

Ranking a Set of Objects using Heterogeneous Workers: QUITE an Easy Problem

Original

Ranking a Set of Objects using Heterogeneous Workers: QUITE an Easy Problem / Nordio, A., Tarable, A., Leonardi, E..
- In: IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING. - ISSN 2327-4697. - ELETTRONICO. -
12:6(2025), pp. 4741-4756. [10.1109/TNSE.2025.3575231]

Availability:

This version is available at: 11583/3004871 since: 2025-11-06T07:45:36Z

Publisher:

IEEE

Published

DOI:10.1109/TNSE.2025.3575231

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Ranking a Set of Objects using Heterogeneous Workers: QUITE an Easy Problem

Alessandro Nordio, *Member, IEEE*, Alberto Tarable, *Member, IEEE*, Emilio Leonardi

Abstract—We focus on the problem of ranking N objects starting from a set of noisy pairwise comparisons provided by a crowd of unequal workers, each worker being characterized by a specific degree of reliability, which reflects her ability to rank pairs of objects. More specifically, we assume that objects are endowed with intrinsic qualities and that the probability with which an object is preferred to another depends both on the difference between the qualities of the two competitors and on the reliability of the worker. We propose QUITE, a non-adaptive ranking algorithm that jointly estimates workers’ reliabilities and qualities of objects. Performance of QUITE is compared in different scenarios against previously proposed algorithms. Finally, we show how QUITE can be naturally made adaptive.

Index Terms—Ranking algorithms, heterogeneous workers, noisy evaluation, applied graph theory, least-square estimation

I. INTRODUCTION AND RELATED WORK

THIS paper focuses on the problem of establishing a reliable ranking among several objects, starting from a set of noisy human evaluations. Such problem emerges in several computer-science contexts, e.g. when web pages are ranked by search engines, when hotels and restaurants are ranked by applications like Tripadvisor, or when products are ranked by on-line sellers [1], [2].

Often, a ranking algorithm receives as input a set of noisy preferences between pairs of objects and infers an estimated order relation among them. Comparisons are sometimes made by human workers, whose behavior cannot be deterministically predicted. Indeed, outcomes of comparisons depend on how objects are “perceived” by human workers, rather than on their intrinsic quality.

This specific problem has attracted a significant bulk of attention in the last few years [3]–[10], and a set of stochastic laws has been proposed to represent the behavior of human workers, such as the very popular Bradley-Terry-Luce [11]–[13] and Thurstone [14], [15], [16] models. Most of them are based on the assumption that an intrinsic quality can be associated to every object and that the probability that object i is preferred to object j consistently depends on the associated qualities q_i and q_j . Few works, on the contrary, explore the problem in the more general and challenging setting in which no intrinsic quality can be associated to objects [8], [17]. The vast majority of previous works [3]–[6], [8]–[10], [18]–[20] assumes workers to be homogeneous, i.e., to behave exactly according to the same law. For such simplified scenario,

several ranking algorithms have been proposed in [3]–[6] and its asymptotic performance has thoroughly been analysed, typically within the (ε, δ) -PAC (probably approximately correct) framework [18]–[20]. A more detailed description of previous work on ranking with homogeneous workers can be found in Appendix A in the Supplemental Material.

In particular, by strengthening and generalizing previous results ([3]–[5], [18]–[20]), in [6], we have recently proposed a class of non-adaptive ranking algorithms that rely on a least-squares (LS) optimization criterion for the estimation of qualities. In the scenario of homogeneous workers with known reliability, such LS algorithms exploit the structure of a graph \mathcal{G} of cardinality N , whose edges are in one-to-one correspondence with the evaluated object pairs. The LS algorithms in [6] are shown to be asymptotically optimal (i.e., they require $O(\frac{N}{\varepsilon^2} \log \frac{N}{\delta})$ comparisons to meet (ε, δ) -PAC constraints), as long as the graph edges are properly selected. Also, they operate by receiving in input the set of estimated distances between object pairs and by returning the quality estimates as well as the estimated ranking.

Only recently, the assumption that all workers obey the same law has been loosened in [7], [21], [22]. The work in [7] presents a non-adaptive ranking algorithm, which receives as input a set of pairwise preferences expressed by heterogeneous users, each one obeying either a generalized BTL or a Thurstone model and characterized by a different reliability, and outputs a suitable ranking of evaluated objects. To the best of our knowledge, [7] is the only paper that addresses this specific problem, proposing an algorithm which is based on an approximate Maximum-Likelihood (ML) estimation of worker reliabilities and object qualities. The work in [7] also provides a theoretical analysis of the algorithm convergence, from which, however, asymptotic properties can hardly be obtained. Details of [6] and [7] will be given in Section IV. In [21], [22], instead, a fully adaptive approach is chosen, according to which object-pair comparisons are performed one-by-one, by selecting each time the most suitable worker/object-pair on the base of the previous history. We believe that in many contexts fully adaptive algorithms are hardly implementable in practice due to the large delay needed to obtain results, which may be induced by strictly sequential human evaluations¹.

The very recent work [17] focuses on a different scenario where no intrinsic qualities can be associated to objects, and the probability with which object i is preferred to j by worker w only depends on w and on the true ranking between i and j . For such scenario, [17] proposes an algorithm for ranking

A. Nordio and A. Tarable are with CNR-IEIIT (Institute of Electronics, Telecommunications and Information Engineering of the National Research Council of Italy), Italy, email: firstname.lastname@ieiit.cnr.it.

E. Leonardi is with DET, Politecnico di Torino, Torino, Italy and is also associate researcher with CNR-IEIIT.

¹In the literature, several similar problems are considered, as for example, the top- k ranking treated in [23], which is however different from our goal.

aggregation, whose performance is analytically evaluated.

These are the major contributions of the present work.

- We propose a non-adaptive algorithm called QUITE, specifically tailored to deal with heterogeneous workers. QUITE can be successfully employed under a rather general class of worker behavior models, e.g., the generalized Thurstone model considered in [7] as well as a generalized version of the Bradley-Terry-Luce model. Our approach is based on an iterative procedure. Given worker reliability estimates, our algorithm first exploits an MAP approach to compute distance estimates of evaluated object pairs, and then resorts to the graph-based LS method proposed in [6] for the scenario with homogeneous workers, to estimate the object qualities. It then employs an MAP approach to refine the worker reliability estimates, given the estimates of the object qualities. The procedure iterates until convergence is reached. As better explained in Section IV, QUITE significantly differs from the algorithm proposed in [7], which implements a Heterogeneous-worker Alternate-Gradient Descent (HAGD) optimization, to approximate a joint ML estimation of objects qualities and worker reliabilities. Moreover, contrarily to the algorithm in [7], at least in its simpler form, QUITE is amenable to a theoretical asymptotical analysis, which shows that, under mild conditions, it is asymptotically optimal (i.e., it requires $O(N/\varepsilon^2 \log(N/\delta))$ comparisons to comply (ε, δ) -PAC requirements).
- We derive a Bayesian Cramér-Rao lower bound (BCRB) for the mean-square error achievable by any estimation technique of quality differences and/or worker reliabilities.
- We test the performance of QUITE against the BCRB, and compare it with the algorithm in [7].
- We extend QUITE, by proposing a partially adaptative version of it, which works in a multistage fashion where the assignment of object pairs to workers is made in a few stages, by exploiting previous partial estimates of workers' reliabilities and qualities of objects. In such a case, we propose a simple recipe for assigning workers to object pairs, on the base of estimations performed at previous stages.
- Finally, we provide an extensive set of simulation and experimental results, proving the effectiveness of our approach and its superiority with respect to the alternating-gradient scheme proposed in [7].

The rest of the manuscript is organized as follows: Section II describes the system model and discusses possible approaches for obtaining a ranking. Section III derives the expression of the BCRB on the variance of the quality-reliability estimates. Section IV provides a review of ranking algorithms available in the literature and introduces our proposed joint quality-reliability estimation algorithm, i.e., QUITE. Section V deals with a multistage version of QUITE and gives also a heuristic rule of assignment of objects pairs to workers, on the basis of previous output of the algorithm. In Section VI, we provide numerical results showing the performance of QUITE in sev-

Variable	Description
N	number of objects to rank
q_i	intrinsic quality object i
$r(i)$	rank of object i
\mathcal{E}	set of object pairs to be evaluated
\mathcal{E}_k	set of object pairs evaluated by worker k
(i_e, j_e)	e -th object pair in \mathcal{E}
$d_e = q_{i_e} - q_{j_e}$	quality difference in object pair e
$q_{i_e, k}^{\text{perceived}}$	quality of object i_e perceived by worker k
$w_{e, k} \in \{0, 1\}$	evaluation outcome of pair e provided by worker k
ρ_k	reliability of worker k

TABLE I
SUMMARY OF MAIN NOTATION

eral scenarios, while in Section VII we discuss the complexity of QUITE. Finally, in Section VIII, we draw the conclusions of our work.

A. Mathematical notation

Boldface uppercase and lowercase letters denote matrices and vectors, respectively. The (i, j) -th entry of matrix \mathbf{A} is denoted by $[\mathbf{A}]_{i, j}$ and its transpose is denoted by \mathbf{A}^\top . Calligraphic letters denote sets or graphs. The symbols $\mathbb{E}_{\mathbf{x}}[\cdot]$, $\mathbb{V}_{\mathbf{x}}[\cdot]$ and $\nabla_{\mathbf{x}}$ are, respectively, the expectation operator, the variance operator and the gradient w.r.t. vector \mathbf{x} . The probability of an event x is denoted by $\mathbb{P}(x)$ while $f_x(x)$ indicates the probability density function of the random variable x . A uniform probability distribution with support $[a, b]$ is denoted by $\mathcal{U}[a, b]$ whereas the Gaussian distribution with mean μ and variance σ is denoted by $\mathcal{N}(\mu, \sigma)$.

