Active Sampling for Pairwise Comparisons via Approximate Message Passing and Information Gain Maximization

Aliaksei Mikhailiuk*, Clifford Wilmot*, Maria Perez-Ortiz‡*, Dingcheng Yue*, Rafał K. Mantiuk*

*Department of Computer Science University of Cambridge Cambridge, United Kingdom
Email: am2442@cam.ac.uk, cliffw29@outlook.com, {dy276, rafal.mantiuk}@cam.ac.uk
‡Department of Computer Science, University College London, London, United Kingdom
Email: maria.perez@ucl.ac.uk

Abstract—Pairwise comparison data arise in many domains with subjective assessment experiments, for example in image and video quality assessment. In these experiments observers are asked to express a preference between two conditions. However, many pairwise comparison protocols require a large number of comparisons to infer accurate scores, which may be unfeasible when each comparison is time-consuming (e.g. videos) or expensive (e.g. medical imaging). This motivates the use of an active sampling algorithm that chooses only the most informative pairs for comparison. In this paper we propose ASAP, an active sampling algorithm based on approximate message passing and expected information gain maximization. Unlike most existing methods, which rely on partial updates of the posterior distribution, we are able to perform full updates and therefore much improve the accuracy of the inferred scores. The algorithm relies on three techniques for reducing computational cost: inference based on approximate message passing, selective evaluations of the information gain, and selecting pairs in a batch that forms a minimum spanning tree of the inverse of information gain. We demonstrate, with real and synthetic data, that ASAP offers the highest accuracy of inferred scores compared to the existing methods. We also provide an open-source GPU implementation of ASAP for large-scale experiments.

I. INTRODUCTION

The fields of subjective assessment and preference aggregation are concerned with measuring and modeling human judgments. Participants usually *rate* a set of stimuli or conditions according to some criteria, or *rank* a subset of them. Rating is inherently more complex for participants than ranking [1], [2], [3], [4]. Thus, comparative judgment experiments are gaining attention in subjective assessment and crowd-sourced experiments, e.g. for image quality assessement. The simplest form of ranking experiments is comparing conditions in pairs (pairwise comparison protocol), and hence it is the most common ranking choice. Here observers are asked to choose one out of two conditions according to some criteria. As opposed to rating, in which conditions are mapped directly to a scale by computing mean opinion scores, we need to model and infer the latent scores from pairwise comparisons. This problem is known as psychometric scaling. Models used for scaling typically rely on the assumptions of Thurstone's model [5] or Bradley-Terry's model [6]. The main limitation of

pairwise comparison experiments is that for n conditions there are $\binom{n}{2} = n(n-1)/2$ possible pairs to compare, which makes collecting all comparisons too costly for large n. However, *active sampling* can be used to select the most informative comparisons, minimizing experimental effort while maintaining accurate results.

The need for an efficient active sampling algorithm for preference aggregation is motivated by the recent spread of applications reliant on: i) user preferences (i.e. recommendation systems, information retrieval and relevance estimation) [7]; ii) matchmaking in gaming systems such as TrueSkill for Xbox Live [8] and Elo for chess and tennis tournaments [9]; iii) psychometric experiments for behavioural psychology [10] and iv) quality of experience (e.g. image and video quality) [11], [12], [13], [4].

State-of-the-art active sampling methods are typically based on information gain maximization [14], [15], [16], [4], [17], [18], where pairs in each trial are selected to maximize the weighted change of the posterior distribution of the scale. However, these are computationally expensive for a large number of conditions (n), as they require computing the posterior distribution for $n(n-1)/2$ pairs at every iteration of the algorithm. To make active sampling computationally feasible, most existing techniques update the posterior distribution only for the pairs that were selected for the next comparison. We show that this leads to a sub-optimal choice of pairs and worse accuracy as the number of measurements increases. To address this problem, we substantially reduce the computational cost of active sampling by using approximate message passing for inference, and by computing the expected information gain only for the subset of the most informative pairs. The reduced computational overhead allows us to update the full posterior distribution at every iteration, thus greatly improving the accuracy. To ensure balanced design and allow for a batch sampling mode, we sample the pairs from a minimum spanning tree as in [17]. The proposed technique (ASAP - Active SAmpling for Pairwise comparisons) results in the most accurate psychometric scale, especially for a large number of measurements. Moreover, the algorithm has

a structure that is easy to parallelize, allowing for a fast GPU implementation. We show the benefit of using full posterior update by comparing to an approximate version of the algorithm (ASAP-approx), that, similar to other methods relies on online posterior update. Our main contributions are: A) an analysis of existing active sampling methods for pairwise comparison experiments under a range of condition score distributions, using both synthetic and real image and video quality assessment data; B) a novel active sampling method (ASAP), offering the highest accuracy of the scale; and C) along with the paper we include an implementation of 9 algorithms[1], providing an open-source software for active sampling in pairwise comparison experiments and including the first GPU implementation of such a method.

II. RELATED WORK

Comparative judgment experiments arise in ranking (ordering conditions) and scaling applications (putting conditions on a scale where distances convey meaning). Suppose we aim to compare a set of n conditions $S = \{o_1, \ldots, o_n\}$ (conditions being images, players, etc.) that are evaluated according to a feature or characteristic (subjective measurements such as aesthetics, relevance, quality, etc.) with unknown underlying ground truth scores $\mathbf{s} = (s_1, \ldots, s_n)$, $s_i \in \mathbb{R}$. In this paper, we simply refer to these as quality scores. The simplest experimental protocol is to compare pairs (o_i, o_j), $o_i, o_j \in S$, $i \neq j$ (referred to as pairwise comparisons). Although other works exist, e.g. estimating total or partial order [19], [20], [21], [22], [23], this paper is focused on active sampling for psychometric scale construction, which uses pairwise comparisons to estimate quality scores $\hat{\mathbf{s}}$ that approximate $\mathbf{s}$. This section discusses related work, divided into four groups, based on the type of approach: passive, sorting, information-gain and matchmaking. The methods tested in the experiments are highlighted in bold face. We also distinguish between *sequential* methods —where the next pair is generated only upon receiving the outcome for the preceding pair —and batch, or parallel methods —where a batch of comparison pairs is generated and outcomes can be obtained in parallel. Batch methods are preferred in crowd-sourcing, where multiple conditions are distributed to participants in parallel.

a) Passive approaches: When every condition is compared to every other condition the same number of times, the experimental design is referred to as *full pairwise comparisons* (FPC). Such an approach is impractical, as it requires $n(n-1)/2$ comparisons per participant. Another approach, *nearest conditions* (NC), relies on the idea that conditions that are similar in quality are more informative for constructing the quality scale [24]. Thus, if the approximate ranking is known in advance, one can compare only the conditions that are neighbours in the ranking. Such initial ranking, however, may not be available in practice.

b) Sorting approaches: Similar to NC, sorting-based methods rank the conditions, then compare those that are of

similar quality. Authors in [25] proposed an active sampling algorithm using a binary tree. Every new condition descends down the tree, branching depending on whether it is better or worse than the condition in the current node. Authors in [26] applied **Quicksort** [27] using pairwise comparisons as the comparison operator.

