

Crowd Access Path Optimization: Diversity Matters

Besmira Nushi¹ Adish Singla¹ Anja Gruenheid¹ Erfan Zamanian²
 Andreas Krause¹ Donald Kossmann^{1,3}

¹ETH Zurich, Department of Computer Science, Switzerland

²Brown University, Providence, USA

³Microsoft Research, Redmond, WA, USA

Abstract

Quality assurance is one of the most important challenges in crowdsourcing. Assigning tasks to several workers to increase quality through redundant answers can be expensive if asking homogeneous sources. This limitation has been overlooked by current crowdsourcing platforms resulting therefore in costly solutions. In order to achieve desirable cost-quality tradeoffs it is essential to apply efficient crowd access optimization techniques. Our work argues that optimization needs to be aware of diversity and correlation of information within groups of individuals so that crowdsourcing redundancy can be adequately planned beforehand. Based on this intuitive idea, we introduce the Access Path Model (APM), a novel crowd model that leverages the notion of access paths as an alternative way of retrieving information. APM aggregates answers ensuring high quality and meaningful confidence. Moreover, we devise a greedy optimization algorithm for this model that finds a provably good approximate plan to access the crowd. We evaluate our approach on three crowdsourced datasets that illustrate various aspects of the problem. Our results show that the Access Path Model combined with greedy optimization is cost-efficient and practical to overcome common difficulties in large-scale crowdsourcing like data sparsity and anonymity.

Introduction

Crowdsourcing has attracted the interest of many research communities such as database systems, machine learning, and human computer interaction because it allows humans to collaboratively solve problems that are difficult to handle with machines only. Two crucial challenges in crowdsourcing independent of the field of application are (i) *quality assurance* and (ii) *crowd access optimization*. Quality assurance provides strategies that proactively plan and ensure the quality of algorithms run on top of crowdsourced data. Crowd access optimization then supports quality assurance by carefully selecting from a large pool the crowd members to ask under limited budget or quality constraints. In current crowdsourcing platforms, redundancy (*i.e.* assigning the same task to multiple workers) is the most common and straightforward way to guarantee quality (Karger, Oh, and Shah 2011).

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Access Path	Error rate	Cost
Pediatricians	10%	\$20
Logopedists	15%	\$15
Other parents	25%	\$10

Table 1: Access path configuration for Example 1

Simple as it is, redundancy can be expensive if used without any target-oriented approach, especially if the errors of workers show dependencies or are correlated. Asking people whose answers are expected to converge to the same opinion is neither efficient nor insightful. For example, in a sentiment analysis task, one would prefer to consider opinions from different non-related groups of interests before forming a decision. This is the basis of the diversity principle introduced by (Surowiecki 2005). The principle states that the best answers are achieved from discussion and contradiction rather than agreement and consensus.

In this work, we incorporate the diversity principle in a novel crowd model, named **Access Path Model** (APM), which seamlessly tackles quality assurance and crowd access optimization. It explores crowd diversity not on the individual worker level but on the common dependencies of workers while performing a task. In this context, an *access path* is a way of retrieving a piece of information from the crowd. The configuration of access paths can be based on various criteria depending on the task: (i) workers’ demographics (*e.g.* profession, group of interest, age) (ii) the source of information or the tool that is used to find the answer (*e.g.* phone call vs. web page, Bing vs. Google) (iii) task design (*e.g.* time of completion, user interface) (iv) task decomposition (*e.g.* part of the answers, features).

Example 1. *Peter and Aanya natively speak two different languages which they would like to teach to their young children. At the same time, they are concerned how this multilingual environment affects the learning abilities of their children. More specifically, they want to answer the question “Does raising children bilingually cause language delay?”. To resolve their problem, they can ask three different groups of people (access paths):*

Figure 1 illustrates the given situation with respect to the Access Path Model. In this example, each of the groups ap-

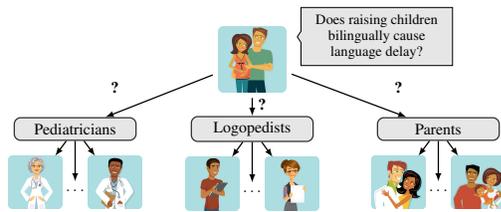


Figure 1: APM for crowdsourcing a medical question

proaches the problem from a different perspective and has different associated error rates and costs. Considering that Peter and Anya have a limited budget to spend and can ask more than one person on the same access path, they are interested in finding the optimal combination of access paths that will give them the most insightful information for their budget constraints. Throughout this paper, a combination of access paths will be referred to as an *access plan* and it defines how many different people to ask on each available access path. Our model aims at helping general requesters in crowdsourcing platforms to find optimal access plans and appropriately aggregate the collected data. Results from experiments on real-world crowdsourcing show that a pre-planned combination of diverse access paths indeed overperforms pure (*i.e.* single access path) access plans, random selection, and equal distribution of budget across access paths.

Contributions

Previous work on quality assurance and crowd access optimization focuses on two different approaches: majority-based strategies and individual models. Majority voting is oblivious to personal characteristics of crowd workers and is therefore limited in terms of optimization. Individual models instead base their decisions on the respective performance of each worker targeting those with the best accuracy (Dawid and Skene 1979; Whitehill et al. 2009). These models are useful for spam detection and pricing schemes but do not guarantee answer diversity and might fall into partial consensus traps.

As outlined in Table 2, the APM is a middle-ground solution between these two choices and offers several advantages. First, it is aware of answer diversity which is particularly important for requests without an established ground truth. Second, since it manages group-based answer correlations and dependencies, it facilitates efficient optimization of redundancy. Third, the APM is a practical model for current crowdsourcing marketplaces where due to competition the availability of a particular person is never guaranteed or authorships may be hidden for privacy reasons. Last, its predictions are mapped to meaningful confidence levels which can simplify the interpretation of results.

In summary, this work makes the following contributions:

- **Modeling the crowd for quality assurance.** We design the Access Path Model as a Bayesian Network that through the usage of latent variables is able to capture and utilize crowd diversity from a non-individual point of view. The APM can be applied even if the data is sparse and crowd workers are anonymous.

	Majority Voting	Individual Models	Access Path Model
Diversity awareness	✗	✓	✓
Cost-efficient optimization	✗	✗	✓
Sparsity Anonymity	✓	✗	✓
Meaningful confidence	✓	✗	✓

Table 2: Comparison of APM with current approaches.