II. SYSTEM DESCRIPTION

We consider the problem of ranking a set of N objects having unknown intrinsic qualities q_1, \dots, q_N , where $q_i \in \mathbb{R}$. A ranking is a relationship between the elements in the set, such that an object ranked higher is considered “better”, “superior”, or “preferred” to an object ranked lower. For example, the qualities q_i induce a true ranking, r , in which $r(i) \succ r(j)$ (i.e., object i is better than object j) if and only if $q_i > q_j$. In other words, the true ranking corresponds to a permutation, π , of the integers $\{1, \dots, N\}$ defined by sorting object qualities, such that $q_{\pi_1} > q_{\pi_2} > \dots > q_{\pi_N}$ ².

A ranking algorithm provides an estimated ranking, \hat{r} , by resorting to some information obtained by comparing the objects. Such estimated ranking corresponds to a permutation $\hat{\pi}$ of the integers $\{1, \dots, N\}$, with the meaning that object $\hat{\pi}_1$ is ranked the best, and object $\hat{\pi}_N$ the worst. We say that \hat{r} is an ϵ -quality ranking if $\hat{r}(i) \succ \hat{r}(j)$ whenever $q_i > q_j + \epsilon$. Moreover a ranking algorithm is (ε, δ) -PAC [18]–[20] if it returns an ϵ -quality ranking with a probability larger than $1 - \delta$.

A ranking algorithm processes information obtained, e.g., through a set of observations or comparisons made by a pool of workers. In particular, we assume that a set \mathcal{E} of distinct object pairs, having cardinality $E = |\mathcal{E}|$, is given to a pool of K workers for evaluation; for each assigned pair, a worker gives a binary answer, indicating the object in the pair he/she

²We suppose ties happen with probability zero.

ranks higher. Since workers are not fully reliable, the same pair is given to several of them. Workers' answers can be considered as independent random variables and modeled as described in the following.

Let $(i_e, j_e) \in \mathcal{E}$ be the e -th object pair, $e = 1, \dots, E$ and denote by $d_e = q_{i_e} - q_{j_e}$ the quality difference between objects in pair e . Also, let $\mathcal{E}_k \subseteq \{1, \dots, E\}$ be the subset of pairs assigned to worker k , $k = 1, \dots, K$ and $\mathcal{K}_e \subseteq \{1, \dots, K\}$ be the subset of workers assigned to pair e , $e = 1, \dots, E$. The set of binary outcomes of the evaluation is denoted by $\mathcal{W} = \{w_{e,k}\}$, where $w_{e,k} \in \{0, 1\}$ represents the outcome of the evaluation provided by worker k on object pair $e \in \mathcal{E}_k$. In particular $w_{e,k} = 0$ if worker k prefers object i_e when evaluating pair e and $w_{e,k} = 1$ otherwise. The preferred object is chosen in accordance to quality *as perceived* by worker k and defined as $q_{i_e, k}^{\text{perceived}} = q_{i_e} + \eta_{k,e,i_e}$ where η_{k,e,i_e} is a random evaluation noise, modeled as independent for each object. In practice the worker outputs $w_{e,k} = 0$ if $q_{i_e, k}^{\text{perceived}} > q_{j_e, k}^{\text{perceived}}$. In order to characterize workers' evaluations, we assign to worker k an unknown reliability parameter $\rho_k \in \mathbb{R}^+$, $k = 1, \dots, K$. Workers with large ρ_k provide highly reliable answers, i.e., they are very sensitive to small quality differences between objects. Instead, answers provided by workers characterized by small ρ_k tend to be less correlated to the actual qualities of the objects being evaluated.

Summarizing, given d_e and ρ_k , the conditional probability of the outcome $w_{e,k} = 0$ is defined as

$$\mathbb{P}(w_{e,k} = 0 | \rho_k, d_e) \triangleq F(\rho_k, d_e), \quad (1)$$

where $F(\rho, d)$ is an increasing function of d satisfying $F(\rho, 0) = 1/2$ and $F(\rho, -d) = 1 - F(\rho, d)$. In the following, we will consider two well-known evaluation models [12], [14]:

- **Thurstone model**, where η_k is a Gaussian random variable with standard deviation $1/\sqrt{2\rho_k^2}$. Therefore

$$F(\rho_k, d_e) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{\rho_k d_e}{\sqrt{2}} \right) \right] \quad (2)$$

- **BTL model**, where η_k is a Gumbel random variable with scale parameter $1/\rho_k$. Since the difference of two independent Gumbel rv's has a logistic distribution, we obtain

$$F(\rho_k, d_e) = \frac{e^{\rho_k d_e}}{1 + e^{\rho_k d_e}}. \quad (3)$$

Since, for both models, $F(\rho_k, d_e)$ is a function of the product $\rho_k d_e$ with a slight abuse of notation, in the following we will write $F(\rho_k d_e)$ instead of $F(\rho_k, d_e)$. Also, since $F(-\rho_k d_e) = 1 - F(\rho_k d_e)$, we can write in both cases

$$\mathbb{P}(w_{e,k} | \rho_k, d_e) = F((1 - 2w_{e,k})\rho_k d_e). \quad (4)$$

Note that Thurstone and BTL models defined above, coincide with the HTVC and HBTL model defined in [7]. A summary of the main notation introduced above is provided in Table I.

III. BAYESIAN CRAMÉR-RAO BOUND

In this section, we derive the Bayesian Cramér-Rao bound (BCRB) on any algorithm that jointly estimates object qualities and worker reliabilities. The BCRB allows to compute a

lower bound to the variance of any estimator of the unknown system parameters, when we assume to have access to prior knowledge.

To do so, we first denote by $\mathbf{q} = [q_1, \dots, q_N]^T$, $\mathbf{d} = [d_1, \dots, d_E]^T$ and $\boldsymbol{\rho} = [\rho_1, \dots, \rho_K]^T$ the vectors of object qualities, quality differences and worker reliabilities, respectively. We assume that $\boldsymbol{\rho}$ and \mathbf{q} are random vectors with i.i.d. entries whose a-priori distributions are

$$f_{\boldsymbol{\rho}}(\boldsymbol{\rho}) = \prod_{k=1}^K f_{\rho}(\rho_k), \quad \text{and} \quad f_{\mathbf{q}}(\mathbf{q}) = \prod_{i=1}^N f_q(q_i), \quad (5)$$

respectively, where $f_{\rho}(\rho)$ and $f_q(q)$ are the a-priori distributions of a generic worker reliability and of an object quality, respectively. We then observe that \mathbf{q} and \mathbf{d} satisfy the linear relationship

$$\mathbf{d} = \boldsymbol{\Gamma}^T \mathbf{q}, \quad (6)$$

where $\boldsymbol{\Gamma}$ is an $N \times E$ matrix whose e -th column, $e = 1, \dots, E$, has a 1 in row i_e , a -1 in row j_e , and 0 elsewhere. Let \mathcal{G} be the undirected graph whose N nodes are in one-to-one correspondence with objects and whose E edges are in one-to-one correspondence with the object pairs in \mathcal{E} . Then, $\boldsymbol{\Gamma}$ can be seen as a (directed) incidence matrix of \mathcal{G} . The joint a-priori distribution of the distances, $f_{\mathbf{d}}(\mathbf{d})$, and the corresponding marginals, $f_d(d_e)$, $e = 1, \dots, E$, follow trivially from (6).

The maximum a posteriori (MAP) joint estimate of quality distances and worker reliabilities can be written as

$$\begin{aligned} & \{\hat{\boldsymbol{\rho}}, \hat{\mathbf{d}}\} \\ &= \arg \max_{\boldsymbol{\rho}, \mathbf{d}} \log \mathbb{P}(\boldsymbol{\rho}, \mathbf{d} | \mathcal{W}) \\ &= \arg \max_{\boldsymbol{\rho}, \mathbf{d}} \log [\mathbb{P}(\mathcal{W} | \boldsymbol{\rho}, \mathbf{d}) f_{\boldsymbol{\rho}}(\boldsymbol{\rho}) f_{\mathbf{d}}(\mathbf{d})] \\ &= \arg \max_{\boldsymbol{\rho}, \mathbf{d}} \left(\sum_{k=1}^K \sum_{e \in \mathcal{E}_k} \log F(x_{e,k}) + \log f_{\boldsymbol{\rho}}(\boldsymbol{\rho}) + \log f_{\mathbf{d}}(\mathbf{d}) \right) \\ &= \arg \max_{\boldsymbol{\rho}, \mathbf{d}} \left(\sum_{e=1}^E \sum_{k \in \mathcal{K}_e} \log F(x_{e,k}) + \log f_{\boldsymbol{\rho}}(\boldsymbol{\rho}) + \log f_{\mathbf{d}}(\mathbf{d}) \right), \end{aligned} \quad (7)$$

where we defined $x_{e,k} \triangleq (1 - 2w_{e,k})\rho_k d_e$, we have exploited (4) and assumed that the evaluation outcomes are independent, so that

$$\mathbb{P}(\mathcal{W} | \boldsymbol{\rho}, \mathbf{d}) = \prod_{k=1}^K \prod_{e \in \mathcal{E}_k} F(x_{e,k}) = \prod_{e=1}^E \prod_{k \in \mathcal{K}_e} F(x_{e,k}).$$

Let $\boldsymbol{\theta} = [\mathbf{q}^T, \boldsymbol{\rho}^T]^T$ be the length- $(N+K)$ vector of unknown parameters. The BCRB allows to lower-bound the mean-square error (MSE) between $\boldsymbol{\theta}$ and its estimate $\hat{\boldsymbol{\theta}}$ achieved by MAP estimation and, a fortiori, by any other conceivable algorithm. The covariance matrix $\boldsymbol{\Sigma}$ of such estimate is given by

$$\boldsymbol{\Sigma} = \mathbb{E}_{\mathcal{W}, \boldsymbol{\theta}} \left[(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T \right]. \quad (8)$$

where we recall that \mathcal{W} is the set of random workers' answers.