Recently, [13] used the **Swiss system** in chess to rank subjective assessment of visual quality. The Swiss system first chooses random conditions to compare, then sorts the conditions to find pairs that are similar. A related method is the Adaptive Rectangular Design (ARD) [28] which allows comparison of conditions far apart on the quality scale in later stages of an experiment. The work of [29] takes a different approach, where active sampling (**AKG**) is based on the Bayesian decision process maximising Kendall's tau rank correlation coefficient [30].

Sorting approaches are praised for their simplicity and low computational complexity and are thus often employed in practice. However, these approaches use heuristics that often result in suboptimal comparison choices, and in general perform worse than the methods that rely on information gain.

c) Information-gain approaches: These methods are based on information maximization. That is, the posterior distribution of quality scores is computed and the next comparison is selected according to a utility function, e.g. Kullback-Leibler (KL) divergence [31] between the current distribution and the distribution assuming any possible comparison [32]. This group is the most relevant to our new method. Methods listed in this section are sequential, unless stated otherwise.

A greedy Bayesian approach, **Crowd-BT**, was proposed in [16]. The entropy for every pair of conditions is computed using the posterior distribution of each pair individually rather than jointly. The method also explicitly accounts for reliability of each annotator: scores and annotator quality are updated using an alternating optimization strategy.

Authors in [14] derive the score distribution from the maximum likelihood estimation and the negative inverse of the Hessian of the log likelihood function. Since the original implementation was not provided by the authors and our implementation suffered from numerical instability, we did not include it in our tests.

Authors in [15], [4] develop a fully Bayesian framework to compute the posterior distribution of the quality scores. **Hybrid-MST** [17] extends this idea by selecting batches of comparisons (instead of single pairs) to maximize the information gain in the minimum spanning tree [33] of a comparison graph. The time efficiency of the method over its predecessor is improved by computing the information gain locally —within the compared pair.

A different approach is taken by [18], where authors propose to solve a least-squares problem to elicit a latent global rating of the conditions using the Hodge decomposition of pairwise comparison data. Like other methods, the information gain is computed using the posterior of only the pair of compared conditions. We refer to this approach as **HR-active**.

d) Matchmaking: A matchmaking system was proposed for gaming, together with the TrueSkill algorithm [8]. The aim is to find the pairs of players with the most similar skill. The skill distribution of a pair of players is used to predict the match outcome. We refer to this approach as **TS-sampling**.

e) Our Work: In contrast to the previous work, our method (i) allows for batch and sequential modes; (ii) estimates the posterior using the entire set of comparison outcomes that has been collected so far; and (iii) computes the utility function for a subset of pairs, saving computations without compromising on performance.

III. Methodology

Our algorithm consists of two main steps: (i) computing the posterior distribution of score variables $\mathbf{r}$ using the pairwise comparisons collected; (ii) using the posterior of $\mathbf{r}$ to estimate the next comparison to be performed. In this section we first describe the score posterior estimation and then explain our active sampling algorithm. We then discuss some features to make it more computationally efficient. Pseudo-code is included in the supplementary.

A. Posterior Estimation

a) Posterior Estimation Model: Our model is similar to Thurstone's model Case V [5], with unobserved normally distributed independent random variables. However, our approach is fully Bayesian, and so instead of point value scores s_i for each condition o_i, we assume that each score is a random variable r_i with distribution $r_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$. Analogous to Thurstone's model, μ_i represents the score value s_i. σ_i^2 represents the uncertainty in an estimate of s_i and is not explicitly expressed in Thurstone's model (it can be obtained, for example by bootstrapping [34]). The probability that o_i is better than o_j is then given by noting that $r_i - r_j \sim \mathcal{N}(\mu_i - \mu_j, \sigma_{ij}^2)$, so that:

$$P(o_i \succ o_j | r_i, r_j) \triangleq \Phi\left(\frac{\mu_i - \mu_j}{\sqrt{2}\sigma_{ij}}\right), \tag{1}$$

where Φ is the cumulative standard normal distribution function and $\sigma_{ij}^2 = \sigma_i^2 + \sigma_j^2 + \beta^2$, with β representing what is referred in the literature to as an observer/comparison noise. We further assume Thurstone Case V model in which β is constant across all conditions. The choice of β determines the relationship between distances in the scale and probabilities of better quality. In our experiments we set $\beta = 1$.

For a pair of compared conditions $A_t = (o_i, o_j)$ for $t \in \{1, \ldots, T\}$, where T is the total number of comparisons measured so far, we denote the comparison outcome as $y_t \in \{-1, 1\}$, where 1 indicates that o_i was preferred and -1 indicates that o_j was preferred, with no draws allowed. In the inference step, we want to estimate the distribution of score variables $\mathbf{r}$ given $\mathbf{y}$ and $\mathbf{A} \triangleq \{A_1, \ldots, A_T\}$. The posterior distribution is:

$$P(\mathbf{r}|\mathbf{y}, \mathbf{A}) = \frac{P(\mathbf{y}, \mathbf{A}|\mathbf{r}) \cdot p(\mathbf{r})}{P(\mathbf{y}|\mathbf{A})}, \tag{2}$$
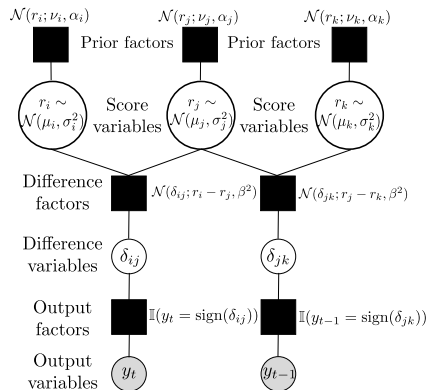

Fig. 1: Factor graph for 2 comparisons of 3 conditions.

where we assume a factorizing Gaussian prior distribution over scores $p(\mathbf{r}) \triangleq \prod_{i=1}^{n} \mathcal{N}(r_i; \nu_i, \alpha_i^2)$, ν_i and α_i^2 being the parameters of the prior, set to 0 and 0.5, respectively. The likelihood $P(\mathbf{y}, \mathbf{A}|\mathbf{r})$ of observing comparison outcomes $\mathbf{y}$ given the ground truth scores is modelled as:

$$P(\mathbf{y}, A|\mathbf{r}) = \prod_{t=1}^{T} P(y_t, A_t|\mathbf{r}), \tag{3}$$

where individual likelihoods can be defined as $P(y_t, A_t|\mathbf{r}) = \mathbb{I}(y_t = \text{sign}(r_i - r_j))$, i.e. equal to 1 if the sign of y_t is the same as that of the difference $r_i - r_j$ and 0 otherwise.

