclusion.

These recent methods mainly assume that each opinion score of each subject follows a probability distribution that can be characterized by a finite number of parameters. These parameters are then estimated using statistical frameworks such as the Maximum Likelihood Estimation (MLE) [21] or the Expectation Maximization Algorithm (EMA) [22]. Specifically, most of these recent parametric approaches assume that the scoring behavior of a subject can be modeled with two parameters, i.e., the subject's bias and inconsistency [1], [5], [7], [8], [23]. The bias is defined as a systematic tendency of a subject to provide smaller (negative bias) or larger (positive bias) opinion scores than the actual subjective quality of the stimulus that is being rated. The inconsistency instead is a measure of the inability of a subject to provide consistent opinion scores when rating the same stimulus more than once.

When relying on bias and inconsistency, the authors usually assume that each raw opinion score of each subject is a realization of a Gaussian random variable. In [1], [8], the mean of such a Gaussian random variable was modeled as the sum of two parameters, i.e., the subjective quality to be recovered and the subject's bias. The variance was also expressed as the sum of two other parameters: the subject's inconsistency and the stimulus's ambiguity. The authors of [5] argued that the Gaussian model proposed in [1] is still valuable if one gets rid of the stimulus's ambiguity, thus the variance of the Gaussian random variable modeling each opinion score of each subject was assumed to be equal only to the subject's inconsistency. The authors then proposed an iterative algorithm, called alternating projection, to recover the subjective quality based on their proposed scoring model. The alternating projection algorithm was implemented in the Netflix SUREAL software [24]. The authors of [7] improved the Gaussian model of [5] by integrating into it the so-called Standard deviation of Opinion Scores (SOS) hypothesis, proposed by the authors of [25], in order to account for the fact that subjects are more accurate when rating very low or very high quality. Finally, the authors of [23] suggested to subtract the MOS from the raw opinion scores and divide the result by the SOS to obtain the so-called Z-scores. They then argued that the bias, the inconsistency and the subjective quality can be estimated more reliably and efficiently from the Z-scores rather than the original opinion scores. This yielded a subjective quality recovery approach called "ZREC".

Another interesting parametric approach is the one proposed in [3]. The authors assumed that each rating of each subject is

a realization of a mixture of two probability distributions. The first probability distribution models accurate opinion scores given by the subject, while the second accounts for the cases in which the subject provides inaccurate opinion scores. The accuracy of the subject is measured with one parameter, i.e., the probability that the subject would score the quality according to the distribution modeling accurate opinion scores. An EMA is proposed by the authors to estimate the model's parameters and thus the subjective quality.

More recently, the authors of [6] introduced the so-called Regularized MLE (RMLE) approach to recover subjective media quality. Basically, they proposed a regularization term to be added to the likelihood function before solving the optimization problem, the solution of which provides an estimate of the parameters that characterize the subjective quality to be recovered. The proposed regularization term is meant to penalize opinion scores that are potentially noisy.

Although the superiority over the MOS of the parametric approaches discussed so far has already been proved in specific applications, the assumptions made by these approaches on the scoring behavior of subjects raises some questions on their general applicability. In fact, as it will be observed in Section V, the performance of some of these approaches with respect to the performance of the MOS can vary significantly from one application scenario to another. In addition, from a theoretical point of view, the optimization problem guiding the parameter estimation process in parametric models is usually computationally challenging and might not even have a unique optimal solution as pointed out in [23]. This introduces additional challenges that parametric models have to overcome.

This paper adopts a different perspective than the recent parametric approaches described in the previous paragraphs. We avoid making assumptions about any particular parametric probabilistic model able to explain all the opinion scores of a subject. Instead, a non-parametric approach to measure the reliability of each opinion score of each subject and recover the subjective quality of each stimulus is proposed. By refraining from imposing stringent assumptions on the subjects' scoring behavior, we aim at proposing an approach whose accuracy is preserved over a broad spectrum of practical situations.

III. A NON-PARAMETRIC MEASURE OF RELIABILITY

This section introduces and motivates our approach to measure the reliability of each opinion score of each subject. The proposed measure will be used later in Section IV to derive a new algorithm to recover the subjective media quality.

As discussed in the previous section, parametric approaches to recover the subjective media quality from raw opinion scores typically consider some parameters that measure the reliability of each subject during the subjective test. For instance, in [5], [23], the inverse of the square of the subject's inconsistency was used as measure of reliability. The authors of [3] estimated the probability of each subject to provide accurate scores, and then used this probability as an indicator of the reliability of the subject. Clearly, the effectiveness of these measures of reliability strongly depends on the assumptions of the underlying parametric model. In this paper, we argue

that the reliability of each opinion score can be effectively measured without resorting to any parametric scoring model.

A. Notation and Hypotheses

Let us introduce the following notation:

- \mathcal{I} , a set of rated stimuli.
- \mathcal{J} , a set of subjects.
- $\mathcal{I}_j \subset \mathcal{I}$ the subset of the stimuli rated by the subject $j \in \mathcal{J}$.
- $\mathcal{J}_i \subset \mathcal{J}$, the subset of subjects that rated the stimulus $i \in \mathcal{I}$ using a discrete scale in the range $\{1, \dots, K\}$.

For a given stimulus $i \in \mathcal{I}$, there might be different opinion scores that accurately characterize its quality. The ground-truth quality can be modeled as a discrete random variable, characterized by its probability mass function (pmf). Specifically, we introduce the following notation:

- V_i the discrete random variable that describes the latent ground-truth quality of a stimulus $i \in \mathcal{I}$ in the range $\{1, \dots, K\}$. In the absence of any noise, subjects sample this random variable, producing opinion scores.
- p_{V_i} denotes the pmf of V_i .

The problem of quality estimation has been formulated in some cases as that of recovering p_{V_i} (refer to, e.g., [3], [26]). In practice, most of the quality assessment literature focuses on predicting point estimates of the quality pmf, such as the mean opinion score. We also follow the latter approach in this paper, i.e., we are interested in estimating $q_i = \mathbf{E}[V_i]$.

In real-world scenarios, the latent random variable describing quality is observed through a set of stochastic, *noisy* observations, i.e., opinion scores. This model effectively captures the intrinsic subjectivity of scores, arising, e.g., from different expectations in terms of quality. We will therefore denote by:

- $R_{j,i}$, the discrete random variable modeling the score of the subject j for the stimulus i on a quality scale in the range $\{1, \dots, K\}$.
- $p_{R_{j,i}}$ denotes the pmf of $R_{j,i}$.

Our goal is to propose a robust estimator of q_i by introducing an approach to measure the reliability of each observed opinion score $\{R_{j,i}\}_{j \in \mathcal{J}_i, i \in \mathcal{I}}$. We will denote as Q_i our proposed estimator of q_i .

B. Measuring the Reliability of an Opinion Score

We propose the following definition of reliability:

Definition 1. *The reliability $W_{j,i}$ of the opinion score $R_{j,i}$ given by the subject j when assessing the quality of the stimulus i is the following ratio:*

$$W_{j,i} = -\frac{1}{\log(p_{V_i}(R_{j,i}))}. \quad (1)$$

An estimate of the distribution p_{V_i} is required to compute the reliability measures $W_{j,i}$. We will explore how to perform such an estimation in Section IV. For the moment, in order to motivate why the formula in Eq. (1) provides a suitable measure of reliability for each opinion score gathered in a subjective test, let us assume that the distribution p_{V_i} is known for each stimulus $i \in \mathcal{I}$.

Note that $W_{j,i}$ is a measure of reliability if and only $W_{j,i}^{-1}$ is a measure of unreliability. To motivate the suitability of $W_{j,i}$ as a measure of reliability, let us introduce the following statistic for each subject:

$$S_j(\mathcal{I}_j) = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} W_{j,i}^{-1}. \quad (2)$$

We will prove that the statistic $S_j(\mathcal{I}_j)$ can be used to measure the average unreliability of the subject j . In fact, the following proposition holds:

Proposition 1. *For each subject j , if there is a constant c such that $\text{var}[W_{j,i}^{-1}] < c \forall i \in \mathcal{I}_j$, then, as $|\mathcal{I}_j| \rightarrow \infty$, $S_j(\mathcal{I}_j)$ converges in the mean square sense to:*

$$\bar{h}_j(\mathcal{I}_j) = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} H(p_{R_{j,i}}) + \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} D_{\text{KL}}(p_{R_{j,i}} \| p_{V_i}), \quad (3)$$

where $H(p_{R_{j,i}})$ is the entropy of the distribution $p_{R_{j,i}}$ and $D_{\text{KL}}(p_{R_{j,i}} \| p_{V_i})$ denotes the Kullback–Leibler (KL) divergence between $p_{R_{j,i}}$ and p_{V_i} .

Proof. We should prove that:

$$\lim_{|\mathcal{I}_j| \rightarrow \infty} \mathbb{E} \left[(S_j(\mathcal{I}_j) - \bar{h}_j(\mathcal{I}_j))^2 \right] = 0. \quad (4)$$

To do this we show that $S_j(\mathcal{I}_j)$ is an unbiased estimator of $\bar{h}_j(\mathcal{I}_j)$ and that its variance goes to zero as $|\mathcal{I}_j| \rightarrow \infty$.