The BCRB states that $\Sigma \succeq \mathbf{M}^{-1}$ (i.e., that matrix $\Sigma - \mathbf{M}^{-1}$ is positive semidefinite), \mathbf{M} being the Bayesian information matrix (BIM) defined by

$$\mathbf{M} \triangleq -\mathbb{E}_{\mathcal{W}, \theta} [\nabla_{\theta} \nabla_{\theta}^{\top} \log \mathbb{P}(\mathcal{W}, \theta)], \quad (9)$$

where ∇_{θ} represents the gradient with respect to the vector of parameters θ . The BCRB implies that

$$[\Sigma]_{i,i} \geq [\mathbf{M}^{-1}]_{i,i}, \quad (10)$$

i.e., it provides a lower bound to the MSE achieved by any estimator of the i -th unknown parameter, $i = 1, \dots, N + K$. The following proposition gives the BCRB for our scenario.

Proposition 1: For the above described scenario the BIM is given by

$$\mathbf{M} = \begin{bmatrix} \Gamma \Delta_q \Gamma^{\top} + \beta_q \mathbf{I}_N & \mathbf{0}_{N \times K} \\ \mathbf{0}_{K \times N} & \Delta_{\rho} + \beta_{\rho} \mathbf{I}_K \end{bmatrix}, \quad (11)$$

where Γ is defined through (6), Δ_q and Δ_{ρ} are diagonal matrices with diagonal elements

$$[\Delta_q]_{e,e} = |\mathcal{K}_e| \mathbb{E}_{\rho,d} \left[\frac{\rho^2 (F'(\rho d))^2}{F(\rho d)(1 - F(\rho d))} \right], e = 1, \dots, E, \quad (12)$$

and

$$[\Delta_{\rho}]_{k,k} = |\mathcal{E}_k| \mathbb{E}_{\rho,d} \left[\frac{d^2 (F'(\rho d))^2}{F(\rho d)(1 - F(\rho d))} \right], k = 1, \dots, K, \quad (13)$$

respectively. Finally:

$$\beta_q = -\mathbb{E}_q \left[\frac{\partial^2}{\partial q^2} \log f_q(q) \right] \quad \text{and} \quad \beta_{\rho} = -\mathbb{E}_{\rho} \left[\frac{\partial^2}{\partial \rho^2} \log f_{\rho}(\rho) \right]. \quad (14)$$

All averages are performed with respect to marginals of qualities, quality distances or reliabilities.

Proof: See Appendix B in the Supplemental Material. ■

Notice that, since the BIM is block diagonal, the BCRB on qualities and the BCRB on reliabilities can be computed separately, by inverting the two diagonal blocks. Moreover, the lower block is diagonal, so the BCRB on worker reliabilities can be explicitly written as

$$\mathbb{E} [|\hat{\rho}_k - \rho_k|^2] \geq ([\Delta_{\rho}]_{k,k} + \beta_{\rho})^{-1}. \quad (15)$$

Regarding the MSE on quality distances, it can be found that

$$\begin{aligned} \mathbb{E} [(\hat{\mathbf{d}} - \mathbf{d})(\hat{\mathbf{d}} - \mathbf{d})^{\top}] &= \Gamma^{\top} \mathbb{E} [(\hat{\mathbf{q}} - \mathbf{q})(\hat{\mathbf{q}} - \mathbf{q})^{\top}] \Gamma \\ &\succeq \Gamma^{\top} \left(\Gamma \Delta_q \Gamma^{\top} + \beta_q \mathbf{I}_N \right)^{-1} \Gamma. \end{aligned}$$

Finally, we particularize the above computations for the two worker models considered in this paper. For the Thurstone model, the argument of the average in (12) and (13) can be rewritten as

$$\frac{(F'(\rho d))^2}{F(\rho d)(1 - F(\rho d))} = \frac{2}{\pi} \frac{e^{-\rho^2 d^2}}{1 - \operatorname{erf}^2(\rho d / \sqrt{2})}, \quad (16)$$

while for the BTL model we have

$$\frac{(F'(\rho d))^2}{F(\rho d)(1 - F(\rho d))} = \frac{e^{\rho d}}{(1 + e^{\rho d})^2}. \quad (17)$$




Fig. 1. Block scheme of the QUIT algorithm. The variable ℓ denotes the iteration index. The algorithm takes as input the set of workers' answers \mathcal{W} and the graph \mathcal{G} and, at each iteration, outputs the quality estimates $\hat{\mathbf{q}}^{(\ell)}$, and the estimates of workers reliabilities $\hat{\rho}^{(\ell)}$. Iterations stop when condition (23) is met or when the maximum number of iterations I_{\max} is reached.

IV. RANKING ALGORITHMS FOR HETEROGENEOUS WORKERS

In this section, first we briefly review the most relevant classes of algorithms proposed in literature: the HAGD algorithm proposed in [7], which is, to the best of our knowledge, the only proposed ranking algorithm explicitly tailored for the heterogeneous case, and the LS algorithm proposed in [6] for a homogeneous case, which constitutes the kernel around which our proposed QUIT algorithm has been developed. Then, we introduce the single-stage version of QUIT algorithm.

A. Alternate-gradient approach (HAGD)

An algorithm that specifically targets ranking with heterogeneous workers is described in [7]. It implements an alternate-gradient optimization, which approximates joint ML estimation of objects qualities and worker reliabilities. More precisely, by exploiting (6), let

$$\mathcal{L}(\rho, \mathbf{q}) = -\log \mathbb{P}(\mathcal{W} | \rho, \Gamma^{\top} \mathbf{q}). \quad (18)$$

The algorithm in [7] starts from an arbitrary point in the (ρ, \mathbf{q}) space, call it $(\hat{\rho}^{(0)}, \hat{\mathbf{q}}^{(0)})$ and works iteratively. At iteration ℓ , $\ell = 1, 2, \dots$ it performs the following three steps:

$$\tilde{\mathbf{q}}^{(\ell)} = \hat{\mathbf{q}}^{(\ell-1)} - \lambda_q \nabla_{\mathbf{q}} \mathcal{L}(\rho, \mathbf{q}) \Big|_{\hat{\rho}^{(\ell-1)}, \hat{\mathbf{q}}^{(\ell-1)}} \quad (19)$$

$$\hat{\mathbf{q}}^{(\ell)} = \tilde{\mathbf{q}}^{(\ell)} - \frac{1}{N} \mathbf{1}_N^{\top} \tilde{\mathbf{q}}^{(\ell)} \quad (20)$$

$$\hat{\rho}^{(\ell)} = \hat{\rho}^{(\ell-1)} - \lambda_{\rho} \nabla_{\rho} \mathcal{L}(\rho, \mathbf{q}) \Big|_{\hat{\rho}^{(\ell-1)}, \hat{\mathbf{q}}^{(\ell-1)}}, \quad (21)$$

where λ_q and λ_{ρ} are positive step sizes. Iterations stop when a predetermined maximum number of iterations have been performed.

B. Least-square approach

In this subsection, we briefly recall a family of algorithms based on Least-Square (LS) estimation of object qualities, which have been introduced in [6] for the homogeneous case, i.e., when all workers are characterized by the same reliability ρ . These algorithms start from an initial unbiased estimate of distance d_e , given by $\hat{d}_e = \frac{1}{\rho} F^{-1}(\hat{p}_e)$, where $\hat{p}_e = 1 - |\mathcal{K}_e|^{-1} \sum_{k \in \mathcal{K}_e} w_{e,k}$ is an unbiased estimate of the probability $F(\rho d_e)$. From the noisy estimates $\hat{\mathbf{d}} = [\hat{d}_1, \dots, \hat{d}_E]^{\top}$, the

estimates $\hat{\mathbf{q}} = [\hat{q}_1, \dots, \hat{q}_N]^T$ of $\mathbf{q} = [q_1, \dots, q_N]^T$ are then obtained by solving the following LS optimization problem

$$\hat{\mathbf{q}} = \arg \min_{\mathbf{x}} \sum_{e \in \mathcal{E}} \omega_e \left(x_{i_e} - x_{j_e} - \hat{d}_e \right)^2, \quad (22)$$

where ω_e are arbitrary positive weights, whose setting is discussed in [6, Sect. 3.1].

C. The QUITe algorithm

Now, we propose an algorithm named QUITe, specifically tailored to deal with heterogeneous workers, which iteratively refines the estimation of quality distances and worker reliabilities. It consists in an alternate optimization of \mathbf{d} given $\boldsymbol{\rho}$ and of $\boldsymbol{\rho}$ given \mathbf{d} , along a total number of I_{\max} iterations, with the possibility of early termination thanks to a stopping condition. Unlike the HAGD algorithm of [7], where each iteration corresponds to a single gradient step, both with respect to distances and with respect to reliabilities, each iteration of QUITe performs a complete MAP optimization with respect to each one of the two sets of variables, given the other. As a consequence, in a single iteration, QUITe outperforms HAGD, at the price of an increased computational complexity. In the following, the symbol ℓ is used to denote the iteration index, and the value of the generic variable v at iteration ℓ is indicated as $v^{(\ell)}$.

Algorithm 1 reports the pseudo-code for QUITe, whereas Figure 1 shows its simplified block scheme. The algorithm takes as inputs the set of workers' answers \mathcal{W} , the graph \mathcal{G} (through its incidence matrix $\boldsymbol{\Gamma}$), the workers' model $F(\cdot)$, the a-priori distributions on object qualities, distances and workers' reliabilities $f_q(\cdot)$, $f_d(\cdot)$ and $f_\rho(\cdot)$, respectively, and their support \mathcal{I}_q , \mathcal{I}_d and \mathcal{I}_ρ . QUITe also takes as input the maximum number of iterations I_{\max} and a threshold τ for the stopping condition.

The QUITe algorithm initializes the estimates of the object qualities, $\hat{\mathbf{q}}^{(0)}$, by randomly drawing it from the distribution $f_q(q)$, $q \in \mathcal{I}_q$. Then, at iteration $\ell = 1, 2, \dots$, the following operations are performed.