- **Crowd access optimization.** We use an information-theoretic objective for crowd access optimization. We prove that our objective is submodular, allowing us to adopt efficient greedy algorithms with strong guarantees.
- **Real-world experiments.** Our extensive experiments cover three different domains: Answering medical questions, sport events prediction and bird species classification. We compare our model and optimization scheme with state of the art techniques and show that it makes robust predictions with lower cost.

Problem Statement

In this work, we identify and address two closely related problems: (1) modeling and aggregating diverse crowd answers which we call the *crowdsourced predictions problem*, and (2) optimizing the budget distribution for better quality referred to as *access path selection problem*.

Problem 1 (CROWDSOURCED PREDICTIONS). *Given a task represented by a random variable Y , and a set of answers from W workers represented by random variables X_1, \dots, X_W , the crowdsourced prediction problem is to find a high-quality prediction of the outcome of task Y by aggregating these votes.*

Quality criteria. A high-quality prediction is not only accurate but should also be linked to a meaningful confidence score which is formally defined as the likelihood of the prediction to be correct. This property simplifies the interpretation of predictions coming from a probabilistic model. For example, if a doctor wants to know whether a particular medicine can positively affect the improvement of a disease condition, providing a raw *yes/no* result answer is not sufficiently informative. Instead, it is much more useful to associate the answer with a trustable confidence score.

Requirements and challenges. To provide high quality predictions, it is essential to precisely represent the crowd. The main aspects to be represented are (i) the conditional dependence of worker answers within access paths given the task and (ii) the conditional independence of worker answers across access paths. As we will show in this paper, modeling such dependencies is also crucial for efficient optimization. Another realistic requirement concerns the support for data *sparsity* and *anonymity*. Data sparsity is common in crowdsourcing (Venanzi et al. 2014) and occurs when the number of tasks that workers solve is not sufficient to estimate their errors which can negatively affect quality. In other cases, the identity of workers may not be available for privacy reasons.

Problem 2 (ACCESS PATH SELECTION). *Given a task represented by a random variable Y , that can be solved by the crowd following N different access paths denoted with the random variables Z_1, \dots, Z_N , using a maximum budget B , the access path selection problem is to find the best possible access plan S_{best} that leads to a high-quality prediction of the outcome of task Y .*

An access plan defines how many different people are chosen to complete the task from each access path. In Example 1, we will ask one pediatrician, two logopedists and three different parents if the access plan is $S = [1, 2, 3]$. Each access plan is associated with a cost $c(S)$ and quality $q(S)$. For example, $c(S) = \sum_{i=1}^3 c_i \cdot S[i] = \80 where c_i is the cost of getting one single answer through access path Z_i . In these terms, the access path selection problem can be generally formulated as:

$$S_{best} = \arg \max_{S \in \mathcal{S}} q(S) \text{ s.t. } \sum_{i=1}^N c_i \cdot S[i] \leq B \quad (1)$$

This knapsack maximization problem is NP-Hard even for submodular functions (Feige 1998). Hence, designing bounded and efficient approximation schemes is useful for realistic crowd access optimization.

Access Path Model

The crowd model presented in this section aims at fulfilling the requirements specified in the definition of Problem 1 (CROWDSOURCED PREDICTION) and enables our method to learn the error rates from historical data and then accordingly aggregate worker votes.

Access Path Design

Due to the variety of problems possible to crowdsource, an important step concerns the design of access paths. The access path notion is a broad concept that can accommodate various situations and may take different shapes depending on the task. Below we describe a list of viable configurations that can be easily applied in current platforms.

- **Demographic groups.** Common demographic characteristics (location, gender, age) can establish strong statistical dependencies of workers' answers (Kazai, Kamps, and Milic-Frayling 2012). Such groups are particularly diverse for problems like sentiment analysis or product evaluation and can be retrieved from crowdsourcing platforms as part of the task, worker information, or qualification tests.
- **Information sources.** For data collection and integration tasks, the data source being used to deduplicate or match records (addresses, business names *etc.*) is the primary cause of error or accuracy (Pochampally et al. 2014).
- **Task design.** In other cases, the answer of a worker may be psychologically affected by the user interface design. For instance, in crowdsourced sorting, a worker may rate the same product differently depending on the scaling system (stars, 1-10 *etc.*) or other products that are part of the same batch (Parameswaran et al. 2014).
- **Task decomposition.** Often, complicated problems are decomposed into smaller ones. Each subtask type can serve

as an access path. For instance, in the bird classification task that we study later in our experiments, workers can resolve separate features of the bird (*i.e.* color, beak shape *etc.*) rather than its category.

In these scenarios, the access path definition natively comes with the problem or the task design. However, there are scenarios where the structure is not as evident or more than one grouping is applicable. Helpful tools in this regard include graphical model structure learning based on conditional independence tests (De Campos 2006) and information-theoretic group selection (Li, Zhao, and Fuxman 2014).

Architectural implications. We envision access path design as part of the quality assurance and control module for new crowdsourcing frameworks or, in our case, as part of the query engine in a crowdsourced database (Franklin et al. 2011). In the latter context, the notion of access paths is one of the main pillars in query optimization for traditional databases (Selinger et al. 1979) where access path selection (*e.g.* sequential scan or index) has significant impact on the query response time. In addition, in a crowdsourced database the access path selection also affects the quality of query results. In such an architecture, the query optimizer is responsible for (i) determining the optimal combination of access paths as shown in the following section, and (ii) forwarding the design to the UI creation. The query executor then collects the data from the crowd and aggregates it through the probabilistic inference over the APM.

Alternative models

Before describing the structure of the Access Path Model, we first have a look at other alternative models and their behavior with respect to quality assurance. Table 3 specifies the meaning of each symbol as used throughout this paper.

Majority Vote (MV). Being the simplest of the models and also the most popular one, majority voting is able to produce fairly good results if the crowdsourcing redundancy is sufficient. Nevertheless, majority voting considers all votes as equal with respect to quality and can not be integrated with any optimization scheme other than random selection.