For the bias we have:

$$\begin{aligned} \mathbb{E}_{R_{j,i}} [S_j(\mathcal{I}_j)] &= -\frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} \mathbb{E}_{R_{j,i}} [\log(p_{V_i}(R_{j,i}))] \\ &= -\frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} \mathbb{E}_{R_{j,i}} \left[\log \left(\frac{p_{V_i}(R_{j,i})}{p_{R_{j,i}}(R_{j,i})} \right) \right] \\ &\quad - \mathbb{E}_{R_{j,i}} [\log(p_{R_{j,i}}(R_{j,i}))] \quad (5) \\ &= \bar{h}_j(\mathcal{I}_j), \quad (6) \end{aligned}$$

where from Eq. (5) to Eq. (6), we use the definitions of entropy and KL divergence. Also, we consider that if there are any terms in the expectation for which $p_{R_{j,i}}(r) = 0$, then those terms are zero through the continuity definition that $0 \log 0 = 0$ [27], so in reality we are only dividing by $p_{R_{j,i}}$ for the positive terms that impact the expectation.

Finally, to compute the variance of $S_j(\mathcal{I}_j)$ we make use of the fact that the opinion scores can be considered independent and the assumption that the variance of $W_{j,i}^{-1}$ is finite:

$$\text{var}[S_j(\mathcal{I}_j)] = \frac{1}{|\mathcal{I}_j|^2} \sum_{i \in \mathcal{I}_j} \text{var}[W_{j,i}^{-1}] \leq \frac{c}{|\mathcal{I}_j|} \xrightarrow{|\mathcal{I}_j| \rightarrow \infty} 0, \quad (7)$$

which completes the proof. \square

Proposition 1 basically states that the statistic $S_j(\mathcal{I}_j)$ converges to the overall unreliability of the subject j , since for a large value of $|\mathcal{I}_j|$, the following approximation can be used:

$$S_j(\mathcal{I}_j) \approx \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} H(p_{R_{j,i}}) + \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} D_{\text{KL}}(p_{R_{j,i}} \| p_{V_i}). \quad (8)$$

Remembering that the subject's inconsistency is a measure of the inability to repeat exactly the same opinion score when

rating several times the same stimulus, it is not difficult to observe that the more the subject j is inconsistent when rating the stimulus i , the larger is the entropy $H(p_{R_{j,i}})$ of the random variable $R_{j,i}$. Hence, the quantity $\frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} H(p_{R_{j,i}})$ in Eq. (8) captures the average *inconsistency* of the subject j . On the other hand, the following $\frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} D_{\text{KL}}(p_{R_{j,i}} \| p_{V_i})$ indicates on average how far the opinion scores of the subject j are expected to be from the accurate opinion scores. Hence, $\frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} D_{\text{KL}}(p_{R_{j,i}} \| p_{V_i})$ measures the average *inaccuracy* of the subject j .

In the light of the above interpretation of Proposition 1 it turns out that asymptotically the value of the statistic $S_j(\mathcal{I}_j) = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} W_{j,i}^{-1}$ measures the average *unreliability* of the subject j . Each single term, i.e., $W_{j,i}^{-1}$, of this statistic can therefore be considered as the measure of the unreliability of the opinion score given by the subject j to the quality of the stimulus i . This motivates our definition of $W_{j,i}$ as a measure of reliability.

It is worth noting that in information theory, the logarithm of the probability of an event is called the self-information of the event. It can be interpreted as a measure of the likelihood of the event, and thus the level of surprise that the observation of the event entails. Therefore, $-\log(p_{V_i}(R_{j,i}))$ can be interpreted as the measure of how surprising (thus potentially unreliable) is the opinion score $R_{j,i}$ if it is supposed to come from p_{V_i} . This is another way to motivate why the $W_{j,i}$ defined in Eq. (1) measures the reliability of the opinion score $R_{j,i}$.

Notice that Proposition 1, which is the main theoretical foundation of the proposed reliability measure, is based on the notion of entropy of a probability distribution (the KL divergence is the relative entropy between two probability distributions). For this reason, the algorithm proposed in the next section, grounded in the proposed reliability measure, shall be denoted as **Entropy-based Subjective Quality Recovery (ESQR)**.

IV. RECOVERING THE SUBJECTIVE QUALITY: THE ESQR ALGORITHM

We now introduce the proposed ESQR algorithm to recover the subjective quality. The key idea of the method is to weight each opinion score in the computation of the subjective media quality using the reliability measure described in Section III.

To measure the reliability of each opinion score using the formula in Eq. (1), the distribution of accurate opinion scores p_{V_i} of each stimulus i is required. In practice, as already mentioned, this distribution is unknown, so an estimate of it is needed. Therefore, the ESQR algorithm implements two main steps:

- 1) For each stimulus i , an estimate \hat{p}_{V_i} of the distribution of accurate opinion scores is obtained from the observed sample of opinion scores $\{R_{j,i}\}_{i \in \mathcal{I}_j, j \in \mathcal{J}}$.
- 2) Then, \hat{p}_{V_i} is used in Eq. (1) to get an estimate $\hat{W}_{j,i}$ of the reliability of each opinion score. $\hat{W}_{j,i}$ is used to define our estimator of the ground-truth subjective quality q_i of each stimulus i as it follows:

$$Q_i = \frac{\sum_{j \in \mathcal{J}_i} \hat{W}_{j,i} \cdot R_{j,i}}{\sum_{k \in \mathcal{J}_i} \hat{W}_{k,i}}. \quad (9)$$

Note that from the definition of Q_i above, the contribution of the subject $j \in \mathcal{J}_i$ to the determination of the ground truth quality of the stimulus $i \in \mathcal{I}$ is weighted by:

$$\omega_{ij} = \frac{\hat{W}_{j,i}}{\sum_{k \in \mathcal{J}_i} \hat{W}_{k,i}}. \quad (10)$$

A. Estimating p_{V_i}

Unlike previous approaches, we do not make assumptions on the shape of the distributions p_{V_i} , $i \in \mathcal{I}$. Therefore, we use a non-parametric approach to estimate the distribution of accurate opinion scores.

The simplest non-parametric estimation of p_{V_i} is based on computing the histogram of opinion scores. In fact, $p_{R_{j,i}}$ can be estimated by the histogram of the opinion scores that the subject j gave for the quality of the stimulus i . For instance, if a single opinion score is collected per stimulus, then $p_{R_{j,i}}$ is a probability mass function that attributes a one to the observed opinion score and 0 to all the other opinion scores on the quality scale. The observed distribution of opinion scores for each stimulus i can then be estimated as:

$$\bar{p}_{V_i}(r) = \frac{1}{|\mathcal{J}_i|} \sum_{j \in \mathcal{J}_i} p_{R_{j,i}}. \quad (11)$$

In practice, this estimate \bar{p}_{V_i} can be influenced by the potential presence of noise in the raw opinion scores collected from the subjects. To obtain a more accurate estimation that closely aligns with the true distribution p_{V_i} we suggest a different approach for weighting different subjects' contributions, as opposed to the uniform weighting ($1/|\mathcal{J}_i|$) used in Eq. (11). Specifically, our proposition is based on the observation that the more a subject's opinion scores correlate with those of other subjects, the more trustworthy that subject is. Consequently, to enhance the precision of our estimate of p_{V_i} we propose applying a weighting coefficient, denoted as ϵ_j to each histogram $p_{R_{j,i}}$ which depends on the overall correlation between subject j and the other subjects.

We employ C_{jk} , the Spearman Rank Order Correlation Coefficient (SROCC) between the opinion scores of the subject j and those of the subject k , as a non-parametric measure of correlation. To compute the average correlation between subject j 's opinion scores and those of all the other subjects, we utilize the Fisher Z-Transformation (FZT), as suggested in [28]. For each subject j , the FZT is applied to the SROCC values C_{jk} $k = 1, 2, \dots, j-1, \dots, j+1, \dots, |\mathcal{J}_i|$. The average of the obtained values is computed. The inverse of the FZT, here denoted by FZT^{-1} , is then applied to the obtained average to obtain the overall correlation \hat{C}_j between the opinion scores of the subject j and those of the other subjects. Finally, the importance ϵ_{ij} of the histogram $p_{R_{j,i}}$ in the estimation of the distribution p_{V_i} is expressed as:

$$\epsilon_{ij} = \frac{|\hat{C}_j|}{\sum_{k \in \mathcal{J}_i} |\hat{C}_k|} \quad i \in \mathcal{I}, \quad j \in 1, 2, \dots, |\mathcal{J}_i|. \quad (12)$$

Therefore, the final estimate \hat{p}_{V_i} of the distribution p_{V_i} , used to recover the subjective quality of each stimulus i as defined in Eq (9), is obtained from the following formula:

$$\hat{p}_{V_i} = \sum_{j \in \mathcal{J}_i} \epsilon_{ij} \cdot p_{R_{j,i}} \quad i \in \mathcal{I}. \quad (13)$$

Algorithm 1: Entropy based Subjective Quality Recovery (ESQR)

Data: $R_{j,i}$, $i \in \mathcal{I}_j$; $j \in \mathcal{J}$ // stimuli i , subjects j