Although the score posterior can be written exactly via Bayes rule, the binary nature of the output factor means that the likelihood in Eq. 3 is not conjugate to the Gaussian prior. This would lead to a non-Gaussian posterior for r_i, and result in challenging, high-dimensional integrals for our information gain metric. A Gaussian approximation to messages yields a multivariate Gaussian posterior with diagonal covariance matrix, resolving both issues.

b) Posterior Estimation Inference: Figure 1 shows a factor graph implementing the distribution $P(\mathbf{r}|\mathbf{y}, A)$, used as the basis for efficient inference, and inspired by TrueSkill [8]. The posterior over r_i is inferred via message passing between nodes on the graph, with messages computed using the sum-product algorithm. In the general case of n conditions and T comparisons, we will have n score variables and prior factors, T difference factors, difference variables, output factors and output variables. Messages from output factors are approximated as Gaussians using expectation propagation via moment matching.

B. Sampling Algorithm: ASAP

The basis of the proposed active sampling algorithm is to compute the posterior distribution over $\mathbf{r}$ that would arise from each possible pairwise outcome in the next comparison, and then use this to choose the next comparison based on a criterion of maximum information gain.

Several utility functions can be used to compute the expected information gain (EIG). Our choice is the commonly

used Kullback-Leibler (KL) divergence [31] between the prior and posterior distributions.

More specifically, our active sampling strategy picks conditions $(o_i, o_j) = A_t$ to compare in measurement t, such that they maximize a measure of information gain I_{t-1}^{ij}:

$$A_t = \underset{(o_i, o_j) \in S^2, i \neq j}{\operatorname{argmax}} I_{t-1}^{ij}, \tag{4}$$

where S is the set of all conditions and subindex $t-1$ indicates that we use all measurements collected up to the point in time t. For simplicity, we define $\hat{\mathbf{r}}_{t-1}$ as the estimated posterior after measurement $t - 1$.

For each possible pair A_t, let $P(\hat{\mathbf{r}}_t | y_t = +1, A_t)$ and $P(\hat{\mathbf{r}}_t | y_t = -1, A_t)$ denote the updated posterior distributions (i.e. including comparison A_t) if o_i is selected over o_j ($y_t = +1$ for $A_t = (o_i, o_j)$) and vice versa. Since we cannot anticipate the outcome of the pairwise comparison, i.e. which condition will be selected, similarly to other active sampling methods [17], [18], [16], [14], [15], we weight the EIG with the probability of each outcome. We compute this probability using Equation 1, $P(o_i \succ o_j | \hat{\mathbf{r}}_{t-1})$; for condition o_i selected over o_j and vice versa, EIG is then defined as:

$$\begin{aligned} I_{t-1}^{ij} =& P(o_i \succ o_j | \hat{\mathbf{r}}_{t-1}) \cdot \mathrm{D_{KL}} \left(P(\hat{\mathbf{r}}_t | y_t = +1, A_t) \parallel p(\hat{\mathbf{r}}_{t-1}) \right) \\ &+ P(o_i \prec o_j | \hat{\mathbf{r}}_{t-1}) \cdot \mathrm{D_{KL}} \left(P(\hat{\mathbf{r}}_t | y_t = -1, A_t) \parallel p(\hat{\mathbf{r}}_{t-1}) \right). \end{aligned} \tag{5}$$

C. Efficiency considerations

At every iteration t, the comparisons to consider is $n(n - 1)/2$, where n is the total number of compared conditions. The complexity of the posterior evaluation is $O(n + t)$, thus the complexity of selecting the next comparison is $O(n^2(n + t))$. This may be very costly when the number of conditions is large. Here, we discuss two modifications that reduce the computational cost, and a batch mode, which also improves the accuracy.

a) Approximate (online) posterior estimation (ASAP-approx): In order to quantify the improvement in accuracy brought by the full posterior update, we follow the common approach, and consider the use of an online posterior update using assumed density filtering (ADF) [35]. That is, the posterior $\hat{\mathbf{r}}_{t-1}$ is used as the prior when computing the information gain for the t^{th} comparison, allowing our algorithm to run in an online manner [36]. Thus, for every o_i and o_j pair, we update only the scores r_i and r_j, resulting in $O(1)$ complexity per pair. No additional ADF-projection step is required since expectation propagation has already yielded a Gaussian approximation to the posterior. The time complexity of selecting the next comparison is thus decreased to $O(n^2)$. However, computational efficiency comes at the cost of accuracy in posterior estimation [36]. We refer to the algorithm using the approximate posterior update as *ASAP-approx*.

b) Selective EIG evaluations: Some comparisons are less informative than others [32], such as conditions far apart on a scale where the outcome y_t is obvious [14], [15]. Therefore we evaluate the EIG only for the most informative pairs. For that we use a simple criterion from Equation 1 to compute

the probability Q_{ij} that conditions o_i and o_j are selected for EIG evaluation. Since Equation 1 is the probability that condition o_i is better than o_j, to identify obvious outcomes we set $Q_{ij} = min(p_{ij}, p_{ji})$. Thus, the probability is large when the difference between the scores and their standard errors are small. To ensure that at least one pair including o_i is selected, we scale Q_{ij} per condition, i.e. $Q_{ij}^* = \frac{Q_{ij}}{\max_{\forall j}(Q_{ij})}$.

c) Minimum spanning tree for the batch mode: When a sampling algorithm is in the sequential mode, one pair of conditions is scheduled in every iteration of the algorithm. However, selecting a batch of comparisons in a single iteration of the algorithm is computationally more efficient and can yield superior accuracy [17]. To extend our algorithm to the batch mode, we treat pairwise comparisons as an undirected graph. Vertices are conditions, and edges are pairwise comparisons. We follow the approach from [17] where the minimum spanning tree (MST) is constructed from the graph of comparisons. The MST is a subset of the edges connecting all the vertices together, such that the total edge weight is minimal. The edges of our graph are weighed by the inverse of the EIG, i.e. for an edge E_{ij} connecting conditions A_i and A_j the weight is given by $w(E_{ij}) = \frac{1}{I^{ij}}$. $n - 1$ pairs are selected for the MST, allowing us to compute the EIG every $n-1$ iterations, greatly improving speed. Since each condition is compared at least once within our batch, detrimental imbalanced designs [37], where a subset of conditions is compared significantly more often than the rest, are eliminated.

IV. EVALUATION

To assess different sampling strategies, we run Monte Carlo simulation on synthetic and real datasets. Spearman rank ordering correlation coefficient (SROCC) and root-mean-squared Error (RMSE) between the ground truth and estimated scores are used for performance evaluation. We report our results as multiples of standard trials, where **1 standard trial** corresponds to $n(n - 1)/2$ measurements (the number of possible pairs for n conditions). For clarity, we present RMSE on a log-scale, and SROCC after a Fisher transformation ($y' = \operatorname{arctanh}(y)$). The same method, based on approximate message passing, was used to produce the scale from pairwise comparisons for each method. We verified that the scaled results are consistent with the MLE-based method from [34].

A. Algorithms compared

We implement and compare different active sampling strategies using original authors' codes where possible: AKG [29], Crowd-BT [16], HR-active [18] and Hybrid-MST [17]. Our own implementation was used for Quicksort [26], Swiss System [13], and TS-sampling [8].