- First, QUITe infers rough estimates of the edge distances, $\boldsymbol{\delta}^{(\ell)} = [\delta_1^{(\ell)}, \dots, \delta_E^{(\ell)}]^T$, and of their variances $\boldsymbol{\sigma}^{(\ell)} = [\sigma_1^{(\ell)}, \dots, \sigma_E^{(\ell)}]^T$, on a per-edge basis. Such procedure is detailed in Section IV-C1; it exploits the workers' answers, \mathcal{W} , and the estimates of the workers' reliabilities at previous iteration, $\hat{\boldsymbol{\rho}}^{(\ell-1)}$, when available.
- Then, the estimates $\boldsymbol{\delta}^{(\ell)}$ and their variances, $\boldsymbol{\sigma}^{(\ell)}$, are combined together through a weighted LS algorithm exploiting the graph, in order to obtain estimates $\hat{\mathbf{q}}^{(\ell)}$ of the object qualities, as described in Section IV-C2.
- Next, new estimates of the distances, $\hat{\mathbf{d}}^{(\ell)}$, are finally obtained from $\hat{\mathbf{q}}^{(\ell)}$ through the incidence matrix $\boldsymbol{\Gamma}$, by applying (6).
- Finally, MAP estimates of the workers' reliabilities, $\hat{\boldsymbol{\rho}}^{(\ell)}$, are obtained from $\hat{\mathbf{d}}^{(\ell)}$, as detailed in Section IV-C3.

The algorithm stops when the normalized difference $\hat{\mathbf{q}}^{(\ell)} - \hat{\mathbf{q}}^{(\ell-1)}$ is sufficiently small, i.e., when

$$\frac{\|\hat{\mathbf{q}}^{(\ell)} - \hat{\mathbf{q}}^{(\ell-1)}\|_2}{N \|\hat{\mathbf{q}}^{(\ell-1)}\|_2} < \tau, \quad (23)$$

or when the maximum number of iterations I_{\max} is reached.

1) *Estimates of the edge distances and of their variance:* Here, we describe how the estimates $\boldsymbol{\delta}^{(\ell)}$ of the distances \mathbf{d} are obtained by the QUITe algorithm. At the first iteration ($\ell = 1$), since estimates of the workers' reliabilities are not yet available, we resort to a generalization of the procedure employed in [6] and briefly recalled in Sect. IV-B, taking into account that workers have different reliabilities. Specifically, for edge e , we first estimate the empirical probability of a 0 answer as

$$\hat{p}_e = 1 - |\mathcal{K}_e|^{-1} \sum_{k \in \mathcal{K}_e} w_{e,k}. \quad (24)$$

If all workers have the same reliability, ρ , using (1), the estimate of the edge distance δ_e can be obtained by $\delta_e^{(1)} = \frac{1}{\rho} F^{-1}(\hat{p}_e)$, $e = 1, \dots, E$. When workers have different random reliabilities, whose priors are $f_\rho(\rho)$, the estimate of the edge distance can be generalized as

$$\delta_e^{(1)} = G^{-1}(\hat{p}_e), \quad (25)$$

where the function $G(\delta)$ is given by

$$G(\delta) = \int_{\mathcal{I}_\rho} F(\rho\delta) f_\rho(\rho) d\rho. \quad (26)$$

The variance of such estimate is computed as [6, Section 3.1]

$$\sigma_e^{(1)} = \left(\frac{dG^{-1}(p)}{dp} \Big|_{p=\hat{p}_e} \right)^2 \frac{\hat{p}_e(1-\hat{p}_e)}{|\mathcal{K}_e|}. \quad (27)$$

The above procedure is denoted by the block labeled "INIT" in Figure 1.

For $\ell > 1$, since estimates $\hat{\boldsymbol{\rho}}^{(\ell-1)}$ of the workers' reliabilities are available, we can apply a MAP approach to obtain the estimates $\boldsymbol{\delta}^{(\ell)}$. Let us focus on object pair e and let \mathcal{W}_e be the set of workers' answers related to such pair. Then, according to (4), the log a-posteriori probability of the distance d_e , given the workers' reliabilities $\boldsymbol{\rho}$ and answers \mathcal{W}_e , is given by

$$\begin{aligned} & \log \mathbb{P}(d_e | \mathcal{W}_e, \boldsymbol{\rho}) \\ &= \log \frac{\mathbb{P}(\mathcal{W}_e | d_e, \boldsymbol{\rho}) f_d(d_e)}{\mathbb{P}(\mathcal{W}_e)} \\ &= \sum_{k \in \mathcal{K}_e} \log F(x_{e,k}) + \log f_d(d_e) - \log \mathbb{P}(\mathcal{W}_e), \end{aligned} \quad (28)$$

where we recall that $x_{e,k} \triangleq (1 - 2w_{e,k})\rho_k d_e$. At iteration ℓ , since the true reliabilities $\boldsymbol{\rho}$ are unknown and only their estimates $\hat{\boldsymbol{\rho}}^{(\ell-1)}$ are available, according to (28) we can write

$$\delta_e^{(\ell)} = \arg \max_{d \in \mathcal{I}_d} \sum_{k \in \mathcal{K}_e} \log F \left((1 - 2w_{e,k}) \hat{\rho}_k^{(\ell-1)} d \right) + \log f_{d,e}^{(\ell)}(d). \quad (29)$$

Moreover, since at step $\ell - 1$ the random variable d_e has been estimated having mean $\delta_e^{(\ell-1)}$ and variance $\sigma_e^{(\ell-1)}$, the prior for d_e has been replaced in (29) with a Gaussian distribution with such mean and variance, i.e., $f_{d,e}^{(\ell)}(d) = \mathcal{N}(d; \delta_e^{(\ell-1)}, \sigma_e^{(\ell-1)})$. The maximization involved in (29) must be performed numerically since, for both Thurstone and BTL models, it is not solvable analytically. However, for both models, the function $F(\cdot)$ is log-concave, so that there is a

single minimum, which can be found by efficient numerical methods.

The variance of the estimate can be obtained starting from (28). Note that, in order to maximize (28), we can set to zero its derivative w.r.t. to d_e , i.e., we compute

$$\begin{aligned} R(d_e) &= \frac{\partial}{\partial d_e} \left(\sum_{k \in \mathcal{K}_e} \log F(x_{e,k}) + \log f_d(d_e) - \log \mathbb{P}(\mathcal{W}_e) \right) \\ &= \sum_{k \in \mathcal{K}_e} (1 - 2w_{e,k}) \rho_k \frac{F'(x_{e,k})}{F(x_{e,k})} + \frac{f'_d(d_e)}{f_d(d_e)} = 0. \end{aligned} \quad (30)$$

Let d_e^* be the solution of (30). Then, d_e^* can be seen as an implicit function of the arguments $w_{e,k}$, $e \in \mathcal{K}_e$. With an abuse of notation, by considering $w_{e,k}$ as a continuous variable, we can approximate d_e^* as $d_e^* \approx \sum_{k \in \mathcal{K}_e} \frac{\partial d_e^*}{\partial w_{e,k}} w_{e,k}$. Since the random variables $w_{e,k}$ are independent, we can write

$$\mathbb{V}[d_e^*] \approx \sum_{k \in \mathcal{K}_e} \left(\frac{\partial d_e^*}{\partial w_{e,k}} \right)^2 \mathbb{V}[w_{e,k}], \quad (31)$$

where we recall that $\mathbb{V}[\cdot]$ is the variance operator and $w_{e,k}$ is a Bernoulli random variable with parameter $F(\rho_k d_e^*)$. Therefore

$$\mathbb{V}[w_{e,k}] = F(\rho_k d_e^*) [1 - F(\rho_k d_e^*)]. \quad (32)$$

Moreover, thanks to implicit differentiation, we can obtain the ‘‘derivative’’ in (31) as

$$\frac{\partial d_e^*}{\partial w_{e,k}} = - \frac{\Delta R(d_e^*) / \Delta w_{e,k}}{\partial R(d_e^*) / \partial d_e^*}. \quad (33)$$

The numerator of (33) can be expressed through the difference quotient, given by

$$\begin{aligned} \frac{\Delta R(d_e^*)}{\Delta w_{e,k}} &= R(d_e^*)|_{w_{e,k}=1} - R(d_e^*)|_{w_{e,k}=0} \\ &= - \frac{\rho_k F'(\rho_k d_e^*)}{F(\rho_k d_e^*) (1 - F(\rho_k d_e^*))}. \end{aligned} \quad (34)$$

The result in (34) has been obtained by using (30) and by observing that, for models as in (2) and (3), we have $F(-x) = 1 - F(x)$ and $F'(x) = F'(-x)$. The denominator of (33) is trivially given by

$$\begin{aligned} \frac{\partial R(d_e^*)}{\partial d_e^*} &= \sum_{k \in \mathcal{K}_e} \rho_k^2 \frac{F''(x_{e,k}^*) F(x_{e,k}^*) - F'(x_{e,k}^*)^2}{F(x_{e,k}^*)^2} \\ &\quad + \frac{f''_d(d_e^*) f_d(d_e^*) - f'_d(d_e^*)^2}{f_d(d_e^*)^2} \end{aligned} \quad (35)$$

$$\triangleq u(d_e^*), \quad (36)$$

where we have defined for brevity $x_{e,k}^* \triangleq (1 - 2w_{e,k}) \rho_k d_e^*$. Substituting (32)-(35) into (31), we obtain

$$\mathbb{V}[d_e^*] \approx \frac{\sum_{k \in \mathcal{K}_e} \rho_k^2 \frac{F'(\rho_k d_e^*)^2}{F(\rho_k d_e^*) (1 - F(\rho_k d_e^*))}}{(u(d_e^*))^2}. \quad (37)$$

The estimated variance for object pair e at iteration ℓ is then given by

$$\sigma_e^{(\ell)} = \mathbb{V}[d_e^*]|_{d_e^* = \delta_e^{(\ell)}, \rho_k = \hat{\rho}_k^{(\ell-1)}}. \quad (38)$$