- 1 $C_{jk} \leftarrow \text{SROCC}(R_{j,\cdot}, R_{k,\cdot})$ $j, k \in \mathcal{J}$ // pairwise subject scores correlation
- 2 $\hat{C}_j \leftarrow \text{FZT}^{-1} \left(\frac{\sum_{k \in \mathcal{J}} \text{FZT}(C_{jk})}{|\mathcal{J}|} \right)$ $j \in \mathcal{J}$ // overall subject-to-subject correlation
- 3 $\epsilon_{ij} \leftarrow \frac{|\hat{C}_j|}{\sum_{k \in \mathcal{J}_i} |\hat{C}_k|}$ $i \in \mathcal{I}$; $j \in \mathcal{J}_i$ // importance of the ratings of subject j in the P_{V_i} estimation
- 4 $\hat{p}_{V_i} \leftarrow \sum_{j \in \mathcal{J}_i} \epsilon_{ij} p_{R_{j,i}}$ $i \in \mathcal{I}$ // estimate the distribution P_{V_i}
- 5 $\hat{W}_{j,i} \leftarrow \frac{1}{-\log(\hat{p}_{V_i}(R_{j,i}))}$ $i \in \mathcal{I}_j$; $j \in \mathcal{J}$ // estimate each opinion's score reliability
- 6 $Q_i \leftarrow \frac{\sum_{j \in \mathcal{J}_i} \hat{W}_{j,i} R_{j,i}}{\sum_{k \in \mathcal{J}_i} \hat{W}_{k,i}}$ $i \in \mathcal{I}$ // estimate the quality

Result: Q_i , $i \in \mathcal{I}$

It is worth noticing that other more sophisticated estimations of \hat{p}_{V_i} might be possible. However, we argue that even a simple and possibly noisy approximation of p_{V_i} , such as the one proposed above, is a good starting point for the second step of the proposed ESQR algorithm, where atypical subjects are further penalized using the reliability measure introduced in Section III. The experiments in Section V support this claim with empirical evidence.

The proposed ESQR algorithm is summarized in Algorithm 1. The input of Algorithm 1 is the set of the observed opinion scores $R_{j,i}$ $i \in \mathcal{I}_j$; $j \in \mathcal{J}$. The output is the recovered subjective quality of each stimulus i . The notation $R_{j,\cdot}$ is used to indicate all the opinion scores of the subject j .

B. Confidence Interval of the Recovered Quality

Using a similar formula as in [23] to compute the standard deviation of a weighted sum of opinion scores, for each stimulus i , an unbiased estimator of the standard deviation of Q_i can be computed as follows:

$$\sigma_{Q_i} = \sqrt{\frac{|\mathcal{J}_i|}{|\mathcal{J}_i| - 1} \frac{\sum_{j \in \mathcal{J}_i} \hat{W}_{j,i} (R_{j,i} - Q_i)^2}{\sum_{k \in \mathcal{J}_i} \hat{W}_{k,i}}}. \quad (14)$$

From the standard deviation in Eq (14), the 95% confidence interval (CI) of the recovered quality stimulus i can be computed as:

$$\text{CI}_{Q_i} = Q_i \pm 1.96 \frac{\sigma_{Q_i}}{\sqrt{|\mathcal{J}_i|}}. \quad (15)$$

Eq. (15) assumes that the estimator Q_i is normally distributed. In Appendix A, we provide theoretical conditions for this to happen, and we show through simulation that these conditions are indeed reasonable in a practical scenario. Also, notice that this normality assumption only applies to the *estimator* of the ground-truth quality, and not the individual opinion scores, which can follow any arbitrary distribution in our framework.

V. NUMERICAL EXPERIMENTS

A. Experimental Settings

The evaluation of subjective media quality recovery methods is challenging, since there is no observable “true” quality of stimuli to be used as ground truth. In related work, e.g., [1], [5]–[7], [23], the effectiveness of quality recovery approaches is assessed in terms of: i) robustness to the insertion of synthetic noise in the quality scores; ii) and uncertainty on the recovered subjective media quality. We will therefore use similar experiments in this paper to evaluate the effectiveness of the proposed ESQR algorithm.

We compare the proposed ESQR algorithm with five state-of-the-art quality recovery approaches, i.e.: the MOS, the algorithm recommended in the ITU-R BT.500 [4], very recent algorithms such as ZREC [23], RMLE [6], and finally the Netflix SUREAL software that implements the so-called “alternating projection” algorithm [5]. The latter has been recommended by the ITU in 2021 as the most comprehensive method for subjective quality recovery (as per Section 12.6 of ITU-R P.913 [20]).

The computational experiments were conducted using the data gathered in six different subjective tests. The related datasets are named: VQEG-HD1 [29], VQEG-HD3 [29], VQEG-HD5 [29], Netflix Public [1], KoNViD-1k [30] and the MoviesLens-1M [31]. While the last two datasets, i.e., the KoNViD-1k [30] and the MoviesLens-1M [31] were obtained from crowdsourcing subjective tests, the others are the results of highly controlled lab experiments. Notice that for crowdsourcing experiments, the matrix of opinion scores is typically sparse, as stimuli are evaluated only by a subset of subjects. Thus, we present results for crowdsourcing datasets in a separate section below.

For the three VQEG experiments, there were 24 participants and each of them rated around 168 stimuli, yielding for each of the three tests, a total of 24×168 opinion scores to be analyzed. The Netflix Public dataset is a relatively small-scale dataset, which includes the opinion scores of 26 subjects on the perceptual quality of 70 processed video sequences and nine source content. The KoNViD-1k subjective test involves 624 participants, who have scored 1200 short video sequences. For the MovieLens-1M, 6040 subjects have expressed their opinion score on 3952 movies. For all the six datasets considered for our experiments, the authors made use of five-point quality scales when gathering the opinion scores from the subjects. Hence, the opinion scores in each dataset range from 1 to 5.

B. Uncertainty of quality estimates

A typical approach to measure the uncertainty of the subjective quality recovered by a given method consists in computing the size of confidence intervals [5], [7], [23]. The larger the CI, the higher the uncertainty on the recovered subjective quality.

Table I shows the comparison between the average size of the CIs of the recovered subjective quality by each method on the four datasets resulting from tests performed in controlled environments. The percentages reported between parenthesis indicate by how much the application of each method reduced on average the size of the CIs that can be computed from the

TABLE I

UNCERTAINTY OF QUALITY ESTIMATES: COMPARISON OF THE SIZE OF CIs ESTIMATED BY THE DIFFERENT QUALITY RECOVERY APPROACHES. PERCENTAGES INDICATE RELATIVE SIZE OF THE CIs WITH RESPECT TO MOS CIs.

Methods	AVG CI SIZE			
	NETF PUB	VQ-HD1	VQ-HD3	VQ-HD5
MOS	0.509 (—)	0.493 (—)	0.565 (—)	0.575 (—)
BT500	0.515 (+1.18%)	0.613 (+24.34%)	0.586 (+3.72%)	0.575 (+0.00%)
ZREC	0.417 (-18.07%)	0.437 (-11.36%)	0.458 (-18.94%)	0.475 (-17.39%)
SUREAL	0.445 (-12.57%)	0.459 (-6.90%)	0.481 (-14.87%)	0.489 (-14.96%)
RMLE	0.453 (-11.00%)	0.417 (-15.42%)	0.472 (-16.46%)	0.483 (-16.00%)
ESQR	0.355 (-30.26%)	0.361 (-26.77%)	0.436 (-22.83%)	0.439 (-23.65%)

raw opinion scores, i.e. the MOS’s CIs (computed for each stimulus i as $\text{MOS} \pm 1.96 \times \text{SOS} / \sqrt{|\mathcal{I}_i|}$, where SOS stands for Standard deviation of Opinion Scores). For instance, as it can be seen in Table I, the average of the sizes of the MOS’s CIs on the Netflix public dataset is 0.509, while the average of the sizes of the CIs of the recovered qualities by the proposed ESQR algorithm is 0.355. Hence, by using the proposed ESQR algorithm instead of the MOS on the Netflix public dataset, on average, the size of the CIs of the recovered subjective qualities is reduced by 30%, i.e., $100 \times (1 - 0.355/0.509)$.

Looking at the results in Table I, it can be noticed that on all datasets, the proposed ESQR algorithm always recovered subjective qualities characterized by smaller CIs than those of all the other approaches on average. Hence, in practice, the proposed algorithm is expected to provide estimates of the subjective quality that are prone to lower uncertainty.

It is interesting to notice that more recent approaches such as ZREC, RMLE and the Netflix SUREAL software offered better performances than the MOS and the algorithm proposed in the ITU-R BT.500. The performances of the Netflix SUREAL software, RMLE and ZREC were however outperformed by that of the proposed ESQR algorithm. In fact, the application of the ESQR algorithm has yielded a reduction of the size of CIs by more than 22% in all cases, while all the other approaches never did better than 19%. The CIs resulting from the output of the algorithm proposed in the ITU-R BT.500 are indeed larger than the MOS’s CIs on average. This is actually not a peculiarity of this work as the same observation was made in [5].

1) *CIs Prediction Accuracy*: When comparing confidence interval sizes, a natural question arises: does a smaller confidence interval actually imply reduced uncertainty, or is it merely a result of underestimating the true uncertainty linked to the quality estimation? In real datasets, there are no ground-truth CIs against which estimated CIs can be benchmarked. Therefore, we must resort to simulations to verify the accuracy the CI estimates of ESQR and competing methods.