B. Simulated Data

In order to control the ground truth distribution underlying the data, we first run a Monte Carlo simulation with synthetic data. In the simulation, we use $P(o_i \succ o_j | r_i, r_j)$ from Equation 1 to draw y_t for comparison at trial t between conditions o_i and o_j, which are determined by each algorithm.

We note that the strongest influence on the results is the proximity of compared conditions in the target scale. When conditions have comparable scores, they are confused more often in comparisons, whereas when conditions are far apart in the scale they are easily distinguished, resulting in different performances for sampling methods. Hence, we consider 3 scenarios for 20 conditions with scores $\mathbf{s}$ sampled uniformly from: (i) large range $[0, 20]$ (scores far apart); (ii) medium range $[0, 5]$; (iii) small range $[0, 1]$ (scores close together). Results for larger numbers of conditions are given in Section IV-D. We run the simulation 100 times for comparisons ranging from 1 to 15 standard trials.

a) Selective EIG evaluations: Figure 2a shows the proportion of saved evaluations with selective EIG computations. Since we initialize our algorithm with all scores set to 0, all possible pairs have their EIG computed at first (0 standard trials in the plot), as all conditions are close to each other. As more data are collected, conditions move away from each other on the scale and the EIG is computed for a subset of pairs only. Computational saving is greater for large-range simulations than for small-range simulations. In small-range simulations, conditions first move away from each other, as in the first few iterations their relative distances are likely to be overestimated, decreasing the overall number of computations; however, with more measurements the conditions move closer, and the proportion of saved evaluations decreases. Figure 2b shows the probability of the EIG being evaluated after 10 standard trials for 20 conditions sampled from the medium range. For visualization purposes, conditions were ordered ascending in the quality scale. Pairs of conditions along the diagonal, i.e. close in the scale, have a higher chance of their EIG being computed. Figure 2c shows performance of ASAP with and without selective EIG evaluations. Thus, selectively evaluating EIG greatly reduces the number of computations while maintaining the same accuracy measured in RMSE and SROCC. In the following sections, we only present the results with selective EIG computations.

b) Minimum spanning tree for the batch mode: Figure 3 shows the results of ASAP with and without batch mode for medium-range simulations. Without MST batch mode, the method is likely to result in an imbalanced sampling pattern, where certain conditions are compared significantly more often than others. This has a detrimental effect on performance, deteriorating the results with growing number of comparisons [37]. Below, we only present results with MST batch mode.

c) Simulation results: Figure 4 shows the results of the simulation for the implemented strategies. At all tested ranges, EIG-based methods have lower RMSE, and therefore higher accuracy, than the sorting methods (Quicksort and the Swiss System). While TS-sampling and Crowd-BT have good accuracy for the large range, these are among the worst methods for the small range. ASAP-approx exerts performance similar to the methods with online posterior update, however offers a modest but consistent improvement in accuracy over Hybrid-MST and HR-active. Of all tested methods, ASAP, employing full posterior update, is the most accurate by a

(a) Saved evaluations   (b) Probability of EIG evaluation
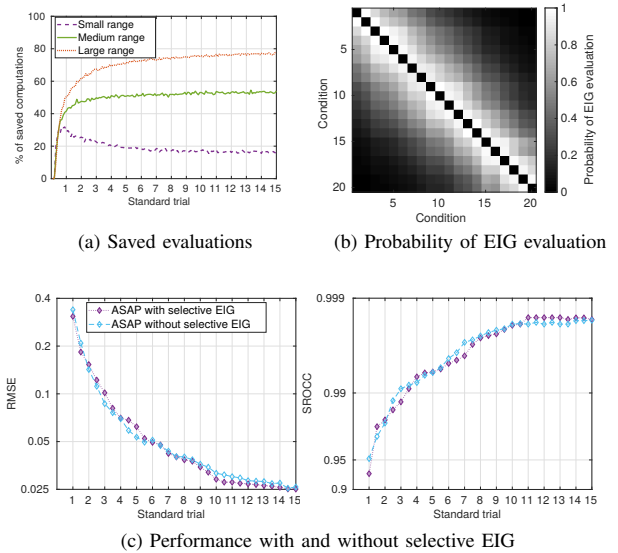

(c) Performance with and without selective EIG

Fig. 2: (a) Percentage of saved evaluations with selective EIG evaluations; (b) probability of EIG evaluation after 10 standard trials for medium range; and (c) RMSE and SROCC with and without selective EIG;
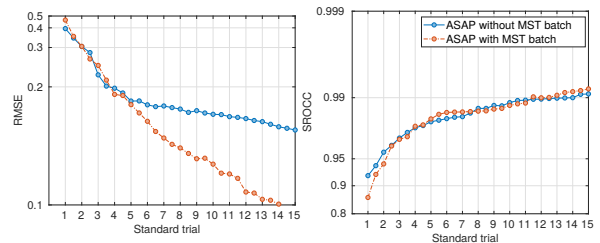

Fig. 3: Simulation with 20 conditions sampled from the medium range with and without MST. We observe similar pattern for conditions sampled from small and large ranges.

substantial margin and across all ranges.

For SROCC, EIG-based methods do not show a clear advantage over sorting methods; however, it should be noted that EIG-based methods are designed to optimize for RMSE rather than ranking. Even so, ASAP still performs the best for small and medium range simulations, and one of the best for large range, reaching SROCC of 0.99 within five standard trials. It should be noted, however, that the problem of ordering conditions from the large range is trivial and the best methods compete at 0.99+ SROCC levels (almost perfect ordering). Because of the poor performance of the sorting-based methods, we do not consider them in the following experiments.

C. Real Data

We validate the performance of sampling strategies on two real-world datasets: i) Image Quality Assessment (IQA) LIVE dataset [38], with pairwise comparisons collected by [4];
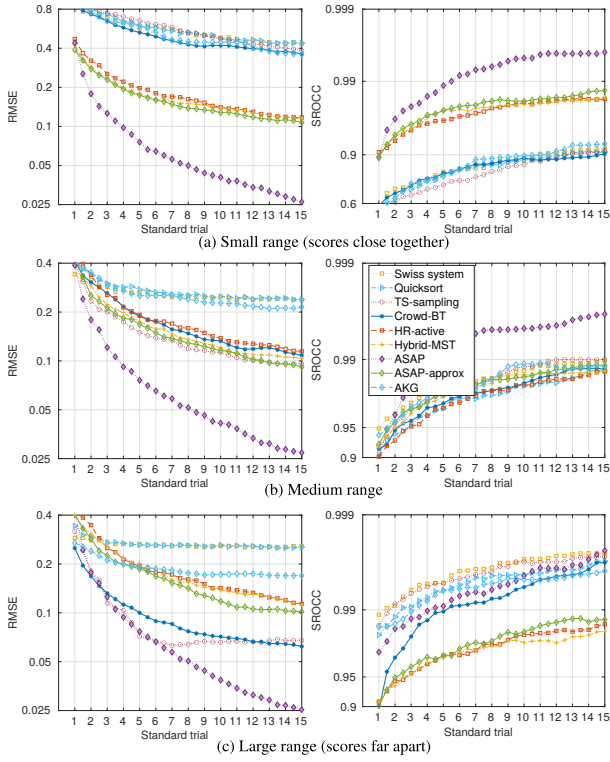
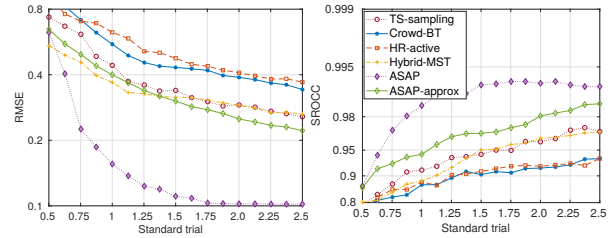Fig. 4: Simulation results with 20 conditions for the compared sampling strategies.