Algorithm 1 QUIT algorithm

Require: \mathcal{W} , Γ , $F(\cdot)$, $f_q(\cdot)$, $f_d(\cdot)$, $f_\rho(\cdot)$, \mathcal{I}_q , \mathcal{I}_d , \mathcal{I}_ρ , I_{\max} , τ

Ensure: $\hat{\mathbf{q}}$

```

1: Initialize  $\hat{\mathbf{q}}^{(0)}$  with i.i.d. random entries according to the
   distribution  $f_q(q)$ ,  $q \in \mathcal{I}_q$ .
2: for  $\ell = 1, 2, \dots, I_{\max}$  do
3:   if  $\ell = 1$  then
4:     for  $e = 1, 2, \dots, E$  do
5:       Compute  $\hat{p}_e$  and  $\delta_e^{(1)}$  according to (24) and (25),
       respectively.
6:       Compute  $\sigma_e^{(1)}$  according to (27).
7:     end for
8:   else
9:     for  $e = 1, 2, \dots, E$  do
10:      Compute  $\delta_e^{(\ell)}$  using (29) and  $\sigma_e^{(\ell)}$  using (38).
11:    end for
12:   end if
13:   Use the weighted LS algorithm of [6], and compute
    $\hat{\mathbf{q}}^{(\ell)} = \text{LS}(\boldsymbol{\delta}^{(\ell)}, \boldsymbol{\sigma}^{(\ell)})$ .
14:   Update the estimate of  $\mathbf{d}$  using (6):  $\hat{\mathbf{d}}^{(\ell)} = \Gamma^T \hat{\mathbf{q}}^{(\ell)}$ .
15:   for  $k = 1, \dots, K$  do
16:     Compute  $\hat{\rho}_k^{(\ell+1)}$  using (42)
17:   end for
18:   if  $\|\hat{\mathbf{q}}^{(\ell)} - \hat{\mathbf{q}}^{(\ell-1)}\|_2 < \tau N \|\hat{\mathbf{q}}^{(\ell-1)}\|_2$  then
19:     break
20:   end if
21: end for
22: return  $\hat{\mathbf{q}} = \hat{\mathbf{q}}^{(\ell)}$ 

```

2) *Graph estimation of quality distances:* We now describe the LS algorithm that is used in Alg. 1, to obtain the quality estimates $\hat{\mathbf{q}}^{(\ell)}$, given current estimates of quality distances $\boldsymbol{\delta}^{(\ell)}$ and of their variances $\boldsymbol{\sigma}^{(\ell)}$. In particular, we extend the approach introduced in [6], where the quality estimates are obtained by solving the weighted LS problem

$$\hat{\mathbf{q}}^{(\ell)} = \arg \min_{\mathbf{x}} \sum_{e \in \mathcal{E}} \omega_e^{(\ell)} \left(x_{i_e} - x_{j_e} - \delta_e^{(\ell)} \right)^2 = \text{LS}(\boldsymbol{\delta}^{(\ell)}, \boldsymbol{\sigma}^{(\ell)}), \quad (39)$$

where $\omega_e^{(\ell)}$ are arbitrary positive edge weights. The problem can be solved by exploiting the graph structure, as described in [6, Equation (13)]. With the assumption that per-edge distance estimates are independent, the weights are chosen in order to minimize the variance of the quality estimates, i.e.,

$$\omega_e^{(\ell)} = \frac{1}{\sigma_e^{(\ell)}}. \quad (40)$$

3) *Estimates of the workers' reliabilities:* The procedure for estimating the reliabilities $\boldsymbol{\rho}$ is similar to that employed to estimate the edge distances. Again, we use a MAP approach, i.e., for worker k , we maximize the log a-posteriori probability

$$\begin{aligned} &\log \mathbb{P}(\rho_k | \mathcal{W}_k, \mathbf{d}) \\ &= \log \frac{\mathbb{P}(\mathcal{W}_k | \rho_k, \mathbf{d}) f_\rho(\rho_k)}{\mathbb{P}(\mathcal{W}_k)} \\ &= \sum_{e \in \mathcal{E}_k} \log F(x_{e,k}) + \log f_\rho(\rho_k) - \log \mathbb{P}(\mathcal{W}_k), \end{aligned} \quad (41)$$

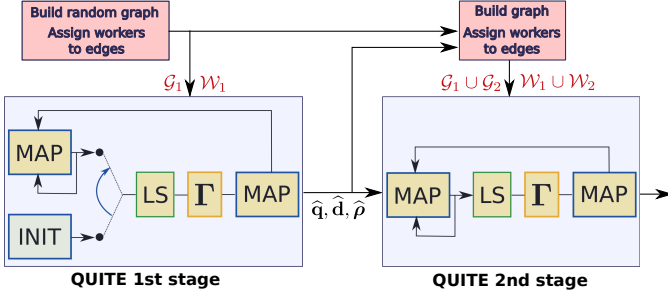


Fig. 2. Block scheme of the two-stage QUITe algorithm. The first stage is as in Figure 1 whereas the second stage has been slightly simplified by removing the block “INIT”. The output of the first stage is used to build a new graph \mathcal{G}_2 , and to initialize the second stage.

where we recall that \mathcal{E}_k is the set of object pairs evaluated by worker k and \mathcal{W}_k are the corresponding evaluation answers. Since the true distance \mathbf{d} are unknown and only their estimates $\mathbf{d}^{(\ell)}$ are available, according to (28), we can write

$$\rho_k^{(\ell)} = \arg \max_{\rho \in \mathcal{I}_\rho} \sum_{e \in \mathcal{E}_k} \log F \left((1 - 2w_{e,k}) d_e^{(\ell)} \rho \right) + \log f_\rho(\rho). \quad (42)$$

D. Theoretical guarantees for QUITe

A ranking algorithm is said asymptotically optimal whenever for any fixed $\epsilon > 0$ and $0 < \delta < 1$, it is able to obtain an (ϵ, δ) -PAC ranking with $O(\frac{N}{\delta^2} \log \frac{N}{\delta})$ comparisons, as N grows large. The asymptotical optimality of the LS algorithm is preserved also in the heterogeneous case. In particular we can claim that:

Proposition 1: A single-iteration version of the QUITe algorithm, according to which initial distances estimates δ_e are computed as in (25), and then quality estimates are obtained by solving the optimization problem (22), with weights $\omega_{i,j} = 1$, is asymptotical optimal as long as $\inf_x \frac{dG(x)}{dx} > 0$. In Appendix C of the Supplemental Material, we report a brief discussion of how proofs in [6] can be extended to the heterogeneous case.

V. TWO-STAGE QUITe ALGORITHM

In this section, we describe a two-stage protocol employing our QUITe algorithm as the fundamental and whose goal is to improve the reliability of the estimated ranking. Its block scheme is depicted in Figure 2. The first stage works exactly as previously described and works on a graph \mathcal{G}_1 and on the set of answers \mathcal{W}_1 . After obtaining an estimate of object qualities, $\hat{\mathbf{q}}$, and worker reliabilities, $\hat{\boldsymbol{\sigma}}$, a second stage consists in the following three steps:

- 1) a new graph \mathcal{G}_2 is created, with objects as nodes and edges in one-to-one correspondence to object pairs to be evaluated;
- 2) object pairs thus obtained are sent out for evaluation to the same set of workers that have evaluated the object pairs in the first stage;
- 3) after collecting the evaluations \mathcal{W}_2 , the QUITe algorithm is executed again on the joint set of evaluations

collected in the two stages, i.e., $\mathcal{W}_1 \cup \mathcal{W}_2$ to yield an improved ranking.

In the second stage, the QUITe algorithm has only a slight difference with respect to the version described in Section IV-C. Indeed, instead of initializing the distance estimates as in line 5 of Algorithm 1, the algorithm simply takes the estimates obtained at the output of the first stage and use them to find an estimate of distances, as in line 10 of the same algorithm. Then such distance estimates are used to obtain an improved estimate of worker reliabilities, and so on. In practice, in the second stage the block labeled “INIT” is removed, since it is unnecessary.

What makes the second stage a breakthrough is the fact that both the graph \mathcal{G}_2 and the assignment of object pairs to workers build upon the results of the first stage, and therefore can be better optimized, since more information is available. In particular, while the construction of graph \mathcal{G}_1 is necessarily blind (no information is available on object qualities), in the construction of graph \mathcal{G}_2 we can exploit available quality estimates. Similarly, available worker reliability estimates can drive the construction of second-stage assignment. In the following subsections, we explain in detail these procedures.

A. Building the second stage graph

While graph \mathcal{G}_1 in the first stage must be chosen as a random (regular) graph, the knowledge obtained in the first stage can be used to drive the choice of \mathcal{G}_2 . The idea is that the most important information for a reliable ranking estimation comes from evaluating pairs of objects that are close in the true ranking. Inspired by such idea, we propose to form graph \mathcal{G}_2 as follows.

Let $\hat{\mathbf{q}}$ be the estimates obtained at the output of the first stage and let $\hat{\boldsymbol{\pi}}$ be the permutation of integers $\{1, \dots, N\}$ such that $\hat{q}_{\hat{\pi}_1} > \hat{q}_{\hat{\pi}_2} > \dots > \hat{q}_{\hat{\pi}_N}$, i.e., $\hat{\boldsymbol{\pi}}$ is essentially the ranking estimated at the first stage. Then, given a maximum even node degree D , to build the graph \mathcal{G}_2 we proceed as follows: for $i = 1, \dots, N$ we define the set

$$\mathcal{S}_i = \left\{ i' \in \{1, \dots, N\} \setminus i \mid |i - i'| \leq \frac{D}{2} \right\} \quad (43)$$

and in \mathcal{G}_2 we connect with an edge objects $\hat{\pi}_i$ with all objects $\hat{\pi}_{i'}$, $i' \in \mathcal{S}_i$. By doing so, node $\hat{\pi}_i$ is connected (in \mathcal{G}_2) with at most D of its closest neighbors in the estimated ranking $\hat{\mathbf{q}}$. Note that the cardinality of \mathcal{S}_i depends on i . Indeed, objects located at the top or at the bottom of the estimated ranking $\hat{\boldsymbol{\pi}}$ are expected to have a smaller number of neighbors with respect to objects in the middle of the ranking list. Therefore, the graph \mathcal{G}_2 is not regular.