Naïve Bayes Individual (NBI). This model assigns individual error rates to each worker and uses them to weigh the incoming votes and form a decision (Figure 2). In cases when the ground truth is unknown, the error estimation is carried out through an EM Algorithm as proposed by (Dawid and Skene 1979). Aggregation (*i.e.* selecting the best prediction) is then performed through Bayesian inference. For example, for a set of votes x_t coming from W different workers X_1, \dots, X_W the most likely outcome among all candidate outcomes y_c is computed as prediction = $\arg \max_{y_c \in Y} p(y_c | x_t)$, whereas the joint probability of a candidate answer y_c and the votes x_t is:

$$p(y_c, x_t) = p(y) \prod_{w=1}^W p(x_{wt} | y_c) \quad (2)$$

The quality of predictions for this model highly depends on the assumption that each worker has solved a fairly sufficient number of tasks. This assumption generally does not hold for open crowdsourcing markets where stable participation of

Symbol	Description
Y	random variable of the crowdsourced task
X_w	random variable of worker w
W	number of workers
Z_i	latent random variable of access path i
X_{ij}	random variable of worker j in access path i
N	number of access paths
B	budget constraint
S	access plan
$S[i]$	no. of votes from access path i in plan S
c_i	cost of access path i
D	training dataset
$s < y, x >$	instance of task sample in a dataset
θ	parameters of the Access Path Model

Table 3: Symbol description

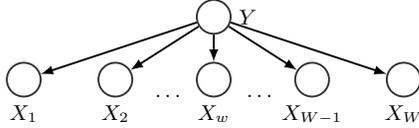


Figure 2: Naïve Bayes Individual - NBI.

workers is not guaranteed. As we show in the experimental evaluation, this is harmful not only for estimating the error rates but also for crowd access optimization because access plans might not be implementable or have a high response time. Furthermore, even in cases of fully committed workers, NBI does not provide the proper logistics to optimize the budget distribution since it does not capture the shared dependencies between the workers. Last, due to the Naïve Bayes inference which assumes conditional independence between each pair of workers, predictions of this model are generally overconfident.

Access Path based models

Access Path based models group the answers of the crowd according to the access path they originate from. We first describe a simple Naïve Bayes version of such a model and then elaborate on the final design of the APM.

Naïve Bayes for Access Paths (NBAP). For correcting the effects of non-stable participation of individual workers we first consider another alternative, similar to our original model, presented in Figure 3. The votes of the workers here are grouped according to the access path. For inference purposes then, each vote x_{ij} is weighed with the average error rate θ_i of the access path it comes from. In other words, it is assumed that all workers within the same access path share the same error rate. This simplification enables the model to support highly sparse data. Yet, due to the similarity with NBI and all Naïve Bayes classifiers, NBAP cannot make predictions with meaningful confidence especially when there exists a large number of access paths.

Access Path Model overview. Based on the analysis of previous models, we propose the Access Path Model as presented in Figure 4, which shows an instantiation for three access paths. We design the triple $\langle \text{task, access path, worker} \rangle$ as a hierarchical Bayesian Network in three layers.

Layer 1. Variable Y in the root of the model represents the

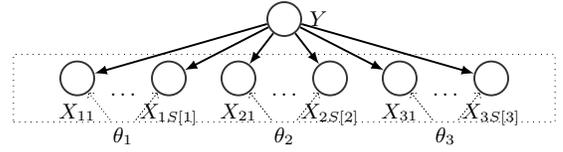


Figure 3: Naïve Bayes Model for Access Paths - NBAP.

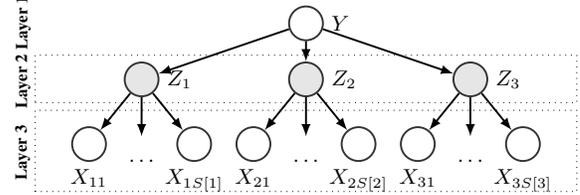


Figure 4: Bayesian Network Model for Access Paths - APM.

random variable modeling the real outcome of the task.

Layer 2. This layer contains the random variables modeling the access paths Z_1, Z_2, Z_3 . Each access path is represented as a latent variable, since its values are not observable. Due to the tree structure, every pair of access paths is conditionally independent given Y while the workers that belong to the same access path are not. The conditional independence is the key of representing diversity by implementing therefore various probabilistic channels. Their purpose is to distinguish the information that can be obtained from the workers from the one that comes from the access path.

Such enhanced expressiveness of this auxiliary layer over the previously described NBAP model avoids overconfident predictions as follows. Whenever a new prediction is made, the amount of confidence that identical answers from different workers in the same access path can bring is first blocked by the access path usage (*i.e.* the latent variable). If the number of agreeing workers within the same access path increases, confidence increases as well but not at the same rate as it happens with NBI. Additional workers contribute only with their own signal, while the access path signal has already been taken into consideration. In terms of optimization, this property of the APM makes a good motivation for combining various access paths within the same plan.

Layer 3. The lowest layer contains the random variables X modeling the votes of the workers grouped by the access path they are following. For example, X_{ij} is the j -th worker on the i -th access path. The incoming edges represent the error rates of workers conditioned by their access paths.

Parameter learning. The purpose of the training stage is to learn the parameters of the model, *i.e.* the conditional probability of each variable with respect to its parents that are graphically represented by the network edges in Figure 4. We will refer to the set of all model parameters as θ . More specifically, $\theta_{Z_i|Y}$ represents the table of conditional error probabilities for the i -th access path given the task Y , and $\theta_{X_{ij}|Z_i}$ represents the table of conditional error probabilities for the j -th worker given the i -th access path.

For a dataset D with historical data of the same type of task, the parameter learning stage finds the maximum likelihood estimate $\theta_{MLE} = \arg \max_{\theta} p(D|\theta)$. According to our model,

the joint probability of a sample s_k factorizes as:

$$p(s_k|\theta) = p(y_k|\theta) \prod_{i=1}^N \left(p(z_{ik}|y_k, \theta) \prod_{j=1}^{S_k[i]} p(x_{ijk}|z_{ik}, \theta) \right) \quad (3)$$

where $S_k[i]$ is the number of votes in access path Z_i for the sample. As the access path variables Z_i are not observable, we apply an Expectation Maximization (EM) algorithm (Dempster, Laird, and Rubin 1977) to find the best parameters. Notice that applying EM for the network model in Figure 4 will learn the parameters for each worker in the crowd. This scheme works if the set of workers involved in the task is sufficiently stable to provide enough samples for computing their error rates (*i.e.* $\theta_{X_{ij}|Z_i}$) and if the worker id is not hidden. As in many of the crowdsourcing applications (as well as in our experiments) this is not always the case, we share the parameters of all workers within an access path. This enables us to later apply on the model an optimization scheme agnostic about the identity of workers. The generalization is optional for the APM and obligatory for NBAP.

