Crowd-based Multi-Predicate Screening of Papers in Literature Reviews

Evgeny Krivosheev
University of Trento, Italy
evgeny.krivosheev@unitn.it

Fabio Casati
University of Trento, Italy
Tomsk Polytechnic University, Russia
fabio.casati@unitn.it

Boualem Benatallah
University of New South Wales,
Australia
boualem@cse.unsw.edu.au

ABSTRACT

Systematic literature reviews (SLRs) are one of the most common and useful form of scientific research and publication. Tens of thousands of SLRs are published each year, and this rate is growing across all fields of science. Performing an accurate, complete and unbiased SLR is however a difficult and expensive endeavor. This is true in general for all phases of a literature review, and in particular for the paper screening phase, where authors filter a set of potentially in-scope papers based on a number of *exclusion criteria*. To address the problem, in recent years the research community has begun to explore the use of the crowd to allow for a faster, accurate, cheaper and unbiased screening of papers. Initial results show that crowdsourcing can be effective, even for relatively complex reviews.

In this paper we derive and analyze a set of strategies for crowd-based screening, and show that an adaptive strategy, that continuously re-assesses the statistical properties of the problem to minimize the number of votes needed to take decisions for each paper, significantly outperforms a number of non-adaptive approaches in terms of cost and accuracy. We validate both applicability and results of the approach through a set of crowdsourcing experiments, and discuss properties of the problem and algorithms that we believe to be generally of interest for classification problems where items are classified via a series of successive tests (as it often happens in medicine).

KEYWORDS

human computation, classification, literature reviews

ACM Reference Format:

Evgeny Krivosheev, Fabio Casati, and Boualem Benatallah. 2018. Crowd-based Multi-Predicate Screening of Papers in Literature Reviews. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3178876.3186036>

1 INTRODUCTION

Systematic literature reviews (SLR) [9, 13, 19] are reviews that follow a predefined process aimed at achieving transparency and impartiality with respect to the sources analyzed, minimizing distortions, biases, and conflicts of interest [35]. They are one of the most important form of publications in science [36], and are the

basis for evidence-based practices and even government policies, from education to healthcare, as they pool results independently obtained from a number of research groups [11]. Recognizing their importance, the number of systematic reviews is growing steadily, with tens of thousands of publications per year in all fields.

The cornerstone of transparency and impartiality in SLRs lies in a formalized paper selection process. This is typically formed by a stated scope and goal of the review (e.g., "study the effect of regular physical exercises on progress of dementia in older adults, focusing only on papers describing randomized controlled trials"), translated by the authors into a corresponding query (a boolean expression that includes relevant keywords) that retrieves candidate papers from a database such as Scopus. To avoid missing papers, the query tends to be inclusive, which means that it returns hundreds or thousands of results [30] that are later screened by researchers based on predefined *exclusion criteria* (e.g., "filter out papers that do not measure cognitive decline"), typically down to a few dozens.

While extremely useful, SLRs are very time consuming in terms of both effort and elapsed time, and this is true also for the paper screening phase [10, 21, 33]. Furthermore, with hundreds of thousands of papers written every year, SLRs rapidly become outdated [4], and although they should be updated periodically, the effort for doing so often represents a barrier [37], so that it is not uncommon for reviews to miss 30% or 40% of relevant papers [4].

In this paper we explore the use of crowdsourcing in systematic reviews, and specifically in the filtering phase, where we screen candidate papers resulting from the initial literature search to identify papers to be included in the analysis. This is an instance of finite pool classification [30] and crowd screening problems [32] where we need to classify a finite set of objects while minimizing cost. The potential benefits of crowdsourcing here are in terms of a faster and cheaper screening (compared to screening by professionals) as well as increased transparency (process and votes can be made public if desired) and reduced risk of author bias. The crowd also brings diversity [40] and, as we experienced first hand, disagreement in the crowd may signal errors or ambiguities in the definition of exclusion criteria. Research in this area is still in its infancy, although a set of recent initial efforts [21, 28, 36, 39] present very encouraging results in terms of both quality and cost reduction with respect to expert screening costs, and show feasibility of crowd-based screening in various domains, including healthcare.

In the following we present a probabilistic model suitable for the criteria-based screening of papers typical of SLRs and propose a set of strategies for crowd-based screening. Our main contribution consists in an adaptive crowdsourcing algorithm that significantly outperforms baselines. The algorithm polls the crowd in small batches and estimates, at each iteration and for each item, i) the

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186036>

criterion for which getting one more crowd vote on the paper can more efficiently lead us to a classification decision, and ii) whether we should give up trying to classify this item, recognizing that the crowd cannot efficiently reach a decision and therefore it should be left to the authors for expert screening. This also means that the algorithm is robust to papers and criteria that are overly difficult for the crowd to classify, in that it does not needlessly spend money on them.

The model is the result of many iterations and variations of experiments on commercial crowdsourcing platforms (Amazon Mechanical Turk (AMT) and *CrowdFlower*¹). We then performed additional experiments to validate the effectiveness of the strategies. While we present the results in the context of SLRs because we validated the model and findings for this case, we believe that results can be generally of interest for classification problems where items are classified via series of successive tests, as it often happens in medicine, as well as for finite pool classification problems and crowd-based query optimization, where the crowd evaluates predicates (analogously to our exclusion criteria) that filter a set of tuples to compute the query results.

2 RELATED WORK

Our work builds on approaches in crowdsourcing in SLR but also more generally on works on crowd-based classification.

Crowdsourcing in Systematic Reviews. Recently, Brown and Allison [3] used crowdsourcing to, among other tasks, classify 689 abstracts based on a set of criteria using AMT. Authors report agreement in 75% of the abstract, based on two raters, and a third rater is used to break the tie in case of disagreement. The paper does not discuss optimal crowdsourcing strategies or algorithms to minimize errors, but points to the potential of crowdsourcing in analyzing literature.

Mortensen and colleagues crowdsourced paper screening [28] in four literature reviews, each with several criteria. Their aim was to explore feasibility and costs of crowdsourcing and they address the problem by measuring workers agreement in a set of tasks run on AMT for papers in the medical field. Their work differs from ours in that it does not propose algorithms to identify optimal crowdsourcing strategies. However, it contains interesting observations related to the importance of task design, to the cost-effectiveness of crowdsourcing even when the task is not optimized, and to the high degree of variability in workers' agreement from paper to paper and criteria to criteria (Fleiss' Kappa ranging from 0.5 to -0.03). This is consistent with our own studies (our papers are in a different scientific area) and we exploit this variability to optimize the cost/error tradeoff.