We simulated the opinion scores of 25 subjects for 100 stimuli. For each stimulus i , we assumed that reliable opinion scores on its quality follow a normal distribution with a mean of q_i and a standard deviation of σ_i . Consequently, q_i represents the ground truth quality for stimulus i . The ground truth CI of the quality of stimulus i is then:

$$\text{CI}_i = q_i \pm 1.96 \frac{\sigma_i}{\sqrt{M_s}}, \quad (16)$$

where $M_s = 25$ is the number of simulated opinion scores

TABLE II
CI PREDICTION ACCURACY

Method	MOS	BT500	ZREC	SUREAL	RMLE	ESQR
Δ^m	0.127	0.062	0.058	0.051	0.087	0.056
ρ^m	1.470	1.263	1.242	1.242	1.256	0.979

for each stimulus. The ground truth quality values q_i were derived by uniformly sampling 100 numbers within the range of [1.5, 4.5]. To simulate the fact that subjects exhibit lower inconsistency at the quality scale’s extremes, as observed in real subjective experiments [25], we employed the SOS hypothesis [25]. More precisely, we set $\sigma_i = 0.2 \times (-q_i^2 + 6q_i - 5)$, ensuring that the standard deviation of the distribution of reliable opinion scores diminishes at the quality scale’s extremes.

We will denote $N(q_i, \sigma_i)$ as the distribution of reliable opinion scores for the stimulus i . In our simulation, each stimulus is assessed by 25 subjects. We followed the scoring model proposed in [3], where each subject could provide a reliable opinion score with a probability of $1 - \eta$ and an unreliable one with probability η . We divided the 25 subjects into two clusters, i.e., a group of 20 accurate subjects, and a group of 5 inaccurate ones. For the accurate subjects, we set $\eta = 0.01$ (1%), meaning that 99% of their opinion scores were sampled from the distribution $N(q_i, \sigma_i)$ of reliable opinion scores and rounded to the closest integer from 1 and 5, while the remaining 1% were randomly selected between 1 and 5. The 5 inaccurate subjects had η randomly chosen between 0.6 and 1, meaning that at least 60% of their opinion scores were randomly selected between 1 and 5, and the rest were drawn from $N(q_i, \sigma_i)$. We conducted this simulation with 30 different seeds, resulting in 30 distinct simulated datasets.

We applied all quality recovery methods to each of the simulated datasets. Let \hat{C}_{id}^m represent the CI estimated by method m for stimulus i in simulated dataset d . To evaluate the accuracy of method m in estimating the ground truth CIs, we compared \hat{C}_{id}^m to C_i using two main indices:

$$\Delta^m = \frac{1}{30 \times 100} \sum_{d=1}^{30} \sum_{i=1}^{100} |ct(\hat{C}_{id}^m) - ct(C_i)| \quad (17)$$

$$\rho^m = \frac{1}{30 \times 100} \sum_{d=1}^{30} \sum_{i=1}^{100} sz(\hat{C}_{id}^m) / sz(C_i) \quad (18)$$

where $ct()$ and $sz()$ stand for center and size of the CI respectively. Δ^m is, therefore, the average distance between the center of the estimated CI of method m and the center of the ground truth CI. Meanwhile, ρ^m is the average ratio between the size of the CI estimated by method m and the size of the ground truth CI. Clearly, the closer Δ^m is to zero, the better; and the closer ρ^m is to 1, the better.

Table II summarizes the results. Regarding Δ^m , the best methods are ESQR, ZREC, and Netflix Soreal software, with the center of the estimated CI differing from that of the ground truth CI by around 0.05. The MOS exhibited the lowest performance ($\Delta^m = 0.127$), followed by RMLE ($\Delta^m = 0.087$). When it comes to predicting CI sizes, ESQR outperformed all other approaches, with a related value of $\rho^m = 0.979$ significantly closer to 1 than that of all the other methods

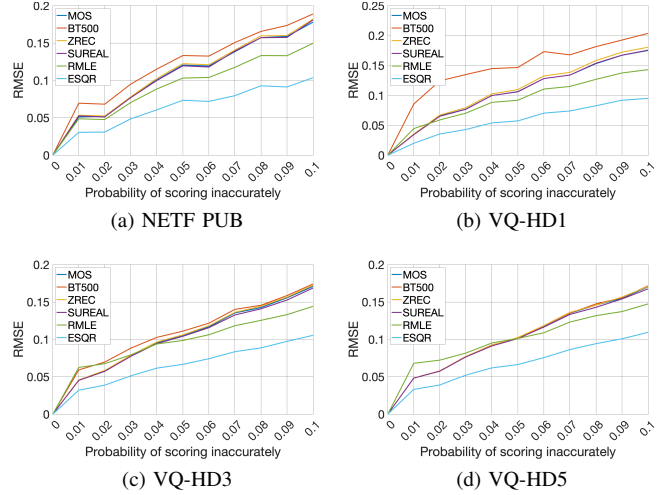


Fig. 2. **Robustness to noise insertion.** RMSE between the quality recovered on the original dataset and under different noisy conditions. The noise was added by replacing a given fraction of the opinion scores (see the x -axis) of each subject with integer numbers sampled at random between 1 and 5. The simulation was run with 30 different seeds and the curve for each quality recovery method reports the average RMSE from the 30 trials.

for which $\rho^m > 1.242$. Thus, ESQR slightly underestimated (by less than 3% of the actual size) the sizes of the ground truth CIs on average, while all the other methods overestimated them significantly (by more than 24% of the actual size).

The results in Table II suggest that the proposed ESQR algorithm can better predict the ground truth CIs and thus better quantify the actual uncertainty characterizing the quality of a stimulus compared to the other quality recovery approaches. We believe this stems from the fact that ESQR makes no restrictive assumptions about the subjects’ scoring behavior.

C. Robustness to Synthetic Noise

Following the same protocol of [1], [5]–[7], [23], all the quality recovery methods are first used to recover the subjective media quality on each dataset. Then, some synthetic noise is added to each dataset. After adding the noise, the subjective quality is estimated again, this time using the noisy dataset. Finally, the recovered quality on the noisy dataset is compared in terms of RMSE to the one obtained before adding noise to the dataset. This allows us to determine which method is more robust to noise.

The noise is added to the datasets using two different approaches that simulate different applications:

- 1) **Noise insertion:** a small fraction of the opinion scores of each subject is replaced by an integer sampled at random in the interval [1, 5];
- 2) **Spammer annotators:** some simulated subjects scoring the quality at random are added to the dataset.

In practice, our first approach to adding noise simulates, for instance, the type of noise that would be generated by the subjects’ fatigue or unexpected subjects’ distraction. In fact, it is reasonable to assume that, due to fatigue or distraction, each subject might inaccurately score the quality of a very small fraction of stimuli. The second noise model simulates situations such as the unexpected crash of the software used

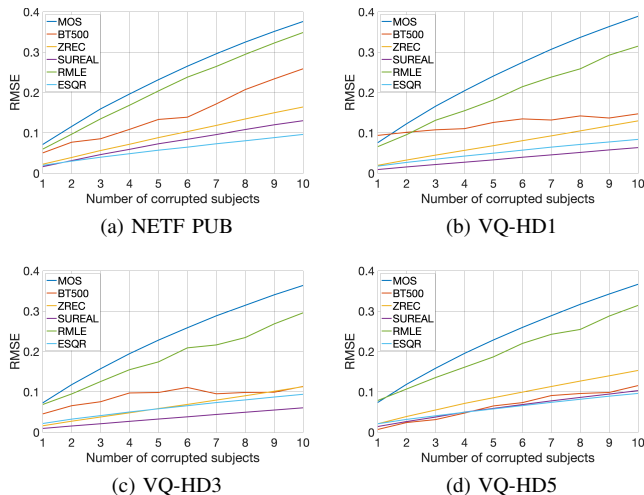


Fig. 3. **Robustness to spammer annotators.** RMSE between the quality recovered on the original dataset and under noisy conditions. The noise was generated by adding simulated subjects (see the x -axis) that score the quality of each stimulus with an integer number sampled at random between 1 and 5. The simulation was run with 30 different seeds and RMSE of the 30 trials for each quality recovery method is shown.

to collect opinion scores, causing a mismatch of the opinion scores of certain subjects (see for instance the Netflix subjective test described in [1]). In that case, the subjects whose opinion scores have been mismatched correspond to subjects rating at random. Another application is the case in which a subject accepts to participate in the subjective test but just provides ratings at random in order to complete the test as quickly as possible. These subjects are referred to in the literature as spammer annotators [32].

Figure 2 reports the results obtained using the first noise model. The fraction of replaced opinion scores is reported on the x -axis. The y -axis reports, for each quality recovery method, the RMSE error between the quality recovered on the original dataset and the one obtained after replacing the fraction of opinion scores on the x -axis with random integers. For instance, looking at the VQEG-HD1 dataset in Figure 2, the RMSE between the values of the recovered qualities by the proposed ESQR algorithm on the original dataset and the values computed after replacing 4% of the opinion scores of each subject in the dataset with random integers is 0.06.

In this first case, as it can be seen from Figure 2, the proposed ESQR algorithm outperformed all the other methods in all testing conditions. In fact, the curve of RMSE values associated to the ESQR algorithm lies below the ones of all the other methods. This result suggests that the proposed ESQR algorithm would guarantee better robustness than the other quality recovery methods to the noise generated for instance by the subjects' fatigue.