Training cost analysis. The amount of data needed to train the APM is significantly lower than what individual models require which results in a faster learning process. The reason is that the APM can benefit even from infrequent participation of individuals X_{ij} to estimate accurate error rates for access paths Z_i . Moreover, sharing the parameters of workers in the same access path reduces the number of parameters to learn from W for individual models to $2N$ for the APM which is typically several orders of magnitude lower.

Inference. After parameter learning, the model is used to infer the outcome of a task using the available votes on each access path. As in previous models, the inference step computes the likelihood of each candidate outcome $y_c \in Y$ given the votes in the test sample x_t and chooses the most likely candidate as prediction = $\arg \max_{y_c \in Y} p(y_c|x_t)$. As the test samples contain only the values for the variables X , the joint probability between the candidate outcome and the test sample is computed by marginalizing over all possible values of Z_i (Eq. 4). For a fixed cardinality of Z_i , the complexity of inferring the most likely prediction is $\mathcal{O}(NM)$.

$$p(y_c, x_t) = p(y_c) \prod_{i=1}^N \left(\sum_{z \in \{0,1\}} p(z|y_c) \prod_{j=1}^{S_t[i]} p(x_{ijt}|z) \right) \quad (4)$$

The confidence of the prediction maps to the likelihood that the prediction is accurate $p(\text{prediction}|x_t)$. Marginalization in Equation 4 is the technical step that avoids overconfidence by smoothly blocking the confidence increase when similar answers from the same access path are observed.

Crowd Access Optimization

Crowd access optimization is crucial for both paid and non-paid of crowdsourcing. While in paid platforms the goal is to acquire the best quality for the given monetary budget, in non-paid applications the necessity for optimization comes from the fact that highly redundant accesses might decrease user satisfaction and increase latency. In this section, we describe

how to estimate the quality of access plans and how to choose the plan with the best expected quality.

Information Gain as a measure of quality

The first step of crowd access optimization is estimating the quality of access plans before they are executed. One attempt might be to quantify the *accuracy* of individual access paths in isolation, and choose an objective function that prefers the selection of more accurate access paths. However, due to statistical dependencies of responses within an access path (*e.g.*, correlated errors in the workers' responses), there is diminishing returns in repeatedly selecting a single access path. To counter this effect, an alternative would be to define the quality of an access plan as a measure of *diversity* (Hui and Li 2015). For example, we might prefer to equally distribute the budget across access paths. However, some access paths may be very uninformative / inaccurate, and optimizing diversity alone will waste budget. Instead, we use the joint *information gain* $IG(Y; S)$ of the task variable Y in our model and an access plan S as a measurement of plan quality as well as an objective function for our optimization scheme. Formally, this is defined as:

$$IG(Y; S) = H(Y) - H(Y|S) \quad (5)$$

An access plan S determines how many variables X to choose from each access path Z_i . $IG(Y; S)$ measures the entropy reduction (as measure of uncertainty) of the task variable Y after an access plan S is observed. At the beginning, selecting from the most accurate access paths provides the highest uncertainty reduction. However, if better access paths are exhausted (*i.e.*, accessed relatively often), asking on less accurate ones reduces the entropy more than continuing to ask on previously explored paths. This situation reflects the way how information gain explores diversity and increases the prediction confidence if evidence is retrieved from independent channels. Based on this analysis, information gain naturally trades accuracy and diversity. While plans with high information gain do exhibit diversity, this is only a means for achieving high predictive performance.

Information gain computation. The computation of the conditional entropy $H(Y|S)$ as part of information gain in Equation 5 is a difficult problem, as full calculation requires enumerating all possible instantiations of the plan. Formally, the conditional entropy can be computed as:

$$H(Y|S) = \sum_{y \in Y, x \in X_S} p(x, y) \log \frac{p(x)}{p(x, y)} \quad (6)$$

X_S refers to all the possible assignments that votes can take according to plan S . We choose to follow the sampling approach presented in (Krause and Guestrin 2005a) which randomly generates samples satisfying the access plan according to our Bayesian Network model. The final conditional entropy will then be the average value of the conditional entropies of the generated samples. This method is known to provide absolute error guarantees for any desired level of confidence if enough samples are generated. Moreover, it runs in polynomial time if sampling and probabilistic inference can also be done in polynomial time. Both conditions are satisfied by our

model due to the tree-shaped configuration of the Bayesian Network. They also hold for the Naïve Bayes baselines as simpler tree versions of the APM.

Submodularity of information gain. Next, we derive the submodularity property of our objective function based on information gain in Equation 5. The property will then be leveraged by the greedy optimization scheme in proving constant factor approximation bounds. A submodular function is a function that satisfies the law of diminishing returns which means that the marginal gain of the function decreases while incrementally adding more elements to the input set.

Let \mathcal{V} be a finite set. A set function $F : 2^{\mathcal{V}} \rightarrow \mathbb{R}$ is submodular if $F(S \cup \{v\}) - F(S) \geq F(S' \cup \{v\}) - F(S')$ for all $S \subseteq S' \subseteq \mathcal{V}$, $v \notin S'$. For our model, this intuitively means that collecting a new vote from the crowd adds more information when few votes have been acquired rather than when many of them have already been collected. While information gain is non-decreasing and non-negative, it may not be submodular for a general Bayesian Network. Information gain can be shown to be submodular for the Naïve Bayes Model for Access Paths (NBAP) in Figure 3 by applying the results from (Krause and Guestrin 2005a). Here, we prove its submodularity property for the APM Bayesian Network shown in Figure 4. Theorem 1 formally states the result and below we describe a short sketch of the proof¹.

Theorem 1. *The objective function based on information gain in Equation 5 for the Bayesian Network Model for Access Paths (APM) is submodular.*

Sketch of Theorem 1. For proving Theorem 1, we consider a generic Bayesian Network with N access paths and M possible worker votes on each access path. To prove the submodularity of the objective function, we consider two sets (plans) $S \subset S'$ where $S' = S \cup \{v_j\}$, i.e., S' contains one additional vote from access path j compared to S . Then, we consider adding a vote v_i from access path i and we prove the diminishing return property of adding v_i to S' compared to adding to S . The proof considers two cases. When v_i and v_j belong to different access paths, i.e., $i \neq j$, the proof follows by using the property of conditional independence of votes from different access paths given Y and using the “information never hurts” principle (Cover and Thomas 2012). For the case of v_i and v_j belonging to the same access path we reduce the network to an equivalent network which contains only one access path Z_i and then use the “data processing inequality” principle (Cover and Thomas 2012). \square

This theoretical result is of generic interest for other applications and a step forward in proving the submodularity of information gain for more generic Bayesian networks.