Krivosheev and colleagues [21] also present a model and strategies for crowdsourcing SLR. An interesting aspect of the model and approach here is that the authors model cost and loss (error) resulting from crowdsourcing task, attempt to estimate them at the start, and provide authors with a price/error trade-off that can be used to decide how much to invest in the task. We borrow several concepts from this work, such as the ability to provide an estimate and a set of alternatives to SLR author, although the model of this paper is limited to screening based on one criterion.

Nguyen et al. [30] adopt a mixed crowd+expert+machine learning approach with an active learning classifier, where papers to be labeled are iteratively chosen to minimize overall cost and loss, by comparing estimated loss of crowd classification versus expert classification. This paper is part of a trend trying to leverage AI in literature reviews, which we do not discuss further as it is not the focus of this paper.

In general all papers reports positive results and complement them with insights and guidelines for task design of even for the design of a dedicated crowdsourcing platform for SLR [3, 40] as well as investigate the use of crowd for other phases of interest for SLR such as information extraction [36]. Interestingly, the only exception is represented by a study performed with medical students as screeners rather than online crowd workers, which reports rather poor accuracy [29].

From these studies we also learn that workers' accuracy vary across criteria, which points to the need of adapting to the characteristics of each SLR, criterion, and crowd. Indeed, one of the main differences of our approach lies in the ability to focus the crowd on "low hanging fruits", that is, items and criteria that are statistically more efficient from the perspective of correctly excluding papers.

Although not focused on paper screening, we also mention a fascinating analysis by Law and colleagues trying to understand under which conditions do researchers resort to crowdsourcing [22]. Among the many interesting considerations lies the observation that crowdsourcing is viable only if both authors *and* reviewers find it acceptable. Paper screening in SLRs seem to fit the requirements for being acceptable by authors but it is equally important for the scientific community to provide solid evidence of the quality of crowdsourced screening if we want it to be accepted by reviewers - especially in fields where SLRs may form the base of policies and practices.

Crowdsourced Classification The problem discussed here is an instance of a finite pool classification problem [30] and specifically of crowdsourcing-based classification. This problem has been studied for hundreds of years now, dating back at least to the end of the 18th century, when the Marquis de Condorcet presented his *Jury Theorem*², stating that if each juror in a jury has an error rate lower than 0.5 and if *guilty vs innocent* votes are independent, larger juries reach more accurate results, and approach perfection as the jury grows.

From there, researchers from the AI, database, and human computation communities have proposed many of classification algorithms, mostly based on variations of majority voting where votes are counted differently based on estimated worker's accuracy. The seminal work of Dawid and Skene [5] and refinement by, among others, Whitehill [41], Dong et al [7], Li et al [23], and Liu et al [24, 25] model workers' accuracy - often with a confusion matrix - and then adopt variants of Expectation Maximization [6] to iteratively refine prior estimates of workers' accuracy and of labels. Approaches based on spectral methods [18] and maximum entropy [43] have also been proposed, and belief propagation has been recently shown [31] to be optimal under certain assumptions.

Prior work also addresses the issue of optimizations in terms of costs for obtaining labels and techniques to reduce cheating [8,

¹www.mturk.com and www.crowdfunder.com

²<http://www.stat.berkeley.edu/~mossel/teach/SocialChoiceNetworks10/ScribeAug31.pdf>

14, 15, 17, 34]. For example, Hirth and colleagues [15] recommend specific cheating detection and task validation algorithms based on the cost structure of the task.

We build over many of these approaches, and in fact we adopt prior art algorithms for estimating workers' accuracy and for assigning labels. Although classification algorithms are central to our overall problem, to a large extent they are for us a swappable component: Our goal is to, given a task design and a classification algorithm, identify how to efficiently query the crowd to minimize the number of labels needed to achieve the desired precision and recall in screening problems.

3 MODEL AND OBJECTIVE

We model the SLR screening problem as a set of papers (items) I to be classified by the screening phase as *included* (in scope) or excluded based on a set of exclusion criteria (predicates) $C = \{c_1, c_2, \dots, c_m\}$. A paper is excluded if at least one exclusion criterion applies, otherwise it is included. A typical SLR screens hundreds or thousands of papers with a handful of exclusion criteria. We focus on screening based on title and abstract, which is a classical first step screening, consistent with SLR guidelines [27].

In a crowdsourcing approach, we ask each crowd worker to look at one or more pairs (i, c) and state if exclusion criteria c applies to paper i . Following the mentioned literature, we model a worker's accuracy with a confusion matrix $A_{c,w}$ defining the probability of making correct and wrong classifications for each criterion c , thereby allowing us to model different accuracies when the true label is inclusion vs exclusion. Criteria can differ in *difficulty*. Some are easier to assess than others. Following Whitehill [41], we model difficulty as a positive real number d_c that, given an expected accuracy α_w of a worker w , skews the accuracy as follows: $\alpha_{c,w} = 0.5 + (\alpha_w - 0.5) * e^{-d_c}$. As the difficulty d_c grows, $\alpha_{c,w}$ goes to 0.5, corresponding to random selection, which we consider to be the lowest accuracy level³. Each criterion also has a *power* (also called selectivity) θ_c , defined as the percentage of papers to which the criterion applies (and hence need to be excluded). For each SLR and criterion, both accuracy and power are unknown a priori.

We assume the adoption of a general purpose crowdsourcing system with limited control on the kind of crowd we attract but with a near infinite pool of workers. We can however test workers by providing a number N_t of test questions (with gold answers provided by SLR authors), and count as valid only votes of workers who pass the test, thereby exercising some control over worker's accuracy (at a cost, as we specify later).

A crowdsourcing *strategy* is a set K of runs, where each run R^k collects $J_{i,c}^k$ votes for criterion c on item i . A run may seek votes on all criteria and all papers, or focus on a subset (that is, $J_{i,c}^k$ might be 0 for some items).

Tasks also have a cost, which is the unit cost UC for a (non-test) vote multiplied by the number of votes obtained. Although many systems allow not to pay for test answers, consistently with [21], we believe it is fair and ethical to also pay for test questions for workers who pass them. Furthermore, placing unreasonably many test questions is likely to result in low reputation scores for us and

hence in our ability to crowdsource. Concretely, this translates into considering a price per label PPL as follows (N_l is the number of valid judgments that a worker gives on non-test papers)

$$PPL = UC \cdot \frac{N_l + N_t}{N_l} \quad (1)$$

The correction factor approaches 1 when N_l is large compared to N_t . In practice our control on N_l can be limited by many factors (also depending on the crowdsourcing platform policies), such as dropouts, the presence of many concurrent workers that exhaust the available tasks, and more. We observe that tests are for us "simply" a knob to turn when trading costs for accuracy. Any other knob that accomplishes the same effect can be equivalently used in what follows.