Fig. 5: Compared sampling strategies on LIVE dataset.

and ii) Video Quality Assessment (VQA) dataset [39]. Each dataset contains complete and balanced matrices of pairwise comparisons, with each condition compared to every other condition the same number of times. The empirical probability of one condition being better than another is obtained from the measured data and used throughout the simulation. We compute RMSE and SROCC between scores produced by each method, and scores obtained by scaling the original matrices of all comparisons.

a) IQA dataset: To allow multiple runs of the Monte Carlo simulation, we randomly select 40 conditions from the 100 available. In the original matrix, each condition is compared 5 times with each other (5 standard trials), yielding 24750 comparisons.

Figure 5 shows the results. The performance trends are

(a) Reference 1

(b) Reference 2
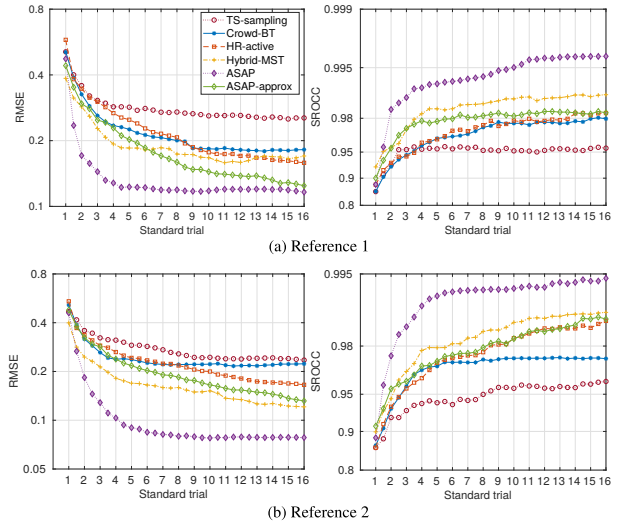
Fig. 6: Compared sampling strategies on VQA dataset.

consistent with the results for the simulated data for the medium range. ASAP has the best performance both in terms of SROCC and RMSE. It is followed by ASAP-approx, Hybrid-MST, and TS-sampling, each having roughly the same performance in terms of both RMSE and SROCC. Crowd-BT and HR-active have the worst performance in terms of both RMSE and SROCC.

b) VQA dataset: The dataset contains 10 reference videos with 16 distortions. Each 16×16 matrix contains 3840 pairwise comparisons, i.e. each pair was compared 32 times.

Figure 6 shows the results of running simulations on the first two reference videos. The performance trends are again, in general, consistent with the results for the simulated data sampled from the medium range, except that TS-sampling performs substantially worse, and Hybrid-MST outperforms ASAP-approx for small numbers of trials. ASAP consistently outperforms other methods. The results for the remaining eight reference videos are given in the supplementary.

D. Large Scale Experiments

It is often considered that 15 standard trials is the minimum requirement for FPC to generate reliable results [40], [41], however, this is rarely feasible in practice. Real-world large-scale datasets barely reach 1 standard trial. To make experiments with large number of conditions feasible, individual reference scenes or videos are often measured and scaled independently, missing important cross-content comparisons. However, the lack of cross-content comparisons yields less accurate scores [24]. Active sampling techniques, such as ASAP, should accurately measure a large number of conditions, while saving substantial amount of experimental effort. To test such a scenario, we simulate the comparison of 200 conditions with scores distributed in the medium range. The results, shown in Figure 7, demonstrate that even with a small
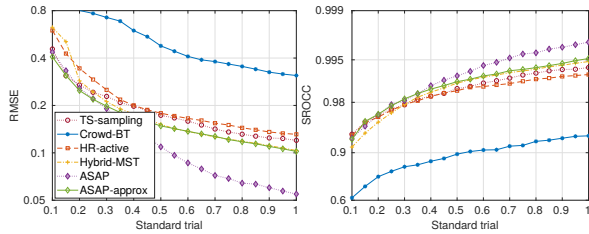
Fig. 7: Large scale experiment simulation with 200 conditions sampled from the medium range.

number of standard trials ASAP outperforms existing methods; it is followed by ASAP-approx and Hybrid-MST.

E. Running Time and Experimental Effort

A practical active sampling method must generate new samples in an acceptable amount time. Hence, in Figure 8 we plot the time taken by each method as the number of conditions grows. The reported times are for generating a single pair of conditions, assuming that 5 standard trials have been collected so far. CPU times were measured for MATLAB R2019a code running on a 2.6GHz Intel Core i5 CPU and 8GB 1600MHz DDR3 RAM. GPU time was measured for Pytorch 1.4 with CUDA 9.2, running on GeForce GTX1080. We omit sorting methods as they do not offer sufficient accuracy. Although ASAP is the slowest method when running on a CPU, it can be effectively parallelized on a GPU and deliver the results in a shorter time than other methods running on a CPU.

In Figure 9 we show the experimental effort required to reach an acceptable level of accuracy for 20 and 200 conditions, where we define experimental effort as the time required to reach an RMSE of 0.15. We assume that each comparison takes 5 seconds, which is typical for image quality assessment experiments [42], [13]. ASAP offers the biggest saving in experimental effort for both small and large scale experiments. In an experiment with 200 conditions ASAP achieves an accuracy of 0.15 RMSE in 0.355 standard trials. The total experimental time is thus 9.8h (7065 comparisons), which is significantly better than the 14.6h (10550 comparisons) for Hybrid-MST. Similarly, for 20 conditions the entire experiment would take 40 min for ASAP and 120 min for Hybrid-MST to reach the same accuracy of score estimates. For experiments with longer comparison times (e.g. video comparison) or high comparison cost (e.g. medical images) ASAP's advantage is even greater.

V. CONCLUSIONS

In this paper, we showed the importance of choosing the right sampling method when collecting pairwise comparison data, and proposed a fully Bayesian active sampling strategy for pairwise comparisons – ASAP.