Optimization scheme

After having determined the joint information gain as an appropriate quality measure for a plan, the crowd access optimization problem is to compute:

$$S_{best} = \arg \max_{S \in \mathcal{S}} \text{IG}(Y; S) \text{ s.t. } \sum_{i=1}^N c_i \cdot S[i] \leq B \quad (7)$$

¹Full proofs available at <http://arxiv.org/abs/1508.01951>

ALGORITHM 1. GREEDY Crowd Access Optimization

```

1 Input: budget  $B$ 
2 Output: best plan  $S_{best}$ 
3 Initialization:  $S_{best} = \emptyset$ ,  $b = 0$ 
4 while ( $\exists i$  s.t.  $b \leq c_i$ ) do
5    $U_{best} = 0$ 
6   for  $i = 1$  to  $N$  do
7      $S_{pure} = \text{PurePlan}(i)$ 
8     if  $c_i \leq B - b$  then
9        $\Delta_{IG} = \text{IG}(Y; S_{best} \cup S_{pure}) - \text{IG}(Y, S_{best})$ 
10      if  $\frac{\Delta_{IG}}{c_i} > U_{best}$  then
11         $U_{best} = \frac{\Delta_{IG}}{c_i}$ 
12         $S_{max} = S_{best} \cup S_{pure}$ 
13       $S_{best} = S_{max}$ 
14       $b = \text{cost}(S_{best})$ 
15 return  $S_{best}$ 

```

where \mathcal{S} is the set of all plans. An exhaustive search would consider $|\mathcal{S}| = \prod_{i=1}^N \frac{B}{c_i}$ plans out of which the ones that are not feasible have to be eliminated. Nevertheless, efficient approximation schemes can be constructed given that the problem is an instance of submodular function maximization under budget constraints (Krause and Guestrin 2005b; Sviridenko 2004). Based on the submodular and non-decreasing properties of information gain we devise a greedy technique in Algorithm 1 that incrementally finds a local approximation for the best plan. In each step, the algorithm evaluates the benefit-cost ratio U between the marginal information gain and cost for all feasible access paths. The marginal information gain is the improvement of information gain by adding to the current best plan one pure vote from one access path. In the worst case, when all access paths have unit cost, the computational complexity of the algorithm is $\mathcal{O}(GN^2MB)$, where G is the number of generated samples for computing information gain.

Theoretical bounds of greedy optimization. We now employ the submodularity of information gain in our Bayesian network to prove theoretical bounds of the greedy optimization scheme. For the simple case of unit cost access paths, the greedy selection in Algorithm 1 guarantees a utility of at least $(1 - 1/e)$ ($= 0.63$) times the one obtained by optimal selection denoted by OPT (Nemhauser, Wolsey, and Fisher 1978). However, the greedy selection scheme fails to provide approximation guarantees for the general setting of varying costs (Khuller, Moss, and Naor 1999).

Here, we exploit the following realistic property about the costs of the access paths and allocated budget to prove strong theoretical guarantees about our Algorithm 1. We assume that the allocated budget is large enough compared to the costs of the access paths. Formally stating, we assume that the cost of any access path c_i is bounded away from total budget B by factor γ , i.e., $c_i \leq \gamma \cdot B \forall i \in \{1, \dots, N\}$, where $\gamma \in (0, 1)$. We state the theoretical guarantees of the Algorithm 1 in Theorem 2 and below we show a brief sketch of the proof¹.

Theorem 2. The GREEDY optimization in Algorithm 1 achieves a utility of at least $\left(1 - \frac{1}{e^{(1-\gamma)}}\right)$ times that obtained by the optimal plan OPT, where $\gamma = \max_{i \in \{1, \dots, N\}} \frac{c_i}{B}$.

For instance, Algorithm 1 achieves an approximation ratio of at least 0.39 for $\gamma = 0.5$, and 0.59 for $\gamma = 0.10$.

Sketch of Theorem 2. We follow the structure of the proof from (Khuller, Moss, and Naor 1999; Sviridenko 2004). The key idea is to use the fact that the budget spent by the algorithm at the end of execution when it can not add an element to the solution is at least $(B - \max_{i \in [1, \dots, N]} c_i)$, which is lower-bounded by $B(1 - \gamma)$. This lower bound on the spent budget, along with the fact that the elements are picked greedily at every iteration leads to the desired bounds. \square

These results are of practical importance in many other applications as the assumption of non-unit but bounded costs with respect to budget often holds in realistic settings.

Experimental Evaluation

We evaluated our work on three real-world datasets. The main goal of the experiments is to validate the proposed model and the optimization technique. We compare our approach with other state of the art alternatives and results show that leveraging diversity through the Access Path Model combined with the greedy crowd access optimization technique can indeed improve the quality of predictions.

Metrics. The comparison is based on two main metrics: *accuracy* and *negative log-likelihood*. Accuracy corresponds to the percentage of correct predictions. Negative log-likelihood is computed as the sum over all test samples of the negative log-likelihood that the prediction is accurate. Hence, it measures not only the correctness of a model but also its ability to output meaningful confidence.

$$-\log\text{Likelihood} = - \sum_{s_t} \log p(\text{prediction} = y_t | x_t) \quad (8)$$

The closer a prediction is to the real outcome the lower is its negative log-likelihood. Thus, a desirable model should offer low values of negative log-likelihood.

Dataset description

All the following datasets come from real crowdsourcing tasks. For experiments with restricted budget, we repeat the learning and prediction process via random vote selection and k-fold cross-validation.

CUB-200. The dataset (Welinder et al. 2010) was built as a large-scale data collection for attribute-based classification of bird images on Amazon Mechanical Turk (AMT). Since this is a difficult task even for experts, the crowd workers are not directly asked to determine the bird category but whether a certain attribute is present in the image. Each attribute (e.g., yellow beak) brings a piece of information for the problem and we treat them as access paths. The dataset contains 5-10 answers for each of the 288 available attributes. We keep the cost of all access paths equal as there was no clear evidence of attributes that are more difficult to distinguish than others. The total number of answers is approximately 7.5×10^6 .

	Description	Forums
(1)	Answers from doctors	www.webmd.com www.medhelp.org
(2)	Answers from patients	www.patient.co.uk www.ehealthforum.com
(3)	General Q&A forum	www.quora.com www.wiki.answers.com

Table 4: Access Path Design for MedicalQA dataset.