In terms of outcome, key measures are the precision and recall of papers to exclude. We also borrow the concept of loss function from [21, 30] because it summarizes well the subjective perspective of the SLR authors. The *loss* = $lr * FE + FI$ is represented by the sum of false inclusions FI (papers that survived the screening phase but that should have been excluded instead) and false exclusions FE (filtered-out papers that should have instead been left in), where FE are weighed by a loss ratio lr denoting that false exclusion are lr times more "harmful" than false inclusion (filtering out a paper is often considered a much more serious error than a false inclusion which "simply" requires extra work by the authors). The loss ratio is the only parameter we ask the authors to set.

Many variations of the model and of loss function are possible, but these suffice for our purposes. Given the model, our objective is to identify and evaluate a set of efficient crowdsourcing strategies for each SLR that correspond to estimated pareto-optimal price/loss curves. With infinite money we can always arrive at a perfect classification (if workers' accuracy is above random and votes are independent), but the challenge is to classify efficiently and at a price/loss point that is acceptable to authors, who decide what price they are willing to pay and which loss they can tolerate. Based on this preference, the algorithm should set the relevant parameters of the crowdsourcing tasks and classification function. We next discuss how this can be done.

4 ALGORITHMS

4.1 Baseline single-run algorithms

Our set of baseline algorithms follows the methods applied in recent literature for crowdsourced classification in finite pool contexts and SLRs in particular. Specifically, as we are in the presence of incomplete information (we know neither the classification of the papers nor the accuracy of the workers), we leverage approaches such as TruthFinder [7] and Expectation Maximization (EM, [5]) to iteratively refine estimates of accuracy and class until convergence. In addition, simple majority voting is also commonly used as its performances are actually reasonable in finite pool classification [30].

Applying them to our problem, we proceed in a single run where we ask each worker to vote on all criteria C for a set of papers. Each worker provides at most N_l labels, and we collect J votes per criteria and per paper. Classification proceeds by evaluating each criterion $c \in C$ on each paper i and, based on the responses received,

³In this paper we do not consider the problem of accuracies below random, but we stress that they can occur in rare cases, for example if criteria are erroneously specified.

estimating with one of the mentioned algorithms the probability $P(i \in OUT_c)$ that paper i is classified as out by criterion c .

Once we have probabilities for each criterion, we compute the probability $P(i \in OUT)$ that a paper i should be excluded as the probability that at least one criterion applies (we assume criteria application is independent):

$$P(i \in OUT) = 1 - \prod_{c \in C} P(i \in IN_c) \quad (2)$$

The loss ratio skews our classification decision to err on the side of inclusion (for $lr > 1$). The expected loss per paper we suffer for an erroneous inclusion of a paper i is $P(i \in OUT)$, while for an erroneous exclusion it is $lr \cdot (1 - P(i \in OUT))$. This means that our threshold for classifying a paper as OUT is when these quantities are the same, that is, $P(i \in OUT) = lr / (lr + 1)$.

Altering the number of votes per worker N_l , votes per item J , and number of tests N_t will modify the expected price and loss. More tests ideally lead to more accurate workers, more labels mean more accurate classification, and more votes per person enable a more accurate estimation of a worker's accuracy. To analyze price vs loss, we simulate the behavior of the model with various values of N_l , J , N_t , and apply EM, TruthFinder (TF) or Majority Voting (MV) to classify papers, and compute the estimated loss. Since values of N_t and J correspond to a cost, we can also get the price tag corresponding to this loss. Out of this set of price/loss points, we can take the pareto-optimal ones and plot them so that authors can decide which one best fits their needs. As discussed there are cost penalties and practical constraints that do not allow us to set these parameters to arbitrarily high values, and values of N_t and J above 10 do not generate significant improvements [21], so the number of reasonable alternatives is fairly small. To simulate the data we need either to make assumptions on the crowd accuracy as well as on criteria power and difficulty, possibly based on prior knowledge, or to estimate these parameters [21] by crowdsourcing labels for a few papers (fifty papers already enable a good estimate as shown later). Figure 1 shows the results of applying the three mentioned algorithms for 3 and 5 labels per item and criterion (the caption describes simulation parameters). The impact of choosing a specific algorithm is relatively small with the exception of MV performing better when labels per paper and per worker are few, which is a known behavior [16]. The dots represent different number of tests (from 1 to 10) and the arrows shows the direction of growth, from top-left to bottom-right. Some points are Pareto-optimal, so in an interaction with SLR authors we would only show those points and ask for the preferred loss/price point.

The results vary slightly if the parameters of the problem and algorithms are different (such as different power, difficulty distributions across criteria, proportion of papers to be excluded, number of papers per worker). We discuss how quality and cost vary later in the paper when we compare and discuss algorithms.

4.2 Multi-Run Strategy by Criteria

The multi-run strategy follows the footsteps of the above-mentioned approaches for query optimization in crowd databases that identify the most selective criteria and query based on those first. The difference here is that we also estimate and consider accuracy (we do not want to query the crowd if this brings high disagreement,

Figure 1: Performance of classification algorithms. Simulation with 1000 papers, four criteria of power = $[c_1 = 0.14, c_2 = 0.14, c_3 = 0.28, c_4 = 0.42]$, $N_t = [2, 3, \dots, 10]$, $lr = 5$. Workers are assumed to be cheaters with probability 0.3, and the rest has uniform accuracy in (0.5-1). Accuracy on OUT papers are 10% higher, as seen in experiments.

Algorithm 1: M-Runs Algorithm

Input: Items I , Criteria C , loss ratio lr

Output: Classified items CI

- (1) $CI \leftarrow \{\}$, $UI \leftarrow I$, $thr = \frac{lr}{lr+1}$, $I^0 = 100$ randomly selected papers from I
- (2) **# Baseline iteration (Run 0)**
- (3) $V^0 \leftarrow$ collect J votes on I^0 for all criteria C
- (4) $CI^0 \leftarrow \text{classify_items}(V^0, thr)$
- (5) $CI \leftarrow CI \cup CI^0$, $UI \leftarrow UI - CI^0$
- (6) **foreach** $c \in C$
- (7) $\hat{\theta}_c \leftarrow \frac{\sum_{i \in I^0} P(i \in OUT_{T_c})}{|I^0|}$, $\hat{\alpha}_c = \frac{\sum_{w \in W_c^0} \alpha_{w,c}}{|W_c^0|}$ (W_c^0 is the set of workers who passed test questions and provided at least one label in baseline iteration)
- (8) **#Ranking**
- (9) $\text{criteria_order} \leftarrow \text{estimate_best_order}(\hat{\theta}, \hat{\alpha})$
- (10) **#M-Runs iterations**
- (11) **foreach** $c \in \text{criteria_order}$
- (12) $V^c \leftarrow$ collect J votes on UI on c
- (13) $CI_{out}^c \leftarrow \text{exclude_items}(V^c, thr)$
- (14) $CI \leftarrow CI \cup CI_{out}^c$, $UI \leftarrow UI - CI_{out}^c$
- (15) $CI_{in} \leftarrow$ tag UI as "IN items"
- (16) $CI \leftarrow CI \cup CI_{in}$
- (17) **return** CI

as it is less cost-effective), and that we work with a specified loss

function and a price vs loss trade-offs that are based on the authors' choice. The algorithms proceeds as follows.