Commonly used sorting methods perform poorly compared to the state-of-the-art methods based on the EIG, and even EIG-based methods are sub-optimal, as they rely on a partial update of the posterior distribution. ASAP computes the full posterior distribution, which is crucial to achieving accurate
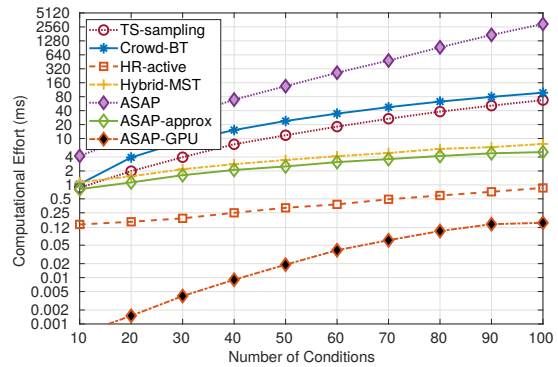

Fig. 8: Average time to select the next comparison for a varied number of conditions and 5 standard trials.
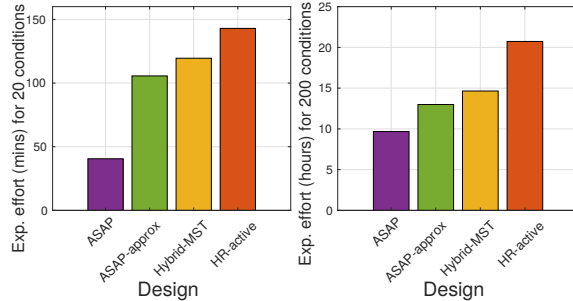

Fig. 9: Experimental effort (amount of time, assuming 5 second decision time, required to reach 0.15 RMSE) for experiments with 20 and 200 conditions.

EIG estimates, and thus the accuracy of active sampling. Fast computation of the posterior, important for real-time applications, was made possible by using fast and accurate factor graph approach, which is new to the active sampling community. In addition, ASAP only computes the EIG for the most informative pairs, reducing the computational cost of ASAP by up to 80%, and selects batches using a minimum spanning tree method, allowing to avoid imbalanced designs.

We recommend ASAP, as it offered the highest accuracy of inferred scores compared to existing methods in experiments with real and synthetic data. The computational cost of our technique is higher than for other methods in the CPU implementation, but is still in the range that makes the technique practical, with a substantial saving of experimental effort. For large-scale experiments, in GPU implementation ASAP offers both accuracy and speed.

REFERENCES

[1] K. Tsukida and M. R. Gupta, "How to Analyze Paired Comparison Data," Department of Electrical Engineering University of Washington, Tech. Rep. UWEETR-2011-0004, 2011. 1

[2] N. Shah, S. Balakrishnan, J. Bradley, A. Parekh, K. Ramchandran, and M. Wainwright, "Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence," in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, vol. 38. PMLR, 2015, pp. 856–865. 1

[3] R. K. Mantiuk, A. Tomaszewska, and R. Mantiuk, "Comparison of four subjective methods for image quality assessment," *Computer Graphics Forum*, vol. 31, no. 8, pp. 2478–2491, 2012. 1

[4] P. Ye and D. Doermann, "Active sampling for subjective image quality assessment," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, June 2014, pp. 4249–4256. 1, 2, 5

[5] L. L. Thurstone, "A law of comparative judgment." *Psychological Review*, vol. 34, no. 4, pp. 273–286, 1927. 1, 3

[6] R. A. Bradley and M. E. Terry, "Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons," *Biometrika*, vol. 39, no. 3/4, p. 324, dec 1952. 1

[7] Y. Kim, W. Kim, and K. Shim, "Latent ranking analysis using pairwise comparisons," in *2014 IEEE International Conference on Data Mining*, Dec 2014, pp. 869–874. 1

[8] R. Herbrich, T. Minka, and T. Graepel, "Trueskill™: A bayesian skill rating system," in *Advances in Neural Information Processing Systems 19*. MIT Press, 2007, pp. 569–576. 1, 3, 4

[9] M. E. Glickman, "Parameter estimation in large dynamic paired comparison experiments," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 48, no. 3, pp. 377–394, 1999. 1

[10] R. A. Cribbie and H. J. Keselman, "Pairwise multiple comparisons: A model comparison approach versus stepwise procedures," *British Journal of Mathematical and Statistical Psychology*, vol. 56, no. 1, pp. 167–182, 2003. 1

[11] E. Prashnani, H. Cai, Y. Mostofi, and P. Sen, "Pieapp: Perceptual image-error assessment through pairwise preference," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1

[12] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp. 586–595. 1

[13] N. Ponomarenko, L. Jin, O. Ieremeiev, V. Lukin, K. Egiazarian, J. Astola, and Benoit, "Image database tid2013: Peculiarities, results and perspectives," *Signal Processing: Image Communication*, vol. 30, pp. 57 – 77, 2015. 1, 2, 4, 7

[14] T. Pfeiffer, X. Gao, Y. Chen, A. Mao, and D. Rand, "Adaptive polling for information aggregation," *AAAI Conference on Artificial Intelligence*, 2012. 1, 2, 4

[15] M. E. Glickman and S. T. Jensen, "Adaptive paired comparison design," *Journal of Statistical Planning and Inference*, vol. 127, no. 1, pp. 279 – 293, 2005. 1, 2, 4

[16] X. Chen, P. N. Bennett, K. Collins-Thompson, and E. Horvitz, "Pairwise ranking aggregation in a crowdsourced setting," *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pp. 193–202, 2013. 1, 2, 4

[17] J. Li, R. Mantiuk, J. Wang, S. Ling, and P. Le Callet, "Hybrid-mst: A hybrid active sampling strategy for pairwise preference aggregation," *NIPS, 31st Conference on Neural Information Processing Systems*, 2018. 1, 2, 4

[18] Q. Xu, J. Xiong, X. Chen, Q. Huang, and Y. Yao, "Hodgerank with information maximization for crowdsourced pairwise ranking aggregation," *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pp. 4326–4334, 2018. 1, 2, 4

[19] R. Heckel, M. Simchowitz, K. Ramchandran, and M. J. Wainwright, "Approximate ranking from pairwise comparisons," *arXiv preprint arXiv:1801.01253*, 2018. 2

[20] R. Heckel, N. B. Shah, K. Ramchandran, and M. J. Wainwright, "Active ranking from pairwise comparisons and when parametric assumptions do not help," *Ann. Statist.*, vol. 47, no. 6, pp. 3099–3126, 12 2019. 2

[21] K. G. Jamieson and R. D. Nowak, "Active ranking using pairwise comparisons," in *Proceedings of the 24th International Conference on Neural Information Processing Systems*, ser. NIPS11. Red Hook, NY, USA: Curran Associates Inc., 2011, p. 22402248. 2

[22] Y. Yue and T. Joachims, "Beat the mean bandit," in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ser. ICML11. Madison, WI, USA: Omnipress, 2011, p. 241248. 2

[23] B. Szörényi, R. Busa-Fekete, A. Paul, and E. Hüllermeier, "Online rank elicitation for plackett-luce: A dueling bandits approach," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS15. Cambridge, MA, USA: MIT Press, 2015, p. 604612. 2