MedicalQA. We gathered 100 medical questions and forwarded them to AMT. Workers were asked to answer the questions after reading in specific health forums categorized as in Table 4 which we then design as access paths. 255 people participated in our experiment. The origin of the answer was checked via an explanation url provided along with the answer as a sanity check. The tasks were paid equally to prevent the price of the task to affect the quality of the answers. For experimental purposes, we assign an integer cost of (3, 2, 1) based on the reasoning that in real life doctors are more expensive to ask, followed by patients and common people.

ProbabilitySports. This data is based on a crowdsourced betting competition (www.probabilitysports.com) on NFL games. The participants voted on the question: “*Is the home team going to win?*” for 250 events within a season. There are 5,930 players in the entire dataset contributing with 1,413,534 bets. We designed the access paths based on the accuracy of each player in the training set which does not reveal information about the testing set. Since the players’ accuracy in the dataset follows a normal distribution, we divide this distribution into three intervals where each interval corresponds to one access path (worse than average, average, better than average). As access paths have a decreasing error rate, we assign them an increasing cost (2, 3, 4).

Model evaluation

For evaluating the Access Path Model independently of the optimization, we first show experiments where the budget is equally distributed across access paths. The question we want to answer here is: “*How robust are the APM predictions in terms of accuracy and negative log-likelihood?*”

Experiment 1: Constrained budget. Figure 5 illustrates the effect of data sparsity on quality. We varied the budget and equally distributed it across all access paths. We do not show results from CUB-200 as the maximum number of votes per access path in this dataset is 5-10.

MedicalQA. The participation of workers in this experiment was stable, which allows for a better error estimation. Thus, as shown in Figure 5(a), for high redundancy NBI reaches comparable accuracy with the APM although the negative log-likelihood dramatically increases. For lower budget and high sparsity NBI cannot provide accurate results.

ProbabilitySports. Figure 5(b) shows that while the improvement of the APM accuracy over NBI and MV is stable, NBAP starts facing the overconfidence problem while budget increases. NBI exhibits low accuracy due to very high sparsity even for sufficient budget. Majority Vote fails to produce accurate predictions as it is agnostic to error rates.

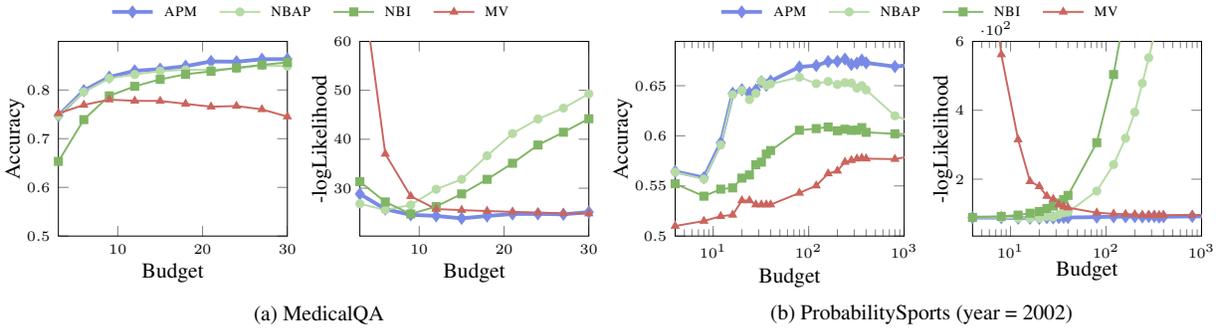


Figure 5: Accuracy and negative log-likelihood for equally distributed budget across all access paths. The negative log-likelihood of Naive Bayes models deteriorates for high budget while for the APM it stays stable. NBI is not competitive due to data sparsity.

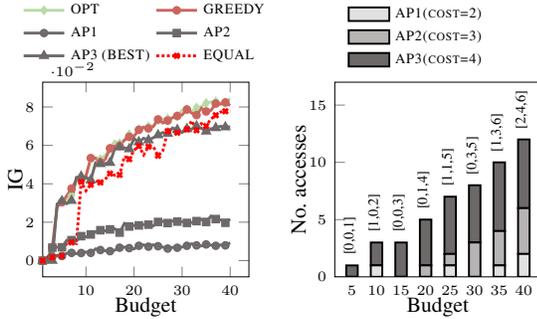


Figure 6: Information gain and budget distribution for ProbabilitySports (year=2002). As budget increases, GREEDY access plans exploit more than one access path.

Optimization scheme evaluation

In these experiments, we evaluate the efficiency of the greedy approximation scheme to choose high-quality plans. For a fair comparison, we adapted the same scheme to NBI and NBAP. We will use the following acronyms for the crowd access strategies: OPT (optimal selection), GREEDY (greedy approximation), RND (random selection), BEST (votes from the most accurate access path), and EQUAL (equal distribution of votes across access paths).

Experiment 2: Greedy approximation and diversity. The goal of this experiment is to answer the questions: “How close is the greedy approximation to the theoretical optimal solution?” and “How does information gain exploit diversity?”. Figure 6 shows the development of information gain for the optimal plan, the greedily approximated plan, the equal distribution plan, and three pure plans that take votes only from one access path. The quality of GREEDY is very close to the optimal plan. The third access path in ProbabilitySports (containing better than average users) reaches the highest information gain compared to the others. Nevertheless, its quality is saturated for higher budget which encourages the optimization scheme to select other access paths as well. Also, we notice that the EQUAL plan does not reach optimal values of information gain although it maximizes diversity. Next, we show that the quality of predictions can be further improved if diversity is instead planned by

using information gain as an objective.

Experiment 3: Crowd access optimization. This experiment combines together both the model and the optimization technique. The main question we want to answer here is: “What is the practical benefit of greedy optimization on the APM w.r.t. accuracy and negative log-likelihood?”

CUB-200. For this dataset (Figure 7(a)) where the access path design is based on attributes, the discrepancy between NBAP and the APM is high and EQUAL plans exhibit low quality as not all attributes are informative for all tasks.

ProbabilitySports. Access Path based models (APM and NBAP) outperform MV and NBI. NBI plans target concrete users in the competition. Hence, their accuracy for budget values less than 10 is low as not all targeted users voted for all events. Since access paths are designed based on the accuracy of workers, EQUAL plans do not offer a clear improvement while NBAP is advantaged in terms of accuracy by its preference to select the most accurate access paths.