B. Assignment of object pairs to workers

Once the graph \mathcal{G}_2 is built, we need to assign object pairs to the pool of workers. For such assignment we must exploit the information we gathered in the first stage about the workers' reliabilities, i.e., $\hat{\boldsymbol{\rho}}$.

Consider a particular object pair, e , in the graph \mathcal{G}_2 and assume that it is assigned to a set of workers \mathcal{K}_e , all having the same reliability ρ . In this simplified case, we recall that

the estimate of the distance d_e based on the workers' answers, $w_{e,k}$, $k \in \mathcal{K}_e$, is given by $\hat{d}_e = \frac{1}{\rho} F^{-1}(\hat{p}_e)$, where $\hat{p}_e = 1 - \frac{1}{|\mathcal{K}_e|} \sum_{k \in \mathcal{K}_e} w_{e,k}$ is the unbiased estimate of the probability p_e , i.e., $\mathbb{E}[\hat{p}_e] = p_e$. Therefore, the estimation error $y_e = \hat{p}_e - p_e$ is a zero-mean random variable with variance

$$\mathbb{V}[y_e] = \mathbb{E}[(\hat{p}_e - p_e)^2] = \mathbb{E}[\hat{p}_e^2] - p_e^2 = \frac{p_e(1-p_e)}{|\mathcal{K}_e|}. \quad (44)$$

Correspondingly, let $z_e = \hat{d}_e - d_e$ the estimation error on the distance d_e . For small y_e , z_e is a random variable with zero mean and variance

$$\begin{aligned} \mathbb{V}[z_e] &\approx \left(\frac{1}{\rho} \frac{dF^{-1}}{dx} \Big|_{x=p_e} \right)^2 \mathbb{V}[y_e] \\ &= \left(\frac{1}{\rho} \frac{dF^{-1}}{dx} \Big|_{x=p_e} \right)^2 \frac{p_e(1-p_e)}{|\mathcal{K}_e|}. \end{aligned} \quad (45)$$

Now, under the BTL model, by using (3) in (45) we have:

$$\mathbb{V}[z_e] \approx H(\rho) = \frac{1}{\rho^2} \frac{[e^{-\rho d_e} + 2 + e^{\rho d_e}]}{|\mathcal{K}_e|} = \frac{2}{\rho^2} \frac{\cosh(\rho d_e) + 1}{|\mathcal{K}_e|}.$$

Note that $H(0) = H(\infty) = \infty$, and therefore the optimal value ρ_e^* of ρ for pair e , which is the one that minimizes $\mathbb{V}[z_e]$, is not extremal and can be obtained by imposing $\frac{dH(\rho)}{d\rho} = 0$, yielding $\rho_e^* \approx \frac{2.399}{|d_e|}$. A similar derivation can be carried out as well for the Thurstone model. Then, denoting $\hat{\rho}_e^* = \rho_e^*|_{d_e=\hat{d}_e}$, (where \hat{d}_e is an estimate of d_e provided by the first QUITe stage) the rule we propose for the assignment of object pairs to workers is as follows: assign workers to edges so to minimize the metric $\sum_{k=1}^K \sum_{e \in \mathcal{E}_k} |\hat{\rho}_k - \hat{\rho}_e^*|$. Such assignment can be easily implemented by complying with the following simple rule: "assign best workers to shortest links".

VI. NUMERICAL RESULTS

In this section, we first provide numerical results of the performance of the single-stage QUITe algorithm and compare it with the HAGD algorithm proposed in [7]. Subsequently, we show the performance improvement achieved by the two-stage QUITe algorithm. All the results have been obtained by simulation.

A. Basic Simulation scenario

In our simulations we consider a number of objects, N , ranging from 40 to 400 and a set of workers of size $K = N$. For each value of N , we build a random regular graph with degree D , which is kept fixed for the whole experiment. The graph has then $E = ND/2$ edges. Each of these corresponds to an object pair evaluated by $M = \alpha K$ ($\alpha \leq 1$) different workers, with $\alpha \in \{0.1, 0.2, 0.5, 1\}$. Allocation of workers to object pairs is regular, so that each worker evaluates exactly $EM/K = \alpha E$ object pairs.

Workers obey either to the BTL or Thurstone model. Object qualities and worker reliabilities are i.i.d. and drawn from the uniform distributions $f_q(q) = \mathcal{U}[0, 1]$ and $f_\rho(\rho) = \mathcal{U}[1, 20]$, respectively, which are supposed to be known by the algorithm. Since such distribution functions have discontinuities, for the computation of the BCRB we have windowed it

with the Planck-taper [24] function, which approximates the uniform distribution $\mathcal{U}[a, b]$ with a finite-support probability density function of class C^∞ , given by

$$f_z^{\text{Planck}}(x) = \begin{cases} \frac{C_p}{1 + \exp\left\{z\left(\frac{1}{x-a} - \frac{1}{a+z-x}\right)\right\}}, & a < x < a+z, \\ C_p, & a+z \leq x \leq b-z, \\ \frac{C_p}{1 + \exp\left\{-z\left(\frac{1}{b-z-x} - \frac{1}{x-b}\right)\right\}}, & b-z < x < b, \\ 0, & \text{elsewhere} \end{cases}$$

where $C_p = 1/(b-a-z)$, the parameter $0 < z \leq (b-a)/2$ represents the smoothness of the function. In the BCRB computation, we have chosen $z = (b-a)/5$ to reasonably approximate the true a-priori distributions.

B. Single-stage QUITe

To show the performance of QUITe we rely on two metrics: the achieved MSE on the estimated object qualities and the error probability of the estimated ranking.

1) *Mean square error on the quality estimates*: In Figures 3 and 4, we set $D = 20$ and we show the MSE provided by QUITe for several different scenarios. To compute the MSE, we have adjusted the estimated quality values by performing an affine transformation, i.e., $\hat{q}_i \rightarrow A\hat{q}_i + B$, where B is the true quality value of the N -th object, which is assumed to be the reference in the LS algorithm, (i.e., imposing $\hat{q}_N = 0$), while A is a strictly positive scaling parameter whose value has been obtained through simulations, as the average over 100 different realizations of the optimal, MSE-minimizing scaling parameters. Notice that the above transformation does not affect the final ranking; indeed the rankings induced by the vectors $\hat{\mathbf{q}}$ and $A\hat{\mathbf{q}} + B$ are the same, for any $A > 0$ and B .

Figure 3 shows the MSE performance of QUITe, with the BTL worker model and $I_{\max} = 30$ iterations. For these experiments, we have not set a stopping threshold $\tau = 0$, so that QUITe always performs all the I_{\max} iterations. Curves are parameterized by the value of $\alpha = M/K$, i.e., the number of evaluations on each edge per worker. Dashed lines represent the respective BCRBs. As it can be seen, the MSE decreases with the number of objects N , and with the parameter α . The dependence on α is somehow trivial: indeed, by increasing α , while maintaining the other parameters constant, we increase the number of evaluations, and hence we reduce the MSE. The behavior with respect to N can be explained as follows. Since in our setting the total number of evaluations is $EM = 10\alpha N^2$ and the number of parameters to be estimated is $K + N = 2N$, the number of evaluations per unknown parameter is $\frac{10\alpha}{2}N$, which increases with N . Hence, the unknown parameters ρ and \mathbf{q} are estimated with increasing reliability as N increases and as the number of evaluations per edge increases. The ratio between the obtained MSE and the relative BCRB is always less than one order of magnitude, with a typical value of 3-5, and slightly improves with N .

Figure 4 shows the MSE performance of the QUITe algorithm, with the Thurstone worker's model. All the other parameters are same as for Figure 3. The behavior of the curves is similar as for the BTL model. However, the Thurstone model shows a slight increase in MSE values, especially for low N , compared to BTL.




Fig. 3. MSE performance of the QUIT algorithm versus the number of objects for $K = N$ workers, and BTL worker model. The graph is regular and has degree $D = 20$. $M = \alpha K$ evaluations are carried out for each object pair, with $\alpha = \{0.1, 0.2, 0.5, 1\}$. The algorithm stops after $I_{\max} = 30$ iterations. The performance of QUIT (solid lines) is compared against the BCRB (dashed lines).




Fig. 4. MSE performance of the QUIT algorithm versus the number of objects for $K = N$ workers, and Thurstone worker model. The graph is regular and has degree $D = 20$. $M = \alpha K$ evaluations are carried out for each object pair, with $\alpha = \{0.1, 0.2, 0.5, 1\}$. The algorithm stops after $I_{\max} = 30$ iterations. The performance of QUIT (solid lines) is compared against the BCRB (dashed lines).

2) *Error probability on the estimated ranking:* Now we show the performance of QUIT in terms of the reliability of the estimated ranking and compare it with the HAGD algorithm introduced in [7] and discussed in Section IV-B. For the same scenario considered before, we have performed Monte Carlo simulations and counted an error whenever the estimated ranking was not ϵ -quality, as defined in Section II, with $\epsilon = 0.06$.

For QUIT, we have set a maximum number of iterations equal to $I_{\max} = 50$ and a stopping threshold $\tau = 10^{-5}$. For the HAGD algorithm, we have performed at most $I_{\max} = 1000$ iterations, with the possibility of stopping iterations in




Fig. 5. Error probability of HAGD and QUIT algorithms, plotted versus the number of objects for $K = N$. Workers obey to the BTL model, with $\alpha = 1/2$ and ϵ -PAC=0.06. The graph degree is $D = 20$.

the same way as for QUIT, and with threshold $\tau = 10^{-5}$, as well. Since the authors in [7] do not provide details on how to set the step sizes specified in (19) and (21), we have set it as $\lambda_q = \lambda_\rho = \frac{N}{5}$, since we have observed that scaling with N is beneficial to performance.