Baseline iteration. We first estimate power and difficulty via a *baseline iteration* (run $k = 0$) on a randomly selected subset I^0 of the set of candidate papers I , as shown in Algorithm 1 (We will get back later in the paper about identifying how large should I^0 be).

In step 4 we classify items and estimate accuracy of each worker with a classification algorithm that also provides accuracy estimates such as TruthFinder (TF) [7]. As TF estimates class probabilities, we estimate the power of a criterion c as the expected value of the probability that the criterion applies, as shown in step 7. We refer instead to the difficulty of a criteria c through the average workers' accuracy on the given criteria, i.e., the average probability that a user, who passed the test questions, gives correct votes on that criteria.

Criteria ranking. Finding the best ordering is trivial if one criterion is more powerful *and* easier than another. Otherwise, different ordering may lead to price/loss points that are on the Pareto frontier and need to be shown to authors for decision. The number of criteria is often low so that considering permutations of all cases where the ordering is not trivial is tractable. We do so in step 8, by computing for each ordering the expected price and loss for different values of N_t and J . The computation of price and loss can be done as for the previous algorithm. Notice that the ordering is very important: given an ordering of criteria $OC = c_0, c_1, \dots, c_n$, the probability of erroneously excluding an item (probability of false exclusion, or PFE) is the probability of erroneously excluding it in the first round (on c_0), plus the probability of correctly including it after c_0 but erroneously excluding it after c_1 , and so on. More formally, denoting with PFE_c the probability of erroneous exclusion when processing criteria c and with PIN_c the probability of classifying a paper as IN on criteria c :

$$PFE = PFE_0 + \sum_{m=1,2,\dots,n} PFE_m \prod_{j=0}^{m-1} PIN_j \quad (3)$$

PFE therefore decreases with PIN, and in practice it decreases sharply if we screen high power criteria first, given that criteria powers over 30% are quite common.

Crowdsourcing iteration. The algorithm iterates through the criteria, excluding items (classified again based on TF in step 12).

The results of M-runs (orange) compared with the baseline single run algorithm (blue) are shown in Figure 2b and 2c, showing loss and precision vs price for different values of N_t and J . The simulation parameters are the same as previously described. The savings are of approximately 20%, and are in general higher if the criteria diversity in terms of power and accuracy is higher.

4.3 Short Adaptive Multi-Run

The previous algorithms apply the same strategy to all papers left to classify. The Short Adaptive Multi-Run algorithm (SM for short)(Algorithm 2) defines instead an *individual* strategy for each item to be labeled, aimed at identifying the shortest path to decision. The idea is that as we collect votes we understand more about the statistical properties of the overall SLR task (such as criteria power and difficulty) and also of each specific paper, based on the votes obtained for that paper so far. Therefore, we can estimate

which is the criterion to test next for each paper by maximizing the probability of (correctly) classifying it as out in the next run, and we can even decide to give up on a paper (leaving it in) because we realize it is too hard (or too expensive) for the crowd to reach consensus or because the probability that we will classify it as out are low. In other words, we aim at excluding the papers for which we can do so cheaply and confidently, and leave the rest to the experts (authors).

At an extreme we would like each run to be composed of one vote on one paper for one criterion (hence the name "short run"). Every time we get a vote we learn something new, and we can use this knowledge to optimize the next vote we ask. In practice a run cannot ask for one vote if we use the basic setup of typical crowdsourcing engines (it would not make sense to take time out of a person to explain a task and a criterion and stop after one vote)⁴.

In the following we introduce SM (see Algorithm 2) by first presenting the intuition behind each step and then showing the related math.

We begin at iteration 0 with an empty set of classified items, both in and out: $CI_{in}^0 \cup CI_{out}^0 = \emptyset$. We assume that authors set thresholds for false inclusion and exclusions, that is, values $\overline{P_{out}}$ and $\overline{P_{in}}$ so that we classify a paper i as out if $P(i \in OUT) \geq \overline{P_{out}}$, and analogously for $P(i \in IN)$. Notice therefore that in SM the authors set the desired precision (as we will see, possibly at the expense of price and recall, but precision is typically non negotiable in SLR as false exclusions are costly).

Baseline estimation. We perform a small baseline run as in the previous approach, to estimate power $\hat{\theta}_c^0$ and difficulty (accuracy) $\hat{\alpha}_c^0$ for each criterion (Algorithm 2, step 2). Experiments have shown us that a baseline of 50 items is often sufficient as an initial estimate (as discussed in the following section), considering also that we revise the estimates as we proceed.

Exclusion probability estimation. Here we begin the iterations. Before each run of crowdsourcing we try to identify, for each item, and given the votes V_i obtained so far for each paper i , which criterion is more likely to *efficiently* filter a paper. In other words, we identify for each criterion c the minimal number $N_{i,c}^{min}$ of successive out votes we need so that *if* we add $N_{i,c}^{min}$ to V_i (resulting in a "imaginary" set of votes V_i') *then* $P(i \in OUT|V_i') > \overline{P_{out}}$, and therefore we exclude the paper and stop working on it. Intuitively, for each item we want to select criteria that have a low $N_{i,c}^{min}$ (low number of votes and therefore low cost) and a high probability $P(N_{i,c}^{min})$ of getting those out votes.

Notice that every vote on (paper i , criterion c) we get will move $P(i \in OUT)$ closer or further away from the threshold $\overline{P_{out}}$. This will change our N^{min} and possibly the selected criterion for the next round. The probability of getting an out vote for (i, c) also changes, and it does so more strongly when the accuracy for that criterion is higher.