[24] 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)*, Burlingame, United States, Jan. 2018. 2, 6

[25] D. A. Silverstein and J. E. Farrell, "Efficient method for paired comparison," *Journal of Electronic Imaging*, vol. 10, pp. 10 – 10 – 5, 2001. 2

[26] L. Maystre and M. Grossglauser, "Just sort it! A simple and effective approach to active preference learning," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70. PMLR, 06–11 Aug 2017, pp. 2344–2353. 2, 4

[27] C. A. R. Hoare, "Quicksort," *The Computer Journal*, vol. 5, no. 1, pp. 10–16, 01 1962. 2

[28] ITU-R, "Subjective assessment methods for 3D video quality," ITU-R Recommendation P.915, Mar 2016. 2

[29] X. Chen, K. Jiao, and Q. Lin, "Bayesian decision process for cost-efficient dynamic ranking via crowdsourcing," *Journal of Machine Learning Research*, vol. 17, no. 216, pp. 1–40, 2016. 2, 4

[30] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1-2, pp. 81–93, 06 1938. 2

[31] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 03 1951. 2, 3

[32] B. Settles, "Active learning literature survey," University of Wisconsin-Madison, Tech. Rep., 2010. 2, 4

[33] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Third Edition*, 3rd ed. The MIT Press, 2009. 2

[34] M. Perez-Ortiz and R. K. Mantiuk, "A practical guide and software for analysing pairwise comparison experiments," *CoRR*, 2017. 3, 4

[35] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*, 1st ed. MIT Press, 2012. 4

[36] T. Minka, R. Cleven, and Y. Zaykov, "Trueskill 2: An improved bayesian skill rating system," Microsoft, Tech. Rep. MSR-TR-2018-8, March 2018. 4

[37] N. B. Shah and M. J. Wainwright, "Simple, robust and optimal ranking from pairwise comparisons," *Journal of Machine Learning Research*, vol. 18, no. 199, pp. 1–38, 2018. 4, 5

[38] H. Sheikh, M. Sabir, and A. Bovik, "A Statistical Evaluation of Recent Full Reference Image Quality Assessment Algorithms," *IEEE Transactions on Image Processing*, vol. 15, no. 11, pp. 3440–3451, 2006. 5

[39] Q. Xu, T. Jiang, Y. Yao, Q. Huang, B. Yan, and W. Lin, "Random partial paired comparison for subjective video quality assessment via hodgerank," in *Proceedings of the 19th ACM International Conference on Multimedia*. ACM, 2011, pp. 393–402. 5

[40] ITU-T P.910, "Subjective video quality assessment methods for multi-media application," *International Telecommunication Union*, 2008. 6

[41] ITU-R BT.500-13, "Methodology for the subjective assessment of the quality of television pictures," *International Telecommunication Union*, 2012. 6

[42] M. Prez-Ortiz, A. Mikhailiuk, E. Zerman, V. Hulusic, G. Valenzise, and R. K. Mantiuk, "From pairwise comparisons and rating to a unified quality scale," *IEEE Transactions on Image Processing*, pp. 1–1, 2019. 7

Supplementary material for: "Active Sampling for Pairwise Comparisons via Approximate Message Passing and Information Gain Maximization"

Aliaksei Mikhailiuk*, Clifford Wilmot*, Maria Perez-Ortiz‡*, Dingcheng Yue*, Rafał K. Mantiuk*

*Department of Computer Science University of Cambridge Cambridge, United Kingdom
‡Department of Computer Science, University College London, London, United Kingdom
Email: am2442@cam.ac.uk, cliffw29@outlook.com, maria.perez@ucl.ac.uk, {dy276, rafal.mantiuk}@cam.ac.uk

I. INTRODUCTION

We include pseudo-code for the ASAP algorithm, sampling patterns of the ASAP, the distribution of the RMSE and SROCC for the conditions sampled from the large range (bottom row of Figure 4 in the main paper) and the results for the 8 reference videos of the Video Quality Assessment (VQA) dataset.

A. Pseudo-code

We provide a detailed description of the ASAP method in Algorithm 1. The notation in the algorithm corresponds to that used in the main paper. The algorithm requires as input a list of comparisons performed so far, $\mathbf{y}$ and a matrix with the probabilities of the EIG being computed —required for the selective evaluations —$\mathbf{Q}$. The algorithm outputs a batch of pairs to be compared $\mathbf{C}$ and an updated probability of being selected for the EIG evaluations, $\hat{\mathbf{Q}}$.
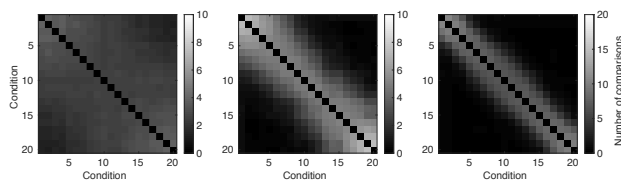
B. Sampling patterns

Fig. 1: Heatmaps for 3 standard trials of comparisons of 20 ordered conditions from small, medium and large ranges

In order to better understand which conditions are favored by ASAP, we produce heatmaps of pairings for conditions sampled from the small, medium and large ranges. We use 3 standard trials. The heatmaps are given in Figure 1. For better visualization, conditions are ordered ascending in their ground truth scores in the consecutive rows and columns. For conditions sampled from the small range, all pairs of conditions are compared approximately the same number of times. However, the number of comparisons gradually decreases for conditions further away in the scale, i.e. further away from the diagonal on the heatmap. For conditions sampled from the medium and large ranges, most comparisons are selected for
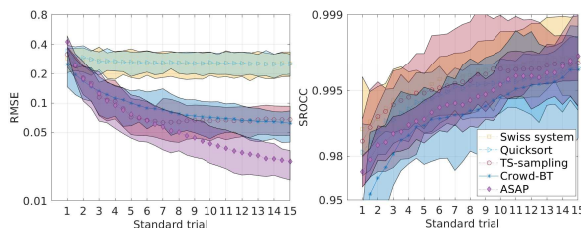

Fig. 2: 75% of the RMSE and SROCC distributions for 20 conditions sampled from the large range and five best performing methods in terms of SROCC.

conditions close in the quality scale, i.e. along the diagonal on the heatmap. This is expected, as pairs of conditions that are far away in the scale are less likely to be confused by observers and are therefore less informative.

### C. Confidence intervals for large range

Figure 2 presents 75% of the RMSE and SROCC distributions for the top five methods in terms of SROCC for conditions sampled from the large range (bottom row of Figure 6 in the main paper). Results for SROCC are noisier than for RMSE, making it hard to identify the best performing method. In terms of RMSE for a number of standard trials less than two ASAP shows similar to others performance, however with the number of standard trials growing, significantly outperforms the compared methods.

D. Results: VQA dataset

This dataset contains 10 reference videos with 16 distortions applied to them. Each 16×16 matrix contains 3840 pairwise comparisons - each pair was compared 32 times. Figure 3 shows the results for reference videos 3 to 10. Results for the first two reference videos are presented in the main paper. Consistent with other tests in the main paper, ASAP shows superior results to other methods. ASAP-approx. has average, similar to other EIG based methods, results. Hybrid-MST tends to perform better for small numbers of standard trials.