One can notice that the Netflix SUREAL software and ZREC showed performances very similar to that of the MOS. The RMLE approach instead showed better performance than the MOS, the BT.500 algorithm, ZREC and the Netflix SUREAL software.

The results related to our second approach to add the noise are summarized in Figure 3. The x -axis reports the number of simulated subjects added to the dataset. These simulated

subjects rate the stimuli by choosing at random an integer between 1 and 5. In this case, the proposed ESQR algorithm, ZREC and the Netflix SUREAL software showed comparable performances. In particular ESQR delivered the best performance on the Netflix public dataset and provided similar performance to that of the Netflix SUREAL software on the VQEG-HD5 dataset. The Netflix SUREAL software showed a better performance than the proposed ESQR algorithm on the VQEG-HD1 and VQEG-HD3 datasets. The RMLE approach showed a performance that is higher than that of the MOS but significantly lower than those of all the other methods.

The joint analysis of Figure 2 and Figure 3 reveals a crucial observation: the proposed ESQR algorithm demonstrates robustness, irrespective of the noise simulation approach. Unlike other quality recovery methods tested, ESQR performance does not exhibit sensitivity to the specific noise simulation employed. For instance, in Figure 2, sophisticated methods like ZREC and the Netflix SUREAL software do not outperform the MOS, yet they perform well in the second noise simulation case, as depicted in Figure 3. Conversely, RMLE excels in the first case in Figure 2 but falters in the second case in Figure 3. ESQR consistently maintains high performance across both cases, standing out as the top performer in the first noise simulation case and one of the best methods in the second case.

This stability in ESQR's performance during transitions between scenarios can be attributed to our avoidance of assumptions about subjects' scoring behavior, a characteristic of parametric approaches. Such assumptions often face challenges due to specific application characteristics. For example, the parametric model in the Netflix SUREAL software assumes a subject permanently possesses bias and inconsistency, making it less effective in capturing the scoring behavior of subjects who only occasionally misjudge quality. This likely explains the similar performance of the Netflix SUREAL software to that of the MOS in the first noise simulation case.

D. Crowdsourcing experiments

This section evaluates the accuracy of ESQR when the matrix of ratings is sparse. This type of matrix is typically obtained from crowdsourcing tests where a very large number of stimuli is employed, but each subject is required to rate only a small subset of them. This yields a stimuli-to-subjects table with numerous empty cells and thus a sparse matrix of ratings.

When the matrix of ratings is sparse, the correlation between the ratings of each pair of subjects cannot always be calculated. This makes it difficult for ESQR to derive a more accurate estimate of the distribution p_{V_i} than the distribution of collected ratings. In this case, our implementation of ESQR considers the distribution \bar{p}_{V_i} (Eq. (11)) of gathered ratings during the test as the estimate of p_{V_i} for each stimulus i in order to compute the reliability of each individual opinion score.

In light of the results discussed in Section V-C, the quality recovery approaches with the most competitive performance with respect to the proposed ESQR algorithm in terms of robustness to noise are the RMLE (see Figure 2), ZREC and the

TABLE III
CROWDSOURCING EXPERIMENTS: CI SIZES OF ESQR VS NETFLIX, SUREAL, ZREC. PERCENTAGE REDUCTION WITH RESPECT TO MOS CI.

Methods	AVG CI SIZE	
	KoNViD-1k	MoviesLens-1M
SUREAL	0.326 (-15.76%)	0.203 (-15.25%)
ZREC	0.318 (-18.04%)	0.205 (-14.58%)
ESQR	0.289 (-25.51%)	0.195 (-18.75%)

Netflix SUREAL software (see Figure 3). The analysis in this section could have therefore been done by considering ZREC, the Netflix SUREAL software and the RMLE approach. Unfortunately the RMLE approach involves an optimization problem whose solution is computationally very demanding on large-scale datasets as the ones considered in this section. This made it impossible to perform the experiments with RMLE on these datasets in a reasonable amount of time. For this reason, we considered only ZREC and the Netflix SUREAL software.

The sizes of the CIs of the recovered subjective qualities by the proposed ESQR algorithm, ZREC and the Netflix SUREAL software are compared in Table III. As in the case of datasets collected in controlled environments, the proposed ESQR algorithm provided a recovered subjective quality prone to lower uncertainty, i.e. smaller CIs. The use of the Netflix SUREAL software and ZREC induced no more than 18% reduction of the size of the raw data CIs, while ESQR achieved 25%. The reduction percentages achieved by the Netflix SUREAL software are slightly higher than the percentages it reached on datasets obtained in controlled environments (see Table I). This is consistent with the fact that greater benefit can be expected from sophisticated quality recovery approaches when used on challenging datasets such as those derived from crowdsourcing subjective tests.

The reduction percentages obtained for the proposed ESQR algorithm in Table III, although being greater than those of the Netflix SUREAL software and ZREC, were in one case smaller than the ones in Table I, which were obtained on datasets collected in controlled environments. This is because the current version of the proposed ESQR algorithm to analyze a sparse matrix of ratings directly uses the distribution of gathered ratings to estimate the reliability of individual opinion scores. We strongly believe that, as for the Netflix SUREAL software, the application of the ESQR algorithm would bring larger benefits on crowdsourcing datasets if an approach to “clean” the distribution of collected opinion scores is employed as in the case of a plain matrix of ratings where pairwise correlations are used. Finding such an approach will thus be one of the main points for a future contribution.

We compared the proposed ESQR algorithm to the Netflix SUREAL software and ZREC in terms of robustness to synthetic noise and spammer annotators added to a sparse matrix of ratings. The results are shown in Figure 4 and Figure 5. As in Section V-C, the quality recovered on the original dataset was compared in terms of RMSE to the one recovered from a noisy version of each dataset. To add noise, a fraction (see the x-axis) of opinion scores was selected at random and replaced with integers randomly sampled between 1 and 5. The ratings of a spammer annotator are simulated

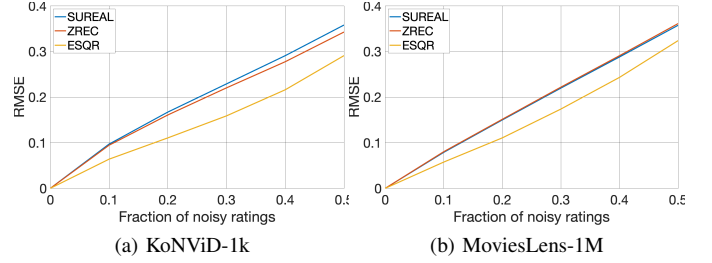


Fig. 4. Crowdsourcing experiments, robustness to noise insertion. Comparing the robustness to synthetic noise of the ESQR algorithm, the Netflix SUREAL software and ZREC on two crowdsourcing datasets. A given fraction of the opinion scores (see the x-axis) in each dataset was replaced with integers sampled at random between 1 and 5. The RMSE between the quality recovered on the original dataset and the one obtained in each noisy situation is reported on the y-axis.

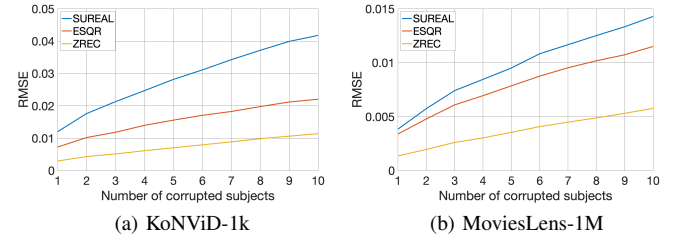


Fig. 5. Crowdsourcing experiments, robustness to spammer annotators. Comparing the robustness of the ESQR algorithm, the Netflix SUREAL software and ZREC to the insertion of spammer annotators on two crowdsourcing datasets. A certain number of spammer annotators (see the x-axis) was added to each dataset. The opinion scores of a spammer annotators are integers sampled at random between 1 and 5. The RMSE between the quality recovered on the original dataset and the one obtained after adding spammer annotators is reported on the y-axis.

by selecting random integers between 1 and 5. As we see in Figure 4, in the case of noise insertion, the proposed ESQR algorithm showed higher robustness to the added noise. In fact, it always recovered a subjective quality from the noisy dataset with the lowest RMSE with respect to the one obtained on the original version of the dataset. This suggests that, by adding additional noise to a challenging dataset, the proposed ESQR algorithm would offer more robustness to it than the Netflix SUREAL software and ZREC. For what concerns the insertion of spammer annotators (see Figure 5), ESQR shows better performance than the Netflix SUREAL software, but this performance is outperformed by that of ZREC.