Experiment 5: Diversity impact. This experiment is designed to study the impact of diversity and conditional dependence on crowd access optimization, and finally answer the question: “How does greedy optimization on the APM handle diversity?”. One form of such dependency is within access path correlation. If this correlation holds, workers agree on the same answer. We experimented by varying the shared dependency within the access path as follows: Given a certain probability p , we decide whether a vote should follow the majority vote of existing answers in the same access path. For example, for $p = 0.4$, 40% of the votes will follow the majority vote decision of the previous workers and the other 60% will be withdrawn from the real crowd votes.

Figure 8(a) shows that the overall quality drops when dependency is high but the Access Path Model is more robust to it. NBAP instead, due to overconfidence, accumulates all votes into a single access path which dramatically penalizes its quality. APM+BEST applies the APM to votes selected from the access path with the best accuracy, in our case doctors’ answers. Results show that for $p > 0.2$, it is preferable to not only select from the best access path but to distribute the budget according to the GREEDY scheme. Figure 8(b) shows results from the same experiment for $p = 0.4$ and varying budget. APM+GREEDY outperforms all other methods reaching a stable quality at $B = 30$.

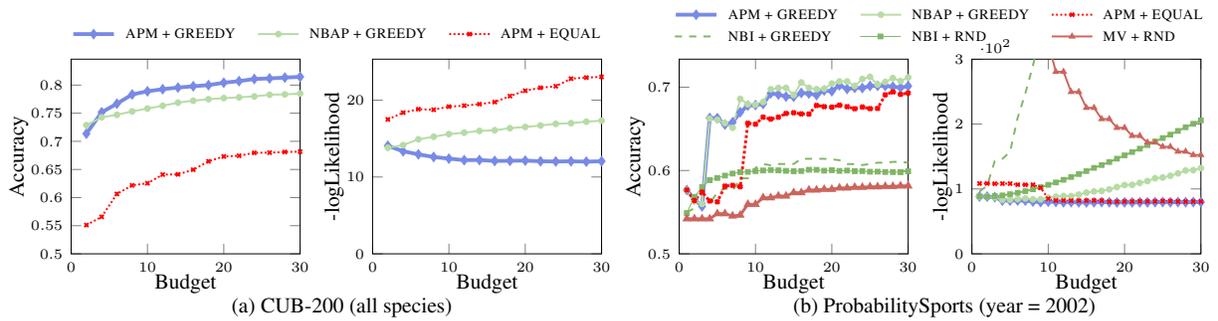


Figure 7: Crowd access optimization results for varying budget. Data sparsity and non-guaranteed votes are better handled by the APM model also for optimization purposes, leading to improved accuracy and confidence.

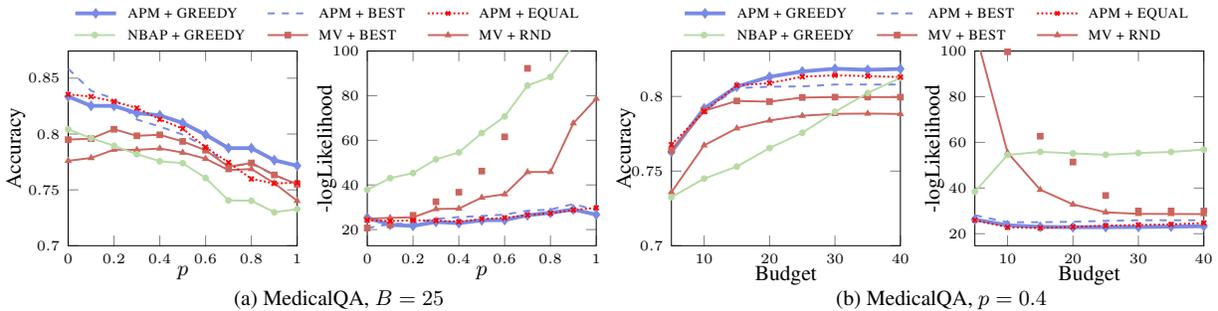


Figure 8: Diversity and dependence impact on optimization. As the common dependency of workers within access paths increases, investing the whole budget on the best access path or randomly is not efficient.

Discussion

We presented experiments based on three different and challenging crowdsourcing tasks. The chosen datasets highlight different yet all important insights on the problem. However, our approach and results are of general purpose and are not tailored to any of the datasets. The main findings are:

- In real-world crowdsourcing the unrealistic assumption of pairwise worker independence poses limitations to quality assurance and increases the cost of crowdsourced solutions based on individual and majority vote models.
- Managing and exploiting diversity with the APM ensures quality in terms of accuracy and more significantly negative log-likelihood. Crowd access optimization schemes on top of this perspective are practical and cost-efficient.
- Surprisingly, access plans that combine various access paths make better predictions than plans which spend the whole budget in a single access path.

Related Work

The reliability of crowdsourcing and relevant optimization techniques are longstanding issues for human computation platforms. The following directions are closest to our study:

Quality assurance and control. One of the central works in this field is presented by (Dawid and Skene 1979). In an experimental design with noisy observers, the authors use an Expectation Maximization algorithm (Dempster, Laird, and Rubin 1977) to obtain maximum likelihood estimates for the observer variation when ground truth is missing or

partially available. This has served as a foundation for several following contributions (Ipeirotis, Provost, and Wang 2010; Raykar et al. 2010; Whitehill et al. 2009; Zhou et al. 2012), placing David and Skene’s algorithm in a crowdsourcing context and enriching it for building performance-sensitive pricing schemes. The APM model enhances these quality definitions by leveraging the fact that the error rates of workers are directly affected by the access path that they follow, which allows for efficient optimization.

Query and crowd access optimization. In crowdsourced databases, quality assurance and crowd access optimization are envisioned as part of the query optimizer, which needs to estimate the query plans not only according to the cost but also to their accuracy and latency. Previous work (Franklin et al. 2011; Marcus et al. 2011; Parameswaran et al. 2012) focuses on building declarative query languages with support for processing crowdsourced data. The proposed optimizers define the execution order of operators in query plans and map crowdsourcable operators to micro-tasks. In our work, we propose a complementary approach by ensuring the quality of each single operator executed by the crowd.

Crowd access optimization is similar to the expert selection problem in decision-making. However, the assumption that the selected individuals will answer may no longer hold. Previous studies based on this assumption are (Karger, Oh, and Shah 2011; Ho, Jabbari, and Vaughan 2013; Jung and Lease 2013). The proposed methods are nevertheless effective for task recommendation, spam detection, and performance evaluation.