More formally we proceed as follows. If we denote with k the number of iterations run thus far, and with $V_{i,c}^k$ the votes obtained in the first k runs, then by applying Bayesian rule we have:

⁴With ad hoc implementations, either stand-alone or on top of commercial engines, and with fast estimation it might be possible to achieve one-vote runs though the key optimization here lies in the personalized strategy: the most important aspect is not so much asking at most one vote in each run, but asking one vote per paper

Algorithm 2: SM-Runs Algorithm

Input: $I, C, lr, \overline{P_{out}}, \overline{P_{in}}$
Output: $CI = \{CI_{in}, CI_{out}\}$

- (1) $CI \leftarrow \{\}, UI \leftarrow I, k \leftarrow 0$
- (2) **# Baseline iteration (Same as Algorithm 1 baseline)**
- (3) $\rightarrow CI, V^0, \hat{\theta}^0, \hat{\alpha}^0$
- (4) **foreach** $i \in UI$: $P^0(i \in IN_c / V_{i,c}^0) \leftarrow (1 - \hat{\theta}_c^0)$
- (5) **#SM-Runs iterations**
- (6) **while** $UI \neq \emptyset$
- (7) $k \leftarrow k + 1$
- (8) **foreach** $i \in UI$
- (9) $c(i) \leftarrow \arg \max_{c \in C} \frac{P(V_{i,c}^{k+1,k+n} = OUT)}{N_{i,c}^{min}}$
- (10) **check_stop_condition_on_i**
- (11) $I^k \leftarrow N$ items with highest $p(i)$
- (12) **foreach** $i \in I^k$
- (13) $v_{i,c}^k \leftarrow$ collect a vote for c on i
- (14) $V_i^k \leftarrow V_i^{k-1} \cup v_{i,c}^k$
- (15) $P^k(i \in IN / V_i^k) \leftarrow \prod_{c \in C} P(i \in IN_c / V_{i,c}^k)$
- (16) $P^k(i \in OUT | V_i^k) \leftarrow 1 - P(i \in IN / V_i^k)$
- (17) **if** $P(i \in IN | V_i^k) > \overline{P_{in}}$
- (18) $CI_{in} \leftarrow CI_{in} \cup \{i\}$
- (19) $UI \leftarrow UI - \{i\}$
- (20) **if** $P(i \in OUT | V_i^k) > \overline{P_{out}}$
- (21) $CI_{out} \leftarrow CI_{out} \cup \{i\}$
- (22) $UI \leftarrow UI - \{i\}$
- (23) **update power as per algorithm 1, step 7**
- (24) $CI^{diff_items} \leftarrow$ tag UI as "IN items"
- (25) $CI \leftarrow CI \cup CI^{diff_items}$
- (26) **return** CI

$$P^k(i \in IN_c | V_{i,c}^k) = \frac{P^k(V_{i,c}^k | i \in IN_c) * (1 - \hat{\theta}_c^{k-1})}{P^k(V_{i,c}^k)} \quad (4)$$

In the formula, after the first run ($k=1$), the term $\hat{\theta}_c^{k-1}$ is the proportion of papers to which criteria c applies, as computed after the baseline. $\hat{\theta}_c$ is then updated after each run.

The two P^k factors on the right side of Equation 4 can be determined as follows, where $J_{i,in}^c$ denotes the number of items i labeled as *in* for criterion c :

$$P^k(V_{i,c}^k | i \in IN_c) = \left(\frac{J_{i,in}^c}{J_{i,in}^c} \right) (\bar{\alpha}_c)^{J_{i,in}^c} * (1 - \bar{\alpha}_c)^{J_{i,out}^c} \quad (5)$$

and

$$P^k(V_{i,c}^k) = P^k(V_{i,c}^k | i \in IN_c) * (1 - \hat{\theta}_c^{k-1}) + P^k(V_{i,c}^k | i \in OUT_c) * \hat{\theta}_c^{k-1} \quad (6)$$

Now that we know how to compute $P^k(i \in IN_c | V_{i,c}^k)$ and therefore $P(i \in OUT | V_{i,c}^k)$ from Equation 2, we can compute how the exclusion probability changes as we add $n=1,2,\dots$ out votes to $V_{i,c}^k$

obtaining a set we denote as $V_{i,c}^{k \leftarrow n}$ and stop when n is such that $P(i \in OUT | V_{i,c}^{k \leftarrow n}) > \overline{P_{out}}$.

To assess the probability of getting $N_{i,c}^{min}$ out votes for criteria c on item i we proceed by first computing the probability that the next vote is out, as follows⁵ (all probabilities are conditional to the votes obtained thus far $V_{i,c}^k$):

$$P(v_{i,c}^{k+1} = OUT) = \alpha_c * (1 - P_i^k(I \in IN_c)) + (1 - \alpha_c) P_i^k(I \in IN_c).$$

We then iterate over this formula for getting the probability for the next out votes, remembering that $P_i^k(I \in IN_c)$ will have changed due to the additional out vote.

Ranking. We rank criteria for each item by weighing cost ($N_{i,c}^{min}$) and probability of success (probability $P(V_{i,c}^{k+1,k+n} = OUT)$ of getting $N_{i,c}^{min}$ consecutive out votes). We define the value of applying a criterion as the price we have to pay for unit of probability of classifying the item as out in the next $N_{i,c}^{min}$ votes, that is: $Value_{i,c} = P(V_{i,c}^{k+1,k+n} = OUT) / N_{i,c}^{min}$. We then borrow ideas from predicate ranking optimization in query processing [12] that essentially ranks based on selectivity/cost (although here we do so per item and assess it at each iteration). Applying the same logic we look for each paper for the criterion with maximum value: $Value_i = \max_{c \in C} P(V_{i,c}^{k+1,k+n} = OUT) / N_{i,c}^{min}$

In developing SM we explored alternative approaches: a main one we explored involves estimating how $P(i \in OUT)$ is likely to change if we ask for one vote on c , as an attempt to drive our choice for which vote to ask next. With relatively simple math, we can estimate the probability of the next vote being in or out, and the impact that this has on $P(i \in OUT)$, and we can select the criteria that leads us closer to the threshold. This initial choice however has an undesired behavior: if there is a low accuracy, high power criterion, it leads us to choosing this criterion. However, the low accuracy means we only take little steps towards our threshold, making the walk long and expensive. Instead, we choose criteria that can provide large variations towards the out threshold.

Stopping. As we iterate, we can see that $Value_i$ may be so low (for example, if we get conflicting votes) that it becomes ineffective to poll the crowd for that item. We can therefore stop working on papers for which $Value_i$ is lower than a threshold based on authors' preferences (Notice that we disregard the money already spent on a paper, as that is a *sunk cost* [1]). The reasonable threshold here depends on the cost ratio cr of the crowd cost for a single vote on one paper and criterion (PPL from Formula 1) divided by the author classification cost. The lower the cost ratio, the more convenient it is to insist with the crowd. For typical cost ratios, considering classification costs as estimated in the literature (see, e.g., [28]) of around 2\$ per abstract (for the US, in the medical field and including overhead), a good empirically set threshold is 100. we do not discuss this threshold further here but refer the interested reader to <http://jointreserch.net> for details.

Crowdsourcing iteration. Ranking determines the priority for the next batch of votes. The batch size is the minimal size that can practically be achieved while ensuring each worker gets value for the time they spend learning and doing the task. In practice, it

⁵To simplify the presentation here we take a single value for criteria accuracy as opposed to a confusion matrix.