E. Reliability of opinion scores

In this section we analyze the effect of the proposed reliability measure (Eq. (1)) as a weight of the contribution of each stimulus to the quality recovery. We compare the proposed reliability weights with the weights defined by Netflix SUREAL and ZREC. Let $\omega_{ij}^{\text{SUREAL}}$, $\omega_{ij}^{\text{ZREC}}$ and $\omega_{ij}^{\text{ESQR}}$ be the weights of the contribution of subject j to the determination of the quality of each stimulus i for SUREAL, ZREC and the proposed ESQR algorithm, respectively. Eq. (10) defines $\omega_{ij}^{\text{ESQR}}$, while SUREAL and ZREC define the contribution as:

$$\omega_{ij}^{\text{SUREAL}} = \frac{(\sigma_j^{\text{SUREAL}})^{-2}}{\sum_{k \in \mathcal{J}} (\sigma_k^{\text{SUREAL}})^{-2}} \quad (19)$$

$$\omega_{ij}^{\text{ZREC}} = \frac{(\sigma_j^{\text{ZREC}})^{-2}}{\sum_{k \in \mathcal{J}} (\sigma_k^{\text{ZREC}})^{-2}} \quad (20)$$

where σ_j^{SUREAL} and σ_j^{ZREC} are the estimated inconsistency of the subject j by the Netflix SUREAL software and ZREC respectively. Notice that, for SUREAL and ZREC, the weights only depend on the subject and are constant across stimuli evaluated by the same subject.

Eq. (19) and Eq. (20) reveal that ZREC and SUREAL weigh opinions similarly in determining ground truth quality. Consequently, we exclusively report results comparing ESQR to SUREAL in the following section, as similar conclusions arise with ZREC.

In Figure 6, we report $\omega_{ij}^{\text{SUREAL}}$ and $\omega_{ij}^{\text{ESQR}}$ computed on the Netflix public dataset. The ability of ESQR weights to modulate the importance of opinion scores *per stimulus* and not only per subject as SUREAL gives higher flexibility and precision in quality estimation. For instance, from the left heatmap in Figure 6 one can notice that, according to SUREAL, all the opinion scores of the subject #7 are considered unreliable. The heatmap of ESQR contradicts that by identifying stimuli for which the opinion scores of subject #7 are still accurate enough to contribute to estimating ground truth quality. In fact, for the following stimuli: #3, #43, #56 and #58 the ESQR contribution weights of subject #7 are rather high. A quick look at the dataset revealed that for these stimuli, subject #7 gave the opinion score chosen by the majority of subjects. Hence, contrarily to SUREAL, ESQR rightly attributed high importance to these opinion scores since they can be considered as accurate. Similarly, looking at the left heatmap in Figure 6, it can be observed that SUREAL attributes very high importance to all the opinion scores of subject #23. The ESQR heatmap however points out some stimuli for which the opinion scores of that subject are less reliable. For instance, subject #23 is the only one that scored the quality of stimulus #44 as being “bad” while all the other subjects found it at least “fair”.

Another interesting example is the situation of stimulus #19 for which all the subjects gave the same opinion scores. Despite all the subject agreed on one opinion score, SUREAL attributed different importance to the subjects when recovering the quality of that stimulus. This is not the case for ESQR that attributed the same contribution weight to all subjects as it can be seen from the right heatmap in Figure 6.

All these examples provide insights of why the proposed ESQR approach achieves in general better robustness and lower uncertainty than SUREAL and ZREC in the numerical experiments presented in this paper.

F. Quality estimates of ESQR vs. other methods

We evaluate the similarity between the subjective quality recovered by ESQR and that of other state-of-the-art quality estimation algorithms on the six considered datasets. ESQR estimates generally align with prior methods, deviating only in specific cases where assumptions are potentially violated. The smallest Pearson correlation found between ESQR and other methods is 0.996 (0.994 for Spearman correlation), indicating

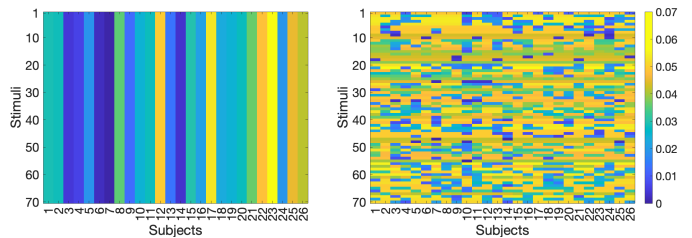


Fig. 6. **Reliability of opinion scores.** The figure shows the contribution weights $\omega_{ij}^{\text{SUREAL}}$ (left) and $\omega_{ij}^{\text{ESQR}}$ (right) of each subject j to the determination of the ground truth quality of each processed video sequence i in the Netflix public dataset. The ESQR weights can capture the reliability of individual opinion scores.

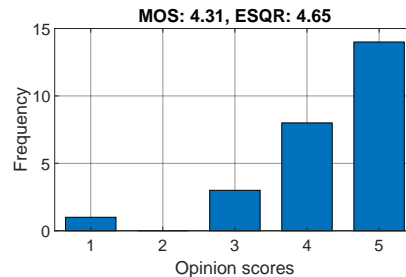


Fig. 7. **Quality estimates of ESQR vs. other methods.** Opinion scores distribution for a stimulus with high discrepancy between Q (ESQR) and the mean opinion score (MOS).

very high consistency on average. The RMSE further confirms these results, with a maximum value of 0.167, notably small on a 5-level quality scale.

It is instructive to analyze cases where the quality estimated by ESQR differs significantly from that of alternative methods. An example of stimulus where the output of ESQR deviates significantly from the MOS (difference = 0.34) is the one whose distribution of scores is showed in Figure 7. In fact, while 14 subjects out of 26 deemed that the quality of that stimulus was excellent and scored it with a 5, there is one subject that found the quality bad and gave a 1 as opinion score. The MOS attributes to this opinion score the same importance that is attributed to the other opinion scores; this yields a $\text{MOS} = 4.31$. The proposed ESQR algorithm instead under-weights that potentially noisy low opinion score and recovers a larger subjective quality (4.65).

To study significant differences between ESQR and the popular Netflix SUREAL, we compare the quality estimates by the two methods for the large-scale dataset MoviesLens-1M through the scatter plot in Figure 8a. We observe that the output of ESQR mostly differs from that of the Netflix SUREAL software at the extremes of the quality scale, i.e., where the distribution of the opinion scores is typically strongly asymmetric. In these cases, the Gaussianity assumption made by Netflix SUREAL concerning the quality distribution is not met, leading the method to introduce substantial bias in the estimates. We observed a similar behavior for KoNViD-1k.

Figure 8b reports the score distribution for the stimulus having the largest estimated quality difference between the ESQR algorithm and Netflix SUREAL. The Netflix SUREAL software recovered a subjective quality equal to the MOS. The ESQR algorithm instead recovered a subjective quality that is significantly lower than the MOS. The difference,

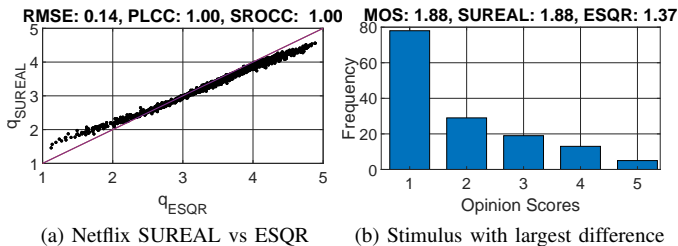


Fig. 8. **Quality estimates of ESQR vs. other methods.** Comparing the output of the ESQR algorithm to the one of the Netflix SUREAL software on the MovieLens-1M dataset. On the right, the distribution of opinion scores of the stimulus for which the outputs of the two approaches differ the most.

TABLE IV
ABLATION STUDY ON THE ESQR ALGORITHM

Index	ESQR	Without Pairwise Correlation	Without Reliability Measure
Rob. to Noise (RMSE)	0.06	0.06	0.10
Rob. to Spammer Annot. (RMSE)	0.06	0.12	0.04
Average Size of CIs	0.36	0.38	0.51
Per Stimulus Reliability	yes	yes	no

0.51, represents more than 10% of the whole quality scale. This confirms that, despite most of the time ESQR provides estimates coherent with previous methods, there are stimuli for which different quality recovery approaches strongly disagree.

G. Ablation Studies

In our ESQR algorithm ablation study we examined two aspects: i) removing pairwise correlation, thus directly using the observed distribution \bar{p}_{V_i} (Eq.(11)) instead of the estimate \hat{p}_{V_i} (Eq.(13)) that involves the weights ϵ_{ij} ; ii) eliminating the reliability measure, and thus using the expected value computed from \hat{p}_{V_i} as the subjective quality estimate. Results are shown in Table IV. We focus only on the Netflix Public dataset since similar conclusions were drawn from other datasets. To evaluate the robustness against noise and spammer annotators, we computed the average RMSE across the noisy scenarios in Figure 2 and Figure 3.

As it can be noticed from Table IV, without pairwise correlation, the algorithm is less robust to spammer annotators, and without the reliability measures, it's less robust to noise and computes a quality estimate with larger CIs and thus prone to more uncertainty. The full ESQR algorithm demonstrates superior performance balance, highlighting the importance of all introduced elements.

H. Computational Time Analysis

Each method was executed 30 times on each dataset, and the average computational time was recorded. These experiments were conducted using MATLAB on a computer equipped with a 2.6 GHz 6-Core Intel Core i7 processor and 16 GB of RAM.

Excluding RMLE, all approaches processed the small-scale datasets (VQEG-HD1, VQEG-HD3, VQEG-HD5, and Netflix Public) in less than 4 ms on average. Specifically, Netflix SUREAL and ZREC took less than 1 ms each, while ESQR required just over 3 ms. A notable increase in computational

time (up to 70 seconds) was observed with RMLE on small-scale datasets. This highlights the computational demands of parametric methods in estimating optimal parameter values. The efficient processing by SUREAL and ZREC is attributed to the approximation of the parameter estimation process through an iterative procedure and utilization of statistical moments, respectively.