Figures 5 and 6 report a comparison between the performance of QUIT and of HAGD in the BTL and Thurstone scenarios, respectively, with $\alpha = \frac{1}{2}$. For QUIT, we also show the performance obtained at the first iteration ($\ell = 1$), i.e., considering the estimates $\hat{q}^{(1)}$ in Algorithm 1.

First of all, we observe that the relative performance of QUIT improves significantly with N , similarly to what observed for the MSE in Figures 3 and 4. For small values of N , QUIT performs similarly to HAGD, with a slight advantage of the latter with respect to the former under the BTL model and a more significant advantage of the former with respect to the latter under the Thurstone model. However, as N increases (i.e., for $N > 200$), the relative performance of the HAGD algorithm significantly worsens, and for $N > 300$ HAGD is outperformed even by the first iteration of QUIT. Moreover, observe that the two QUIT curves referring respectively to the case $I_{\max} = 1$ and $I_{\max} = 50$ are almost parallel, and QUIT with $I_{\max} = 50$ outperforms QUIT with $I_{\max} = 1$ by an approximately constant factor close to 4. The behavior of QUIT with $I_{\max} = 1$ is justified by Proposition 1, which states its asymptotic optimality.

The same conclusions can be drawn from Figure 6, where the system parameters of Figure 5 are employed, but with the Thurstone worker model instead. In this case, QUIT is never outperformed by the HAGD algorithm, while the gap between QUIT and HAGD increases with N (again, for $N \geq 300$ even the first iteration of QUIT provides better error probability than HAGD).

As a final remark please observe that the two models (Thurstone and BTL) are not directly comparable, since the noise term η_k , introduced in Section II, has a different distribution (Gaussian vs Gumbel), and even their variances do not match for the same value of the reliability parameters ρ_k .

3) *Impact of model mismatch:* Our paper, like most of the existing works on the subject, assumes that the workers'




Fig. 6. Error probability of HAGD and QUITE algorithms, plotted versus the number of objects, for $K = N$. Workers obey to the Thurstone model, with $\alpha = 1/2$ and $\epsilon\text{-PAC}=0.06$. The graph degree is $D = 20$.

behavior is known by the system. Clearly, if such behavior is unknown, QUITE (as other similar algorithms) can be still applied by using an unverified model (i.e. making an assumption on $F(\rho, d)$ in (1)). By doing so, however the ranking algorithm may incur in a penalty due to model mismatch. In Fig. 5 we also show the error probability of the QUITE algorithm with $I_{\max} = 50$, if it assumes a Thurstone behavioral model for the workers while they behave according to BTL. As expected, the mismatch impacts negatively on the performance of QUITE. However, we observe that, even in presence of model mismatch, QUITE with $I_{\max} = 50$ significantly outperforms QUITE with $I_{\max} = 1$ and, for N large enough, the HAGD algorithm.

4) *Impact of the graph structure:* Finally, Figure 7 explores the impact of the graph structure on QUITE performance. We set $N = 200$ and, to be fair, we keep fixed the total number of workers' evaluations, which is equal to $\alpha KND/2 = CN^2/2$, where $C = \alpha D$, and we vary the degree D of the regular graph. We show performance curves parameterized by different values of C . Observe that all curves start at $D = C$ since $\alpha = C/D$ cannot be larger than 1. Figure 7 shows that the choice of the graph has a significant impact on the QUITE performance. In particular, the performance improves significantly by increasing the graph degree under both the BTL and Thurstone models. Indeed, the LS algorithm tends to become more efficient as the degree of the graph increases, since to infer the object qualities it can resort on a larger number of distance estimates, even if the accuracy of every individual estimate decreases.

C. Results with real-world datasets

In this subsection, we show the performance of QUITE when considering a real-world scenario, where the “evaluations” are the outcome of experiments, and not synthetically generated by simulations. In particular, following [6], we consider five recent seasons of the English Premier League. For each season, the participating $N = 20$ football teams are the objects. The true ranking is assumed to be the final




Fig. 7. Error probability provided by QUITE versus the graph degree D , for both the Thurstone and BTL worker models. The number of objects is $N = 200$ and the ranking is ϵ -quality with $\epsilon = 0.06$.

season ranking. We build up a complete graph, whose edges are the matches between every pair of football teams. The match between team i and team j is considered as lasting for 180 minutes, since it includes both the round when i is at home and the round where i is away. If team i has scored $g_{i,j}$ goals in the match against team j , we count $K_{i,j} = g_{i,j} + 1$ evaluations in favor of i when compared to j . Adding one evaluation to all teams in all comparisons avoids dropping edges in the graph, so that the algorithm performance is improved, as already noticed in [6]. Out of the total $M_{i,j} = K_{i,j} + K_{j,i}$ evaluations between i and j , those that favor the team that ranks higher in the true ranking are considered correct, while the others are wrong evaluations.

We consider a number of workers equal to the number of teams, i.e., $K = N = 20$. We consider a hammer-spammer worker model, in which 4 workers are hammers (they always correctly evaluate object pairs) and the remaining 16 are spammers (they make an evaluation error with probability $1/2$, independently of the qualities of the evaluated pairs). The assignment of evaluations to workers is performed accordingly, i.e., hammers are assigned only correct evaluations, while spammers are assigned on average 50% correct and 50% wrong evaluations. We also apply the constraint of not assigning more than one evaluation for a given pair to the same worker. The QUITE algorithm assumes uniform a-priori distributions on $[0, 1]$ and $[1, 20]$, for qualities and reliabilities, respectively. QUITE has been run with a stopping threshold $\tau = 10^{-4}$ and for both the Thurstone and BTL models, to see the influence of the underlying worker model. The results have been plotted in terms of the Kendall tau distance, which counts the number of inversions in the estimated ranking with respect to the true ranking, i.e. the number of pairs (i, j) for which i is ranked better than j in the true ranking and worse than j in the estimated one.

Results are shown in Fig. 8 where blue and red bars refer to QUITE performance for the workers' BTL model, while yellow and purple bars refer to Thurstone model. For both




Fig. 8. Distance between true and estimated ranking for five Premier League seasons. The y -axis is the Kendall tau distance (number of inversions) between the final season ranking and the output of the QUITE algorithm, for different choices of the model. Comparison with the first QUITE iteration, where the workers’ reliabilities are not estimated.

models, we report both the performance of QUITE after convergence, and the performance after the first QUITE iteration, without estimates of the workers’ reliabilities, which is akin to the weighted LS algorithm from [6]. As it can be seen, the performance of QUITE is much better after estimating the workers’ reliabilities, with average Kendall distances being always lower than 3. After the first iteration, the estimate considerably changes from season to season, ranging from 6 (season 15/16) to 22 (season 16/17). This is a further confirmation that an estimate of workers’ reliabilities can considerably improve the quality of the final ranking. Finally, the curves show that the QUITE algorithm is insensitive, at least for the considered scenario, to the choice of the worker model.

D. Single-stage QUITE versus two-stage

We now evaluate the performance of the two-stage QUITE algorithm described in Section V, showing the advantage with respect to single-stage QUITE. In Figure 9, we consider the BTL model, with worker reliability drawn from the uniform distribution $\mathcal{U}[1, 20]$. The parameters for the first stage are the same as for Figure 5, except that the graph \mathcal{G}_1 has degree $D = 10$. The same graph degree has also been chosen for building the graph \mathcal{G}_2 , using the procedure detailed in Section V-A³. For both stages the number of evaluations per pair is set to $M = \alpha K$ with $\alpha \in \{0.5, 1\}$. To be fair, in the figure we compare the performance of the two-stage version of QUITE against a single-stage QUITE with graph degree $D = 20$, so as to keep constant the total number of workers’ evaluations. In counting errors, we have set $\epsilon = 0.04$ for $\alpha = 1/2$ and $\epsilon = 0.02$ for $\alpha = 1$.

As for assigning workers to pairs in the second stage, following the discussion of Section V-B, we have opted for

³For simplicity, every object pair that was already evaluated in the first stage is excluded from \mathcal{G}_2 .




Fig. 9. Comparison of single-stage and two-stage QUITE algorithms in the BTL scenario by keeping fixed the total number of edges in the graph. In the single-stage case the graph is regular and has degree $D = 20$. In the two-stage case the graph \mathcal{G}_1 is regular with $D = 10$ and \mathcal{G}_2 is build according to V-A with $D = 10$. $M = \alpha K$ evaluations are carried out for each object pair, with $\alpha = \{0.5, 1\}$.

the following simple assignment rule. By using the estimates $\hat{\rho}$ obtained from the first stage, we sort the workers in decreasing order of their estimated reliabilities and we partition them into $1/\alpha$ subsets of M contiguous (in terms of reliability values) workers each. Notice that $1/\alpha$ is an integer. Then, using \hat{q} obtained from the first stage, we sort the object pairs of \mathcal{G}_2 in increasing order of their estimated distances and we partition them into $1/\alpha$ subsets of contiguous pairs⁴. Finally, we assign the i -th subset of pairs to the i -th subset of workers, $i = 1, \dots, 1/\alpha$. In this way, the shortest-distance pairs are evaluated by the most reliable workers, as estimated by the first stage.

From Figure 9, it is clear how two-stage QUITE improves on single-stage QUITE. This is particularly evident for $\alpha = 1$, and $\epsilon = 0.02$, where the single-stage QUITE exhibits an error probability always larger than 0.7, while the two-stage QUITE, which relies on first-stage estimates to build the second-stage graph, achieves a considerably better performance, with an error probability below 2% for $N = 200$.

Figure 10 shows the comparison between single-stage and two-stage QUITE in the Thurstone scenario. All the other parameters are the same as for Figure 9. Also in this case, the error probability dramatically decreases of even an order of magnitude, by choosing a two-stage strategy. The advantage of using two-stage QUITE is especially relevant for the case $\alpha = 1$, $\epsilon = 0.02$. For both BTL and Thurstone scenarios, the gain provided by the second stage seems to lie in adding new edges between objects with a similar estimated quality, as estimated by the first stage.