Figure 2: Behavior of algorithms. Charts are simulated with 1000 papers, four criteria of power = $[c_1 = 0.14, c_2 = 0.14, c_3 = 0.28, c_4 = 0.42]$, $N_t = [2, 3, \dots, 10]$, $l_r = 5$. Workers are assumed to be cheaters with probability 0.3, and the rest has uniform accuracy in (0.5-1). Accuracy on OUT papers are 10% higher, as seen in experiments. See text for description.

rarely makes sense to offer batches of less than 10 items as they are less attractive. We return to the crowd to ask one more vote for each paper in the batch, determine the probability of exclusion as discussed above and classify paper as out if $P(i \in OUT) > \overline{P}_{out}$. If there are no more paper left to classify we stop, else we iterate.

We next analyze the results of the algorithm and discuss its properties, also in light of crowdsourcing experiments.

5 ANALYSIS AND EXPERIMENTS

Simulations. We first show the behaviors of algorithms via simulations. The strategies presented here have a number of parameters and the behavior varies in interesting ways as we change these parameters. Here we limit to point out some aspects we found particularly interesting and provide an in-depth analysis online for the interested reader, along with the code to replicate both simulation and analysis of experiments⁶. We also remind that authors do not set or estimate any parameter: they simply need to state their loss function and the preference for given loss vs price points when there is no Pareto-optimal value.

Figure 2 shows the result of a simulation run with 1000 papers and parameters as described in the caption. It plots the loss vs price curve for the SM strategy for the same scenario discussed for the other algorithms (The SM variant adopted here has a 1000 papers run, assumes a stopping threshold of 100, and shows an average of 50 simulations). $\overline{P}_{out}$ is 0.99. Figure 2(a) and (b) show that SM can achieve the same loss and precision for a fraction of the cost (both could improve by changing $\overline{P}_{out}$, though increasing the price)⁷. Notice that price and loss both decrease (at least initially) as we increase the number of tests N_t , which is our “knob” to increase accuracy (and cost) of workers. This is because SM detects the increased accuracy and adapts to it by asking for less votes for the same loss and precision. Figure 2c shows the ROC curve where we can see that SM has a greater area for a much smaller price. Charts are analogous in terms of shapes and trends for other values of θ, J (and for N_t as well for the ROC curve) so we do not show them. Figure 2d shows again loss vs price but this time assuming the presence of a very difficult criterion (accuracy of 0.55), and shows the robustness of SM on the loss (even with the very conservative loss ratio of 5 we used here).

⁶<http://jointresearch.net>

⁷We omit plotting the std bars as they would make the chart unreadable

TASK: In the following we ask you to classify an "abstract" (a short summary of a scientific paper). We ask you to state if the paper describes an **intervention** study. An intervention occurs when we perform an experiment with a set of subjects (persons) and then evaluate results. It is different from a survey, where we simply ask people questions without subjecting them to an experiment.

Example of paper that fits the criterion (you should select YES as answer to this question):

[abstract of an example paper from gold dataset]

Example of paper that DOES NOT fit the criterion (you should select NO as answer to this question, because this paper does not describe an experiment, trial, or intervention):

[abstract of an example paper from gold dataset]

If it is not clear from the text how to classify this paper, just select the option "not clear from the text."

We only accept tasks with a minimum of 5 answers and a minimum of 75% correct answers.

Description of the paper (ID \${paper_id}):

\$(abstract)

Does the paper describe an "intervention"?

☐ Yes

☐ No

☐ Not clear from the text

Figure 3: Classification task for SLR

Figure 2e and 2f show the impact of estimation error for accuracy on precision and recall respectively. Notice that if we underestimate (orange line) we achieve higher precision (we are more conservative). For recall, if we underestimate accuracy and accuracy is low (N_t is low) then we get very low recall: we give up rather early, leaving papers to authors to classify. As accuracy increases, the differences smooth out and are within the variance. The charts for power estimation error have essentially the same shape and are not shown.

Baseline runs and numbers of labels per paper affect the estimation errors. The issue is not so much the number of papers in the baseline: 40-50 papers suffice to estimate power within a 5-7% margin of error (consider the problem similar to estimating the fairness of a coin modeled as a Beta distribution, 50 tosses would give a reasonable estimate). Furthermore, estimates are re-assessed as we go. The key here is rather to enable a good accuracy estimation, and experiments have shown that with less than 3 votes per paper the estimation error grows above 10% and with low accuracy criteria this can generate very low recall (as we may believe accuracy to be at 0.5). Experiments with variations of the stopping threshold also produced limited effect. When going from 100 to 150, the recall increased by approximately 0.04%, always keeping the precision threshold at 0.99. The price difference is also negligible.

Experiments. Between January and September 2017 we performed a set of studies and experiments on two commercial crowdsourcing engines (CrowdFlower and Mechanical Turk). We ran a total of 20 experiments with different settings, asking workers to label a total of 174 papers with two to four exclusion criteria (a total of 514 classification decisions) taken from two systematic reviews, one done by us in an interdisciplinary area (computer science and

social sciences) reviewing technology for fighting loneliness (reference omitted for double blind), and the other in medicine [38] having more complex exclusion criteria. We collected votes by over 3200 respondents. These initial studies helped us to understand the nature of the problem, estimate crowd accuracies, get a feeling for latencies and costs, and also refine task design which, although orthogonal to our goals here, is important for getting good results [42]. In the following we focus on the experiments to assess the validity of SM with respect to other algorithms and baselines.

Setup. To this end we classified 374 papers on AMT by posting tasks that were asking crowd workers to classify many papers based on one criteria. We requested workers with an HIT approval rate of 70% or more. The task starts by explaining the criteria to workers, providing a positive and a negative example, and then asking to label the papers as in, out, or unclear from the text (Figure 3). Adding the latter option was a result of previous experiments where many workers complained that this option was missing and were unsure about what to answer to qualify for payment. We also informed workers that they would need to get 75% of the answers as correct in order to be paid. The examples are taken from the "gold" dataset, which are the papers classified by researchers in our team. Each worker saw the same example papers in the instructions for the same criteria. The three criteria we tested involved assessing whether the paper included an *intervention* (as shown in Figure 3), whether it described studies on adults *65 and older*, and whether it involved the use of *technology*.

The task proceeded by criteria, not by paper: we showed instruction for a criterion, and then asked for classification on that criterion. We chose this option as we noticed in initial experiments that explaining and understanding criteria takes effort, and teaching workers the subtleties of several criteria at the same time may lead to increased effort and reduced accuracy. Workers could classify the papers until they wished to do so. We repeated the process for the three criteria getting at least 5 votes per paper (we collected up to 15 votes per paper for the intervention criterion show in the figure as it had low accuracy and we wanted to analyze it more deeply). Before running the task we did some pre-runs to assess the proper pay. In terms of costs, we experimented different payments, always making sure that we stay well over 8USD per hour based on estimated completion times, which results in approximately 10cents per vote. We did not screen workers with test questions, though the experiments gave us a dataset over which we can now use to "simulate" the effect of filtering out workers that did not get 100% accuracy on the first N_t questions. The dataset is publicly available online at jointresearch.net.