Algorithm 1: ASAP

Input: $\mathbf{y}, \mathbf{Q}$
Output: $\mathbf{C}, \hat{\mathbf{Q}}$
 # Calculate the posterior for the given state of the comparison matrix
$\boldsymbol{\mu}, \boldsymbol{\Sigma}$ = approxPosterior($\mathbf{y}$)
 # Iterate over the rows of the expected information gain matrix $\mathbf{I}$
for $i \leftarrow 1$ **to** n **do**

    # Iterate over the columns of the expected information gain matrix $\mathbf{I}$
    for $j \leftarrow 1$ **to** $(i-1)$ **do**

        # Probability of selecting o_i over o_j
        $p_{ij} = P(o_i \succ o_j | r_i, r_j) = \Phi\left(\frac{\mu_i - \mu_j}{\sqrt{2}\sigma_{ij}}\right)$
        # Selective EIG evaluations with roulette
        if $Q_{ij} > U[0,1]$ **then**

            # Posterior given all comparisons and assuming o_i is selected over o_j
            μ_{ij}, Σ_{ij} = approxPosterior($\mathbf{y}; o_i \succ o_j$)
            # Posterior given all comparisons and assuming o_j is selected over o_i
            μ_{ji}, Σ_{ji} = approxPosterior($\mathbf{y}; o_j \succ o_i$)
            # KL divergence between current distribution and distribution assuming the two possible outcomes
            KL_{ij} = KLDivergence($\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \mathcal{N}(\mu_{ij}, \Sigma_{ij})$)
            KL_{ji} = KLDivergence($\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \mathcal{N}(\mu_{ji}, \Sigma_{ji})$)
            # Weighted information gain
            $I_{ij} = p_{ij} \times KL_{ij} + (1 - p_{ij}) \times KL_{ji}$

        end

        # Update the probability of being selected for the comparison
        $\hat{Q}_{ij} = min(p_{ij}, 1 - p_{ij})$

    end

    # Scale q per condition
    $\hat{Q}_{ij} = \frac{\hat{Q}_{ij}}{\max_{\forall j}(\hat{Q}_{ij})}, \ \forall j$

end
 # Make the EIG matrix symmetric and find reciprocal of each entry
$I = \mathbf{1}/(I + I^T)$
 # Create the minimum spanning tree from the matrix $\mathbf{I}$
G = minspantree($\mathbf{I}$)
 # Nodes connected by an edge are pairs of conditions to compare
$\mathbf{C}$ = getConnectedNodes(G)
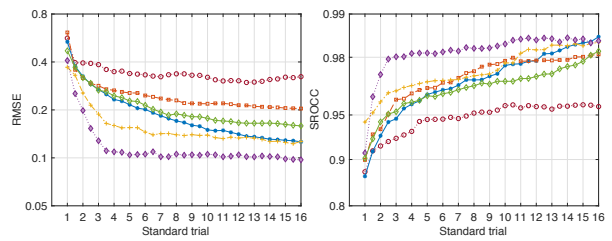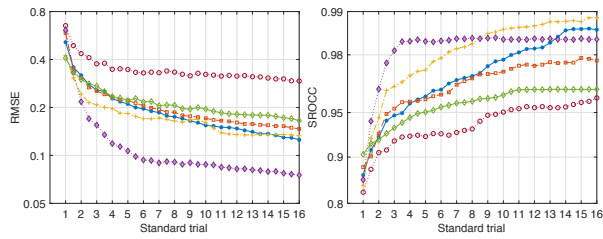*# Note if batch mode is not used pairs to compare are selected by* $\mathbf{C} = argmax(I_{ij})$
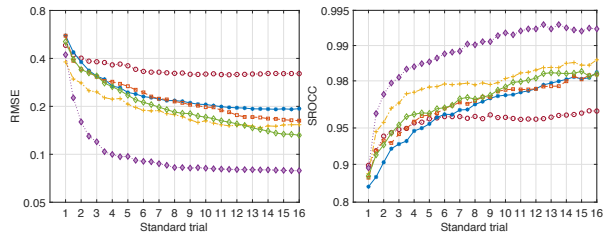
(a) Reference 3
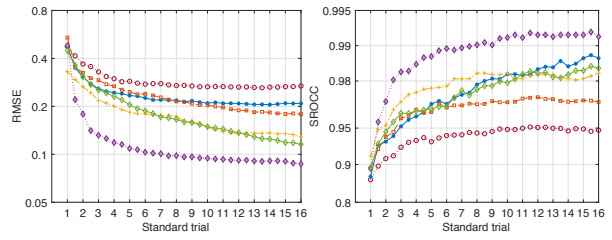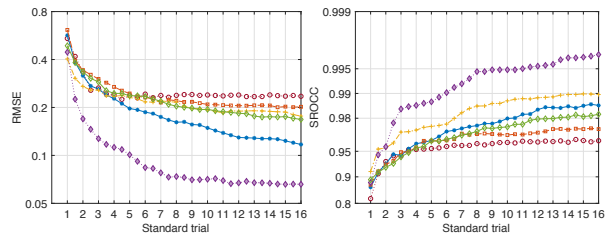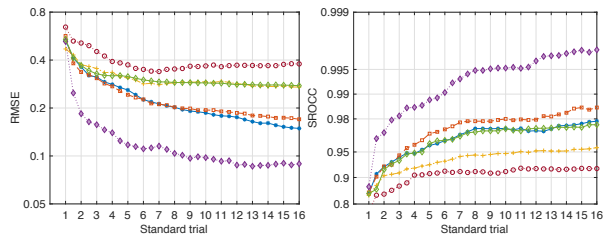
(b) Reference 4
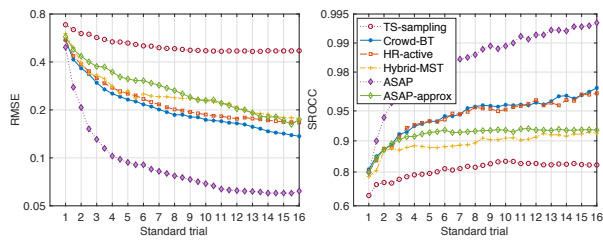
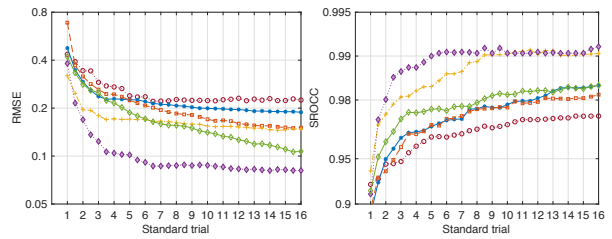(c) Reference 5

(d) Reference 6

(e) Reference 7

(f) Reference 8

(g) Reference 9

(h) Reference 10

Fig. 3: Compared sampling strategies on VQA dataset.