VII. ALGORITHM COMPLEXITY

In this section, we analyze the complexity of QUITE in order sense, with the same choice of parameters of the previous

⁴Since the number of pairs may not be divisible by $1/\alpha$, we may need to add some dummy pairs.




Fig. 10. Comparison of single-stage and two-stage QUITe algorithms in the Thurstone scenario by keeping fixed the total number of edges in the graph. In the single-stage case the graph is regular and has degree $D = 20$. In the two-stage case the graph \mathcal{G}_1 is regular with $D = 10$ and \mathcal{G}_2 is build according to V-A with $D = 10$. $M = \alpha K$ evaluations are carried out for each object pair, with $\alpha = \{0.5, 1\}$.

section. First of all, we notice that the overall process is made up of two distinct phases: the collection of evaluations from the workers, and the actual algorithmic part devoted to the estimation of the object ranking.

- The first phase is typically the most expensive and time-consuming one. Reaching out a pool of workers, submitting the evaluation tasks to each of them and getting feedback takes a time on the order of days, if not weeks. The number of evaluations, which accounts for the monetary cost of the process, is arbitrary (typically at least $\Theta((N/\epsilon^2) \log(N/\delta))$). The computational complexity of this phase is low.
- The second phase, which consists in running Algorithm 1, has a computationally complexity per iteration that can be computed as follows. The operations at Lines 10 and 16 are $O(1)$, and they have to be performed for each object pair and worker, respectively, so that the total complexity is $O(N)$. The LS algorithm of Line 13 requires the inversion of a size- N matrix, an operation with complexity $O(N^3)$. Finally, the matrix multiplication in Line 14 is of order $O(N^2)$. Thus, the overall computationally complexity is $O(I_{\max} N^3)$. Since it is a polynomial complexity, and there is no real-time constraint, such complexity is largely affordable for any practical number of objects N , and translates in a time consumption negligible if compared to the first phase.

If QUITe is implemented in two stages, the time and cost complexity of both phases is roughly doubled. In the multi-stage case, the assignment of object pairs to workers, as described in Section VI-D, requires $O(N \log N) + O(N^2 \log N)$ operations to sort workers and object pairs, respectively.

A. Comparison with the HAGD algorithm proposed in [7]

Comparing in terms of complexity QUITe with the HAGD algorithm proposed in [7] is not immediate, because of the different number of iterations they need to reach a given performance. On the one hand, an iteration of QUITe has complexity $O(N^3)$ which is significantly higher than a single iteration of the HAGD algorithm (i.e., $O(N)$). On the other hand, while the number of iterations required by QUITe to reach convergence is rather limited and independent on N , for the HAGD algorithm in [7] the number of iterations is much larger and exhibits a clear dependence on N . In the provided results, we set $I_{\max} = 50$ for QUITe, although in most cases the number of iterations required to reach convergence is significantly smaller, typically less than 10, whereas for HAGD we have set $I_{\max} = 1000$.

In Appendix D of the Supplemental Material, we show a comparison between the HAGD algorithm and QUITe, both single-stage and double-stage, in terms of run time of a single instance on an ordinary workstation.

VIII. CONCLUSION

In this paper, we have faced with the problem of ranking a set of objects by means of the evaluations provided by a pool of heterogeneous workers. In particular, our work answers positively the following question: “Does an estimate of the workers’ reliabilities improve ranking accuracy?”. Not only: we have proposed QUITe, an iterative algorithm that comes into two flavors: single-stage or two-stage.

Even single-stage QUITe improves on algorithms from the literature. Two-stage QUITe capitalizes on the first stage estimates to ask further partially informed evaluations to the same pool of workers, yielding a dramatic improvement, in terms of ranking correctness.

Our work can be extended in several directions, for example, it may be worth of further investigations the design and analysis of algorithms that take in input partial rankings among multiple objects provided by workers, as well as, algorithms requiring workers to express soft (non binary) preferences between objects in pairs.

REFERENCES

- [1] M. Richardson, E. Dominowska, and R. Ragno, “Predicting clicks: Estimating the click-through rate for new ads,” in *Proceedings of the 16th International Conference on World Wide Web*, ser. WWW ’07. New York, NY, USA: ACM, 2007, pp. 521–530.
- [2] R. Herbrich, T. Minka, and T. Graepel, “Trueskill™: A bayesian skill rating system,” in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. C. Platt, and T. Hoffman, Eds. MIT Press, 2007, pp. 569–576.
- [3] N. B. Shah and M. J. Wainwright, “Simple, robust and optimal ranking from pairwise comparisons,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 7246–7283, 2017.
- [4] A. d’Aspremont, M. Cucuringu, and H. Tyagi, “Ranking and synchronization from pairwise measurements via SVD,” *arXiv preprint arXiv:1906.02746*, 2019.
- [5] M. Falahatgar, Y. Hao, A. Orlitsky, V. Pichapati, and V. Ravindrakumar, “Maxing and ranking with few assumptions,” in *Advances in Neural Information Processing Systems*, 2017, pp. 7060–7070.
- [6] E. Christoforou, A. Nordio, A. Tarable, and E. Leonardi, “Ranking a set of objects: A graph based least-square approach,” *IEEE Transactions on Network Science and Engineering*, vol. 8, no. 1, pp. 803–813, 2021.

- [7] T. Jin, P. Xu, Q. Gu, and F. Farnoud, "Rank aggregation via heterogeneous thurstone preference models," in *thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [8] Y. Liu, E. X. Fang, and J. Lu, "Lagrangian inference for ranking problems," *Operations research*, vol. 71, no. 1, pp. 202–223, 2023.
- [9] H. Lou, T. Jin, Y. Wu, P. Xu, Q. Gu, and F. Farnoud, "Active ranking without strong stochastic transitivity," *Advances in neural information processing systems*, vol. 35, pp. 297–309, 2022.
- [10] W. Li, S. Shrotriya, and A. Rinaldo, " ℓ_∞ -bounds of the MLE in the BTL model under general comparison graphs," in *Uncertainty in Artificial Intelligence*. PMLR, 2022, pp. 1178–1187.
- [11] R. L. Plackett, "The analysis of permutations," *Applied Statistics*, pp. 193–202, 1975.
- [12] 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, pp. 324–345, 1952.
- [13] R. D. Luce, *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.
- [14] L. L. Thurstone, "The method of paired comparisons for social values," *The Journal of Abnormal and Social Psychology*, vol. 21, no. 4, p. 384, 1927.
- [15] G. Yao and U. Böckenholt, "Bayesian estimation of thurstonian ranking models based on the gibbs sampler," *British Journal of Mathematical and Statistical Psychology*, vol. 52, no. 1, pp. 79–92, 1999.
- [16] M. Lee, M. Steyvers, M. DeYoung, and B. Miller, "A model-based approach to measuring expertise in ranking tasks," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 33, no. 33, 2011.
- [17] Y. Wu, T. Jin, H. Lou, P. Xu, F. Farnoud, and Q. Gu, "Adaptive sampling for heterogeneous rank aggregation from noisy pairwise comparisons," in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, vol. 151. PMLR, 28–30 Mar 2022, pp. 11 014–11 036. [Online]. Available: <https://proceedings.mlr.press/v151/wu22f.html>
- [18] B. Szörényi, R. Busa-Fekete, A. Paul, and E. Hüllermeier, "Online rank elicitation for plackett-luce: A dueling bandits approach," in *Advances in Neural Information Processing Systems*, 2015, pp. 604–612.
- [19] M. Falahatgar, A. Orlitsky, V. Pichapati, and A. T. Suresh, "Maximum selection and ranking under noisy comparisons," *arXiv preprint arXiv:1705.05366*, 2017.
- [20] M. Falahatgar, A. Jain, A. Orlitsky, V. Pichapati, and V. Ravindrakumar, "The limits of maxing, ranking, and preference learning," in *International Conference on Machine Learning*, 2018, pp. 1426–1435.
- [21] 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. 217, pp. 1–40, 2016.
- [22] X. Chen, P. N. Bennett, K. Collins-Thompson, and E. Horvitz, "Pairwise ranking aggregation in a crowdsourced setting," in *Proceedings of the sixth ACM international conference on Web search and data mining*. ACM, 2013, pp. 193–202.
- [23] K. Li, X. Zhang, and G. Li, "A rating-ranking method for crowdsourced top-k computation," in *SIGMOD '18: Proceedings of the 2018 International Conference on Management of Data*, May 2018.
- [24] L. W. Tu, "Bump functions and partitions of unity," in *An Introduction to Manifolds*. Springer, 2008, pp. 127–134.



Alessandro Nordio (S'00-M'03) is a Researcher with the Institute of Electronics, Computer and Telecommunication Engineering of the Italian National Research Council. In 2002 he received the Ph.D. in Telecommunications from EPFL, Lausanne, Switzerland. From 2002 to 2009 he was a post-doc researcher with the Electronic Department of Politecnico di Torino, Italy. His research interests are in the field of signal processing, wireless sensor networks, theory of random matrices, and crowdsourcing systems.



crowdsourcing systems.

Alberto Tarable (S'00-M'02) received the Laurea degree (summa cum laude) in 1998 and the Ph.D. degree in Electronic Engineering in February 2002, both from Politecnico di Torino. From 2002 to 2012, he worked as a researcher in the Department of Electronics and Telecommunications of Politecnico di Torino. From 2012, he is a researcher in the Institute of Electronics, Computer and Telecommunication Engineering of the Italian National Research Council. His research interests include physical-layer design of communication systems, QKD systems,



Emilio Leonardi received the Dr. Ing. degree in electronics engineering and Ph.D. degree in telecommunications engineering from the Politecnico di Torino, Turin, Italy, in 1991 and 1995, respectively. He is currently an Associate Professor with the Department of Electronics, Politecnico di Torino, Turin, Italy. His research interests include performance evaluation of computer networks and distributed systems, dynamics over social networks, and human centric computation.