Results. All tasks were completed in a few hours, and we assume that if we had more papers they would have been classified with sublinear increment in time. As expected (see table below), and consistently with the literature [28], power and workers' accuracy vary significantly by criteria. Use of technology has high power as words related to technology are very common and it is hard for keyword-based queries to filter for the specific use of technology we look for.

Criteria:	<i>intervention</i>	<i>use of tech</i>	<i>65 and older</i>
Power	0.24	0.61	0.05
Accuracy	0.60	0.77	0.75

Figure 4: Behavior of SM with data from experiments. Runs of 1000 papers, accuracy and power as described in the text, the other parameters are unchanged.

The presence of a criteria (intervention) with rather low accuracy underscores the importance of an adaptive approach where we focus on high-accuracy criteria and leave the leftover papers to the authors.

Another interesting finding was that nearly half of the papers erroneously classified by the crowd were either errors in the gold data (our error) or they were cases where after reading in detail the abstract we were unsure ourselves. This prompted us to study a bit deeper average agreement among expert raters. Mateen et al report on an experiment that measured agreement on around 96% of the papers [26]. An analysis of SLRs conducted by our team reported 92% agreement among two raters and, in addition, 2% of cases where one rater was unsure. This indicates that the P_{out} precision threshold we picked of 0.99 is in line or even exceeding current standards. We also observed that the mere act of having the need to explain criteria to others (that also demand fairness in job acceptance) forced us to be very precise and indeed, looking back at our own classification, we found errors also due to a certain imprecision in the initial definitions.

Using the experimental data to fuel simulations did not bring significant changes to the charts already discussed, although there are interesting differences and we focus on these in the following, especially to underline the limits of SM. One interesting aspect is that actual data do not precisely and consistently fit the model: workers accuracy cannot always be modeled as i.i.d. variables, and the margin of error in predicting future accuracy from past is rather high, even if we vary the testing patterns. This is not entirely surprising as some workers may improve as they proceed with the task while others may get sloppy or tired, and indeed optimal testing to cater for these issues is an active area of research [2].

Figure 4 shows the results of such experiments-fueled simulations, assuming no additional tests filtering. Figure 4a shows the usual loss vs price chart, and the results are fairly consistent with the simulations. Figures 4b and c break this down by precision and recall. The latter is particularly interesting as the recall for baseline is somewhat comparable to M-runs, or at least it makes for a non obvious choice for lower level of J . This is probably due to the relatively low accuracy we have in these (relatively untested) set of workers.

Finally we observe that in terms of overall cost, even at 10cents per vote we remain well below author cost for the paper we screen out (from 20 to 40%).

6 CONCLUSION

The SM algorithm seems to have the potential to outperform baselines for finite pool classification problems, and especially for SLR. We also confirm initial findings that crowdsourcing is feasible for paper screening in SLR. We have also explored extensions of this approach to general classification problems, including problems combining crowd and machine classification [20]. The work still has many limitations, especially that of improving the estimation of accuracies and extending the model to cover the case where workers' accuracy is very "noisy".

Acknowledgements. This project has received funding from the EU Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690962.

REFERENCES

- [1] H. Arkes and C. Blumer. 1985. The psychology of sunk cost. *Organizational Behavior and Human Decision Processes* 35, 1 (1985).
- [2] Jonathan Bragg, Mausam, and Daniel S. Weld. 2016. Optimal Testing for Crowd Workers. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS '16)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 966–974. <http://dl.acm.org/citation.cfm?id=2937029.2937066>
- [3] Andrew W. Brown and David B. Allison. 2014. Using Crowdsourcing to Evaluate Published Scientific Literature: Methods and Example. *Plos One* 9, 7 (2014).
- [4] Perrine Créquit, Ludovic Trinquart, Amélie Yavchitz, and Philippe Ravaud. 2016. Wasted research when systematic reviews fail to provide a complete and up-to-date evidence synthesis: the example of lung cancer. *BMC Medicine* 14, 1 (2016), 8. <https://doi.org/10.1186/s12916-016-0555-0>
- [5] A. P. Dawid and A. M. Skene. 1979. Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm. *Journal of the Royal Statistical Society. Series C Applied Statistics* 28, 1 (1979).
- [6] A. P. Dempster, N. M. Laird, and D. B. Rubin. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1 (1977).
- [7] Xin Luna Dong, Laure Berti-Equille, and Divesh Srivastava. 2013. Data Fusion : Resolving Conflicts from Multiple Sources. In *Proc of WAIM2013*. Springer. <https://doi.org/10.1007/978-3-642-36257-6>
- [8] Carsten Eickhoff and Arjen P de Vries. 2013. Increasing cheat robustness of crowdsourcing tasks. *Information retrieval* 16, 2 (2013), 121–137.
- [9] M.J. Grant and A. Booth. 2009. A typology of reviews: an analysis of 14 review types and associated methodologies. *Health Info Libr J* 26, 2 (2009), 91–108.