Diversity for quality. Relevant studies in management science (Hong and Page 2004; Lamberson and Page 2012) emphasize diversity and define the notion of *types* to refer to highly correlated forecasters. Another work that targets groups of workers is introduced by (Li, Zhao, and Fuxman 2014). This technique discards groups that do not prove to be the best ones. (Venanzi et al. 2014) instead, refers to groups as *communities* and all of them are used for aggregation but not for optimization. Other systems like CrowdSearcher by (Brambilla et al. 2014) and CrowdSTAR by (Nushi et al. 2015) support cross-community task allocation.

Conclusion

We introduced the Access Path Model, a novel crowd model that captures and exploits diversity as an inherent property of large-scale crowdsourcing. This model lends itself to efficient greedy crowd access optimization. The resulting plan has strong theoretical guarantees, since, as we prove, the information gain objective is submodular in our model. The presented theoretical results are of general interest and applicable to a wide range of variable selection and experimental design problems. We evaluated our approach on three real-world crowdsourcing datasets. Experiments demonstrate that our approach can be used to seamlessly handle critical problems in crowdsourcing such as quality assurance and crowd access optimization even in situations of anonymized and sparse data.

Acknowledgements. This work was supported in part by the Swiss National Science Foundation, and Nano-Tera.ch program as part of the Opensense II project. The authors would also like to thank Brian Galebach and Sharad Goel for providing the Probability Sports dataset.

References

- Brambilla, M.; Ceri, S.; Mauri, A.; and Volonterio, R. 2014. Community-based crowdsourcing. In *WWW*, 891–896.
- Cover, T. M., and Thomas, J. A. 2012. *Elements of information theory*. John Wiley & Sons.
- Dawid, A. P., and Skene, A. M. 1979. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied statistics* 20–28.
- De Campos, L. M. 2006. A scoring function for learning bayesian networks based on mutual information and conditional independence tests. *JMLR* 7:2149–2187.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)* 1–38.
- Feige, U. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM* 45:314–318.
- Franklin, M. J.; Kossmann, D.; Kraska, T.; Ramesh, S.; and Xin, R. 2011. Crowddb: answering queries with crowdsourcing. In *SIGMOD*, 61–72. ACM.
- Ho, C.-J.; Jabbari, S.; and Vaughan, J. W. 2013. Adaptive task assignment for crowdsourced classification. In *ICML*, 534–542.
- Hong, L., and Page, S. E. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc. of the National Academy of Sciences of USA* 101(46):16385–16389.
- Hui, T. W. L. C. P., and Li, C. J. Z. W. 2015. Hear the whole story: Towards the diversity of opinion in crowdsourcing markets. *VLDB*.
- Ipeirotis, P. G.; Provost, F.; and Wang, J. 2010. Quality management on amazon mechanical turk. In *Proc. of the ACM SIGKDD workshop on human computation*, 64–67. ACM.
- Jung, H. J., and Lease, M. 2013. Crowdsourced task routing via matrix factorization. *arXiv preprint arXiv:1310.5142*.
- Karger, D. R.; Oh, S.; and Shah, D. 2011. Budget-optimal crowdsourcing using low-rank matrix approximations. In *49th Annual Allerton Conference*, 284–291. IEEE.
- Kazai, G.; Kamps, J.; and Milic-Frayling, N. 2012. The face of quality in crowdsourcing relevance labels: Demographics, personality and labeling accuracy. In *CIKM*, 2583–2586. ACM.
- Khuller, S.; Moss, A.; and Naor, J. S. 1999. The budgeted maximum coverage problem. *Inform. Process. Lett.* 70(1):39–45.
- Krause, A., and Guestrin, C. 2005a. Near-optimal nonmyopic value of information in graphical models. In *UAI*.
- Krause, A., and Guestrin, C. 2005b. A note on the budgeted maximization of submodular functions.
- Lamberson, P., and Page, S. E. 2012. Optimal forecasting groups. *Management Science* 58(4):805–810.
- Li, H.; Zhao, B.; and Fuxman, A. 2014. The wisdom of minority: Discovering and targeting the right group of workers for crowdsourcing. In *Proc. of the 23rd WWW*.
- Marcus, A.; Wu, E.; Karger, D. R.; Madden, S.; and Miller, R. C. 2011. Crowdsourced databases: Query processing with people. *CIDR*.
- Nemhauser, G.; Wolsey, L.; and Fisher, M. 1978. An analysis of the approximations for maximizing submodular set functions. *Math. Prog.* 14:265–294.
- Nushi, B.; Alonso, O.; Hentschel, M.; and Kandylas, V. 2015. Crowdstar: A social task routing framework for online communities. In *ICWE*, 219–230.
- Parameswaran, A. G.; Park, H.; Garcia-Molina, H.; Polyzotis, N.; and Widom, J. 2012. Deco: declarative crowdsourcing. In *CIKM*.
- Parameswaran, A.; Boyd, S.; Garcia-Molina, H.; Gupta, A.; Polyzotis, N.; and Widom, J. 2014. Optimal crowd-powered rating and filtering algorithms. *VLDB*.
- Pochampally, R.; Sarma, A. D.; Dong, X. L.; Meliou, A.; and Srivastava, D. 2014. Fusing data with correlations. In *SIGMOD*.
- Raykar, V. C.; Yu, S.; Zhao, L. H.; Valadez, G. H.; Florin, C.; Bogoni, L.; and Moy, L. 2010. Learning from crowds. *JMLR*.
- Selinger, P. G.; Astrahan, M. M.; Chamberlin, D. D.; Lorie, R. A.; and Price, T. G. 1979. Access path selection in a relational database management system. In *SIGMOD*, 23–34. ACM.
- Surowiecki, J. 2005. *The wisdom of crowds*. Random House LLC.
- Sviridenko, M. 2004. A note on maximizing a submodular set function subject to knapsack constraint. *Operations Research Letters* v.(32):41–43.
- Venanzi, M.; Guiver, J.; Kazai, G.; Kohli, P.; and Shokouhi, M. 2014. Community-based bayesian aggregation models for crowdsourcing. In *Proc. of the 23rd WWW*, 155–164.
- Welinder, P.; Branson, S.; Mita, T.; Wah, C.; Schroff, F.; Belongie, S.; and Perona, P. 2010. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology.
- Whitehill, J.; Ruvolo, P.; Wu, T.; Bergsma, J.; and Movellan, J. R. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *NIPS*.
- Zhou, D.; Basu, S.; Mao, Y.; and Platt, J. C. 2012. Learning from the wisdom of crowds by minimax entropy. In *NIPS*, 2195–2203.