Finally, it is interesting to note that, despite being slightly slower than ZREC and the Netflix SUREAL software, ESQR completed the recovery of the subjective quality on the MovieLens-1M dataset, involving up to 1 million ratings, in no more than 12 seconds. This clearly suggests that the efficiency of ESQR is not questionable for practical exigencies.

VI. CONCLUSIONS

In this paper, we introduce ESQR, a novel Entropy-based Subjective Quality Recovery algorithm to estimate subjective media quality from noisy opinion scores. The primary idea behind our approach is to treat quality estimation as a non-parametric problem, diverging from the prevalent practice in the literature that involves modeling scoring behavior through predefined and often simplistic distributions. Specifically, we establish a reliability measure for each stimulus capturing the degree of surprise that a given score brings compared to the overall score distribution. We then utilize this measure to weigh the contribution of individual opinion scores to the overall quality of a stimulus.

When comparing ESQR to five state-of-the-art quality recovery methods across six diverse datasets, our results indicate that: i) ESQR produces subjective quality estimates characterized by reduced uncertainty; ii) ESQR demonstrates superior robustness to noise compared to other methods; and iii) ESQR maintains its accuracy across a broader range of applications and datasets.

Future work will explore a refinement of the introduced reliability measure to explicitly consider the ordinal nature of quality scales, as entropy-based approaches may overlook this essential aspect.

APPENDIX A

ASYMPTOTIC DISTRIBUTION OF QUALITY ESTIMATOR

Consider a stimulus that has been scored independently by M subjects. The subject j has given a score $1 \leq R_j \leq K < \infty$ drawn from a distribution p_{R_j} . Based on the scores $\{R_1, \dots, R_M\}$ a quality recovery algorithm is proposed as a weighted average:

$$Q(M) = \frac{\sum_{j=1}^M h(R_j)R_j}{\sum_{l=1}^M h(R_l)}. \quad (21)$$

In the ESQR algorithm we would have $h(R_j) = -\log(\hat{p}_V(R_j))^{-1}$ where \hat{p}_V is an approximation to the scoring distribution. In this section, to simplify our theoretical analysis we assume that we are in fact using the true scoring distribution p_V . Other algorithms (see Eq. (19)) have the same structure so our analysis can also be extended to those cases.

Proposition 2. *If there exists $c > 0$ such that $|h(R_j)| \leq c$ for all j , and as $M \rightarrow \infty$ for any $q \in [1, K]$ we have that:*

$$\sum_{j=1}^M \text{var} [h(R_j)(R_j - q)] \rightarrow \infty, \quad (22)$$

then the asymptotic distribution of Q is:

$$\lim_{M \rightarrow \infty} \mathbb{P}(Q \leq q) = \Phi \left(-\frac{\mu_q}{\sigma_q} \right), \quad (23)$$

where Φ is the distribution function of a standard normal random variable and:

$$\mu_q = \sum_{j=1}^M \mathbb{E} [h(R_j)(R_j - q)], \quad (24)$$

$$\sigma_q^2 = \sum_{j=1}^M \text{var} [h(R_j)(R_j - q)]. \quad (25)$$

Proof. Using (21) we may write the distribution of Q as:

$$F_Q(q) = \mathbb{P}(Q \leq q) = \mathbb{P} \left(\sum_{j=1}^M h(R_j)(R_j - q) \leq 0 \right). \quad (26)$$

Since $|h(R_j)| \leq c$, for some c , and $q \in [1, K]$ then the random variables in the summation are uniformly bounded for all j as $|h(R_j)(R_j - q)| \leq c(K - 1)$. Under (22) we can now apply Lindeberg's central limit theorem (CLT) [33, Example 27.4] to obtain the desired result. \square

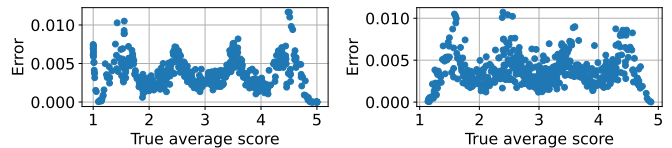
Notice that event though we have a closed form approximation for the distribution of Q using the the CLT we cannot guarantee that Q will be normally distributed since μ_q and σ_q are functions of q . The approximation will yield a normal distribution if and only if μ_q/σ_q is a linear function of q . By looking at (24) we see that μ_q is indeed a linear function of q . Then we have the following corollary:

Corollary 1. *Under the hypotheses of Proposition 2, Q is asymptotically normal if and only if (25) is independent of q .*

It is clear that a good approximation will be retained as long as μ_q/σ_q is approximately linear in q where Φ changes more rapidly. We now perform some numerical simulations to study whether the quality estimate can be assumed to be normal. In order to do this, we need to test two things:

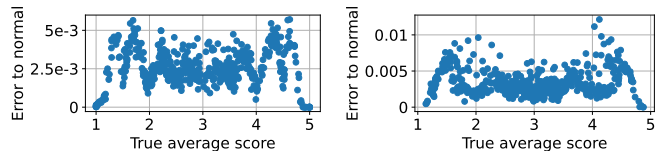
- T1) Using the CLT is a good approximation for F_Q for the number M of subjects typically considered. This would validate that (23) is a good approximation for practical finite values of M .
- T2) The argument of (23) is a linear function of q where Φ changes rapidly, which, together with T1) would validate that the estimator is approximately normal in practice.

For the tests, we perform simulation with $M = 24$ subjects. We consider a model very similar to that of section V-B. We assume that each subject independently rates the same stimulus, giving a score on $\{1, 2, 3, 4, 5\}$. The scores is given according to the true distribution p_V with probability $1 - p_e$, and a uniform score with probability p_e . The probability p_e , different for each subject, is obtained as a uniform random



(a) Normal. Average L1 error is 0.004. (b) Beta. Average L1 error is 0.004.

Fig. 9. **Test T1.** Scatter plot of the L^1 error between (23) for finite M and the empirical estimation of Q as a function of the true average score.



(a) Normal. Average L1 error is 0.003. (b) Beta. Average L1 error is 0.003.

Fig. 10. **Test T2.** Scatter plot of the L^1 error between (23) for finite M and a normal distribution with the same mean and variance as Q , as a function of the true average score.

variable on $(0, 0.05)$. The distribution p_V is obtained by discretizing continuous distributions, namely:

- A normal random variable of mean x_e and deviation a , where x_e is drawn from a continuous uniform random variable on $(1, 5)$ and $a = 0.2 \times (-x_e^2 + 6x_e - 5)$ [25]. The discretization is done considering the points $\{1.5, 2.5, 3.5, 4.5\}$ of the normal variable.
- A beta random variable with parameters a and b drawn as independent continuous uniform variables on $(1, 10)$. The discretization is done by dividing the support $(0, 1)$ into 5 consecutive equally spaced intervals.

To test the validity of T1) and T2) we proceed as follows:

- Let \mathbf{q} be a uniform grid of n_d points in $[1, 5]$, that is:

$$\mathbf{q} = \left\{ \frac{4i + (n_d - 5)}{n_d - 1} : i = 1, \dots, n_d \right\}. \quad (27)$$

These points are where F_Q will be estimated.

- Choose an input distribution for V , normal or beta.
- Do the following experiment n_p times:
 - 1) Choose the parameters θ for the true score, with $\theta = (x_e, a)$ for normal and $\theta = (a, b)$ for beta.
 - 2) For each $q \in \mathbf{q}$ obtain estimates $\hat{\mu}_q$ and $\hat{\sigma}_q$ of (24) and (25), respectively, through samples averages using $n_s = 5000$ independent realizations of the scores of the M subjects. Then obtain the estimate of the CDF of Q using the CLT as: $\hat{F}_{Q, \text{CLT}}(q) = \Phi \left(\frac{\hat{\mu}_q}{\hat{\sigma}_q} \right)$.
 - 3) For each $q \in \mathbf{q}$ obtain an independent estimate of the distribution of Q using n_s independent realizations of the scores as:

$$\hat{F}_{Q, \text{emp}}(q) = \frac{\#\text{Samples of } Q \leq q}{n_s}. \quad (28)$$

Also, compute the sample mean $\hat{\mu}_{Q, \text{emp}}$ and sample variance $\hat{\sigma}_{Q, \text{emp}}^2$.

- 4) **Computations for T1:** $\hat{F}_{Q, \text{CLT}}$ uses Prop. 2 (the CLT) to approximate the distribution of Q , while $\hat{F}_{Q, \text{emp}}$ does not make any modeling assumptions. If the CLT is a good approximation, then both estimators should give similar

values, which validates T1). To check this, we estimate this error using the L^1 norm of the error between the two estimators by using the discrete samples in \mathbf{q} . To do this we compute the error: $e_i = \left| \hat{F}_{Q,CLT}(q_i) - \hat{F}_{Q,emp}(q_i) \right|$ where $q_i = \frac{4i+(n_q-5)}{n_q-1}$ and estimate the L^1 error between the two estimators using the trapezoidal rule. where Δ is the spacing between two values of q .