- [10] JPT Higgins and S. Green. 2011. *Cochrane Handbook for Systematic Reviews of Interventions Version 5.1.0*. The Cochrane Collaboration. Available from www.handbook.cochrane.org.
- [11] A. Haidich. 2010. Meta-analysis in medical research. *Hippokratia* 14, Suppl 1 (2010), 29–37.
- [12] Joseph M. Hellerstein and Michael Stonebraker. 1993. Predicate Migration: Optimizing Queries with Expensive Predicates. In *Proceedings of ACM SIGMOD*. ACM.
- [13] Lorna K Henderson, Jonathan C Craig, Narelle S Willis, David Tovey, and Angela C Webster. 2010. How to write a Cochrane systematic review. *Nephrology* 15, 6 (2010), 617–624.
- [14] Matthias Hirth, Tobias Hofffeld, and Phuoc Tran-Gia. 2011. Cost-optimal validation mechanisms and cheat-detection for crowdsourcing platforms. In *Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS), 2011 Fifth International Conference on*. IEEE, 316–321.
- [15] Matthias Hirth, Tobias Hofffeld, and Phuoc Tran-Gia. 2013. Analyzing costs and accuracy of validation mechanisms for crowdsourcing platforms. *Mathematical and Computer Modelling* 57, 11 (2013), 2918–2932.
- [16] Yuan Jin, Mark Carman, Dongwoo Kim, and Lexing Xie. 2017. Leveraging Side Information to Improve Label Quality Control in Crowd-Sourcing. In *Procs of Hcomp2017*. AAAI.
- [17] David R Karger, Sewoong Oh, and Devavrat Shah. 2011. Budget-optimal crowdsourcing using low-rank matrix approximations. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. IEEE, 284–291.
- [18] David R Karger, Sewoong Oh, and Devavrat Shah. 2011. Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems*. 1953–1961.
- [19] Khalid S Khan, Regina Kunz, Jos Kleijnen, and Gerd Antes. 2003. Five steps to conducting a systematic review. *Journal of the Royal Society of Medicine* 96, 3 (2003), 118–121.
- [20] Evgeny Krivosheev, Boualem Benatallah Bahareh Harandizadeh, and Fabio Casati. 2018. Crowd-based Multi-predicate Screening of Papers in Literature Reviews (poster). In *Proceedings of WWW2018*. International World Wide Web Conferences Steering Committee.
- [21] Evgeny Krivosheev, Valentina Caforio, Boualem Benatallah, and Fabio Casati. 2017. Crowdsourcing Paper Screening in Systematic Literature Reviews. In *Procs of Hcomp2017*. AAAI.
- [22] Edith Law, Krzysztof Z. Gajos, Andrea Wiggins, Mary L. Gray, and Alex Williams. 2017. Crowdsourcing As a Tool for Research: Implications of Uncertainty. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 1544–1561. <https://doi.org/10.1145/2998181.2998197>
- [23] Hongwei Li, Bin Yu, and Dengyong Zhou. 2013. Error Rate Analysis of Labeling by Crowdsourcing. In *Procs of ICML2013*.
- [24] Chao Liu and Yi Min Wang. 2012. TrueLabel + Confusions: A Spectrum of Probabilistic Models in Analyzing Multiple Ratings. In *Procs of ICML2012*. ICML.
- [25] Qiang Liu, Alexander T Ihler, and Mark Steyvers. 2013. Scoring workers in crowdsourcing: How many control questions are enough?. In *Advances in Neural Information Processing Systems*. 1914–1922.
- [26] FJ Mateen, J Oh, AI Tergas, NH Bhayani, and BB Kamdar. 2013. Titles versus titles and abstracts for initial screening of articles for systematic reviews. *Clin Epidemiol.* 5, 1 (2013).
- [27] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G Altman, Prisma Group, et al. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *PLoS med* 6, 7 (2009), e1000097.
- [28] Michael L. Mortensen, Gaelen P. Adam, Thomas A. Trikalinos, Tim Kraska, and Byron C. Wallace. 2016. An exploration of crowdsourcing citation screening for systematic reviews. *Research Synthesis Methods* (2016). RSM-02-2016-0006.R4.
- [29] Lauren Ng, Veronica Pitt, Kit Huckvale, Ornella Clavisi, Tari Turner, Russell Gruen, and Julian H Elliott. 2014. Title and Abstract Screening and Evaluation in Systematic Reviews (TASER): a pilot randomised controlled trial of title and abstract screening by medical students. *Systematic reviews* 3, 1 (2014), 121. <https://doi.org/10.1186/2046-4053-3-121>
- [30] An T Nguyen, Byron C Wallace, and Matthew Lease. 2015. Combining Crowd and Expert Labels using Decision Theoretic Active Learning. *Proceedings of the 3rd AAAI Conference on Human Computation (HCOMP)* (2015), 120–129.
- [31] Jungseul Ok, Sewoong Oh, Jinwoo Shin, and Yung Yi. 2016. Optimality of Belief Propagation for Crowdsourced Classification. In *Procs of ICML2016*.
- [32] Aditya Parameswaran, Hector Garcia-Molina, Hyunjung Park, Neoklis Polyzotis, Aditya Ramesh, and Jennifer Widom. 2012. CrowdScreen: Algorithms for Filtering Data with Humans. In *Proceedings of ACM SIGMOD*. ACM.
- [33] Margaret Sampson, Kaveh G Shojania, Chantelle Garrity, Tanya Horsley, Mary Ocampo, and David Moher. 2008. Systematic reviews can be produced and published faster. *Journal of clinical epidemiology* 61, 6 (2008), 531–536.
- [34] Padhraic Smyth, Usama Fayyad, Michael Burl, Pietro Perona, and Pierre Baldi. 1995. Inferring ground truth from subjective labelling of venus images. *Advances in neural information processing systems* 7 (1995), 1085–1092.
- [35] Earl Steinberg, Sheldon Greenfield, Dianne Miller Wolman, Michelle Mancher, Robin Graham, et al. 2011. *Clinical practice guidelines we can trust*. National Academies Press.
- [36] Yalin Sun, Pengxiang Cheng, Shengwei Wang, Hao Lyu, Matthew Lease, Iain Marshall, and Byron C. Wallace. 2016. Crowdsourcing Information Extraction for Biomedical Systematic Reviews. In *4th AAAI Conference on Human Computation and Crowdsourcing (HCOMP): Works-in-Progress Track*. <http://arxiv.org/abs/1609.01017>. 3 pages. arXiv:1609.01017.
- [37] Yemisi Takwoingi, Sally Hopewell, David Tovey, and Alex J Sutton. 2013. A multicomponent decision tool for prioritising the updating of systematic reviews. *Bmj* 7191, December (2013), 1–8. <https://doi.org/10.1136/bmj.f7191>
- [38] Nicola Veronese, Silvia Facchini, Brendon Stubbs, Claudio Luchini, Marco Solmi, Enzo Manzato, Giuseppe Sergi, Stefania Maggi, Theodore Cosco, and Luigi Fontana. 2017. Weight loss is associated with improvements in cognitive function among overweight and obese people: A systematic review and meta-analysis. *Neuroscience & Biobehavioral Reviews* 72 (2017), 87–94.
- [39] Byron C Wallace, Anna Noel-Storr, Iain J Marshall, Aaron M Cohen, Neil R Smalheiser, and James Thomas. 2017. Identifying reports of randomized controlled trials (RCTs) via a hybrid machine learning and crowdsourcing approach. *J Am Med Inform Assoc* (2017).
- [40] Michael Weiss. 2016. Crowdsourcing literature reviews in new domains. *Technology Innovation Management Review* 6, 2 (2016), 5–14.
- [41] Jacob Whitehill, Ting-fan Wu, Jacob Bergsma, Javier R Movellan, and Paul L Ruvolo. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. (2009), 2035–2043.
- [42] Jie Yang, Judith Redi, Gianluca DeMartini, and Alessandro Bozzon. 2016. Modeling Task Complexity in Crowdsourcing. In *Proceedings of The Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2016)*. AAAI, 249–258.
- [43] D. Zhou, J. Platt, S. Basu, and Y. Mao. 2012. Learning from the wisdom of crowds by minimax entropy. In *Procs of NIPS 2012*.