- 5) **Computations for T2:** validating T1, this does not mean that Q is approximately normal. To verify this we compare $\hat{F}_{Q,CLT}$ with the distribution of a normal with mean $\hat{\mu}_{Q,emp}$ and variance $\hat{\sigma}_{Q,emp}^2$ by computing the L^1 in the same manner as with $F_{Q,emp}$.

After the n_p repetitions we have computed the n_p estimates of $\hat{F}_{Q,CLT}$, $\hat{F}_{Q,emp}$, and the distribution of a normal with mean $\hat{\mu}_{Q,emp}$ and variance $\hat{\sigma}_{Q,emp}^2$, for different parameters of the input distribution. We also computed the L^1 error between the $\hat{F}_{Q,CLT}$ and the other two estimates. If both errors are small then T1 and T2 are validated, which means that, at least for the proposed distributions, the CLT is a good approximation for values as small as $M = 24$ subjects and that Q is approximately normal.

In Fig. 9 we can see the scatter plot of the L1 error of the true distribution of Q and $\hat{F}_{Q,CLT}$, the approximation (23) using the CLT for finite M , for $n_p = 500$ repetitions and the three score distributions. We see that the total worst case error is very small, below 0.012, for all the possible true scores.

In Fig. 10 we see the scatter plot of the L1 error between $\hat{F}_{Q,CLT}$ and a normal distribution with the same mean and variance as Q . Again for all the results the worst error is very small, around 0.006 for Gaussian and 0.012 for the Beta.

REFERENCES

- [1] Z. Li and C. G. Bampis, "Recover Subjective Quality Scores from Noisy Measurements," in *Proc. Data Compression Conference (DCC)*. Snowbird, UT, USA: IEEE, Apr. 2017, pp. 52–61.
- [2] S. Pezzulli, M. G. Martini, and N. Barman, "Estimation of Quality Scores From Subjective Tests-Beyond Subjects' MOS," *IEEE Transactions on Multimedia*, vol. 23, pp. 2505–2519, 2021.
- [3] J. Li, S. Ling, J. Wang, and P. Le Callet, "A Probabilistic Graphical Model for Analyzing the Subjective Visual Quality Assessment Data from Crowdsourcing," in *Proc. 28th Intl. Conf. on Multimedia*. Seattle, WA, USA: ACM, Oct. 2020, pp. 3339–3347.
- [4] ITU-T, "Methodology for the subjective assessment of the quality of television pictures," ITU-T Rec. BT.500, pp. 1–100, Oct. 2019.
- [5] Z. Li, C. G. Bampis, L. Janowski, and I. Katsavounidis, "A Simple Model for Subject Behavior in Subjective Experiments," *Proc. Intl. Symposium on Electronic Imaging*, vol. 32, no. 11, pp. 131–1–131–14, Jan. 2020.
- [6] L. Fotio Tiotsop, A. Servetti, M. Barkowsky, and E. Masala, "Regularized Maximum Likelihood Estimation of the Subjective Quality from Noisy Individual Ratings," in *Proc. 14th Intl. Conf. on Quality of Multimedia Experience (QoMEX)*. Lippstadt, Germany: IEEE, Sep. 2022, pp. 1–4.
- [7] L. Fotio Tiotsop, A. Servetti, and E. Masala, "A scoring model considering the variability of subjects' characteristics in subjective experiments," in *Proc. 15th Intl. Conf. on Quality of Multimedia Experience (QoMEX)*. Ghent, Belgium: IEEE, Jun. 2023, pp. 1–6.
- [8] L. Janowski and M. Pinson, "The Accuracy of Subjects in a Quality Experiment: A Theoretical Subject Model," *IEEE Transactions on Multimedia*, vol. 17, no. 12, pp. 2210–2224, Dec. 2015.
- [9] ITU-T, "Subjective evaluation of media quality using a crowdsourcing approach," ITU-T PSTR-CROWDS, pp. 1–17, May 2018.
- [10] H. Lin, G. Chen, M. Jenadeleh, V. Hosu, U.-D. Reips, R. Hamzaoui, and D. Saupe, "Large-scale crowdsourced subjective assessment of picturewise just noticeable difference," *IEEE transactions on circuits and systems for video technology*, vol. 32, no. 9, pp. 5859–5873, 2022.
- [11] Q. Huynh-Thu, M.-N. Garcia, F. Speranza, P. Corriveau, and A. Raake, "Study of Rating Scales for Subjective Quality Assessment of High-Definition Video," *IEEE Transactions on Broadcasting*, vol. 57, no. 1, pp. 1–14, Mar. 2011.
- [12] J.-S. Lee, L. Goldmann, and T. Ebrahimi, "Paired comparison-based subjective quality assessment of stereoscopic images," *Multimedia tools and applications*, vol. 67, pp. 31–48, 2013.
- [13] Q. Xu, J. Xiong, Q. Huang, and Y. Yao, "Online HodgeRank on Random Graphs for Crowdsourced QoE Evaluation," *IEEE Transactions on Multimedia*, vol. 16, no. 2, pp. 373–386, Feb. 2014.
- [14] Q. Xu, M. Yan, C. Huang, J. Xiong, Q. Huang, and Y. Yao, "Exploring Outliers in Crowdsourced Ranking for QoE," in *Proc. of 25th Intl. Conf. on Multimedia*. Mountain View, CA, USA: ACM, Oct. 2017, pp. 1540–1548.
- [15] J. Li, R. Mantiuk, J. Wang, S. Ling, and P. Le Callet, "Hybrid-mst: A hybrid active sampling strategy for pairwise preference aggregation," *Advances in neural information processing systems*, vol. 31, 2018.
- [16] E. Zerman, G. Valenzise, and A. Smolic, "Analysing the Impact of Cross-Content Pairs on Pairwise Comparison Scaling," in *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)*, Berlin, Germany, Jun. 2019, pp. 1–6.
- [17] M. Perez-Ortiz, A. Mikhailiuk, E. Zerman, V. Hulusic, G. Valenzise, and R. Mantiuk, "From Pairwise Comparisons and Rating to a Unified Quality Scale," *IEEE Transactions on Image Processing*, vol. 29, pp. 1139–1151, 2019.
- [18] E. Zerman, V. Hulusic, G. Valenzise, R. Mantiuk, and F. Dufaux, "The relation between MOS and pairwise comparisons and the importance of cross-content comparisons," in *Human Vision and Electronic Imaging Conference, IS&T International Symposium on Electronic Imaging (EI 2018)*, vol. 30, no. 14, Burlingame, United States, Jan. 2018, pp. 1–6.
- [19] ITU-T, "Subjective video quality assessment methods for multimedia applications," ITU-T Rec. P.910, pp. 1–46, Jul. 2022.
- [20] ITU-T, "Methods for the subjective assessment of video quality, audio quality and audiovisual quality of Internet video and distribution quality television in any environment," ITU-T Rec. P.913, pp. 1–42, Jun. 2021.
- [21] R. B. Millar, *Maximum likelihood estimation and inference: with examples in R, SAS and ADMB*. John Wiley & Sons, 2011.
- [22] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal processing magazine*, vol. 13, no. 6, pp. 47–60, 1996.
- [23] J. Zhu, A. Ak, P. Le Callet, S. Sethuraman, and K. Rahul, "Zrec: Robust recovery of mean and percentile opinion scores," in *2023 IEEE International Conference on Image Processing (ICIP)*, 2023, pp. 2630–2634.
- [24] Netflix, "The Soreal software," <https://github.com/Netflix/soreal>, (retrieved Dec 16th, 2022).
- [25] T. Hoßfeld, R. Schatz, and S. Egger, "SOS: The MOS is not enough!" in *Proc. 3rd Intl. Workshop on Quality of Multimedia Experience (QoMEX)*. Mechelen, Belgium: IEEE, Sep. 2011, pp. 131–136.
- [26] C. Kang, G. Valenzise, and F. Dufaux, "Predicting Subjectivity in Image Aesthetics Assessment," in *21st International Workshop on Multimedia Signal Processing (MMSP'2019)*. Kuala Lumpur, Malaysia: IEEE, Sep. 2019, pp. 1–6.
- [27] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience, 2006.
- [28] N. C. Silver and W. P. Dunlap, "Averaging correlation coefficients: Should fisher's z transformation be used?" *Journal of applied psychology*, vol. 72, no. 1, p. 146, 1987.
- [29] VQEG, "Report on the validation of video quality models for high definition video content (v. 2.0)," <https://bit.ly/2Z7GWDI>, (retrieved Dec 16th, 2022).
- [30] V. Hosu, F. Hahn, M. Jenadeleh, H. Lin, H. Men, T. Szirányi, S. Li, and D. Saupe, "The konstanz natural video database (konvid-1k)," in *2017 Ninth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2017, pp. 1–6.
- [31] F. M. Harper and J. A. Konstan, "The movielens datasets: History and context," *Acm transactions on interactive intelligent systems (tiis)*, vol. 5, no. 4, pp. 1–19, 2015.
- [32] Y. E. Kara, G. Genc, O. Aran, and L. Akarun, "Modeling annotator behaviors for crowd labeling," *Neurocomputing*, vol. 160, pp. 141–156, Jul. 2015.
- [33] P. Billingsley, *Probability and Measure*, 2nd ed. John Wiley and Sons, 1986.