Peer Prediction with Heterogeneous Users

Citation

Agarwal, Arpit, Debmalya Mandal, David Parkes, and Nisarg Shah. 2017. Peer Prediction with Heterogeneous Users. Proceedings of the 2017 ACM Conference on Economics and Computation, Cambridge, MA, June 26 - 30, 2017, 81-98.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41542960>

Terms of Use

This article was downloaded from Harvard University's DASH repository, WARNING: This file should NOT have been available for downloading from Harvard University's DASH repository.

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Peer Prediction with Heterogeneous Users

ARPIT AGARWAL, University of Pennsylvania

DEBMALYA MANDAL, Harvard University

DAVID C. PARKES, Harvard University

NISARG SHAH, Harvard University

Peer prediction mechanisms incentivize agents to truthfully report their signals, in the absence of a verification mechanism, by comparing their reports with those of their peers. Prior work in this area is essentially restricted to the case of homogeneous agents, whose signal distributions are identical. This is limiting in many domains, where we would expect agents to differ in taste, judgment and reliability. Although the Correlated Agreement (CA) mechanism [30] can be extended to handle heterogeneous agents, the new challenge is with the efficient estimation of agent signal types. We solve this problem by clustering agents based on their reporting behavior, proposing a mechanism that works with clusters of agents and designing algorithms that learn such a clustering. In this way, we also connect peer prediction with the Dawid and Skene [5] literature on latent types. We retain the robustness against coordinated misreports of the CA mechanism, achieving an approximate incentive guarantee of ϵ -informed truthfulness. We show on real data that this incentive approximation is reasonable in practice, and even with a small number of clusters.

CCS Concepts: •Theory of computation → Algorithmic game theory and mechanism design;

Additional Key Words and Phrases: Peer Prediction; Information Elicitation; Clustering; Tensor Decomposition

1 INTRODUCTION

Peer prediction is the problem of information elicitation without verification. Peer prediction mechanisms incentivize users to provide honest reports when the reports cannot be verified, either because there is no objective ground truth or because it is costly to acquire the ground truth. Peer prediction mechanisms leverage correlation in the reports of peers in order to score contributions. In recent years, peer prediction has been studied in several domains, including peer assessment in massively open online courses (MOOCs) [Gao et al., 2016, Shnayder and Parkes, 2016], feedback on the local places in a city [Mandal et al., 2016], and in the context of collaborative sensing platforms [Radanovic and Faltings, 2015a].

The simplest peer prediction mechanism is *output agreement*, which pairs up two users and rewards them in the event that their reports agree (the ESP game [von Ahn and Dabbish, 2004] can be interpreted this way). However, output agreement is not incentive aligned for reports of *a priori* unlikely signals. As a result, there has been a lot of attention in recent years to finding methods that work more generally, as well as providing robustness to coordinated misreports.

We thank Adish Singla and Dengyong Zhou for providing us the Stanford Dogs (Dogs) and the Waterbird dataset (WB). We also acknowledge the useful feedback from anonymous reviewers.

This work is supported in part by Google, the SEAS TomKat fund, and NSF grant CCF-1301976.

Authors' email addresses: Arpit Agarwal: arpit@seas.upenn.edu; Debmalaya Mandal: dmandal@g.harvard.edu; David C. Parkes: parkes@eecs.harvard.edu; Nisarg Shah: nisarg@g.harvard.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. XXXX-XXXX/2017/1-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/3033274.3085127>

All existing, general methods are essentially restricted to settings with homogeneous participants, whose signal distributions are identical. This is a poor fit with many suggested applications of peer prediction. Consider for example, the problem of peer assessment in MOOCs. DeBoer et al. [2013] and Wilkowski et al. [2014] observe that students differ based on their geographical locations, educational backgrounds, and level of commitment, and indeed the heterogeneity of assessment is clear from a study of Coursera data [Kulkarni et al., 2015]. A similar problem occurs in determining whether a news headline is offensive or not. Depending on which social community a user belongs to, we should expect to get different opinions [Zafar et al., 2016]. In feedback about places such as restaurants or in responses to marketing surveys, user responses will vary because of idiosyncratic differences in taste.

One obstacle behind designing peer prediction mechanisms for heterogeneous agents is an impossibility result. No mechanism can provide strict incentive for truth-telling to a population of heterogeneous agents without knowledge of their signal distributions [Radanovic and Faltings, 2015c]. This negative result holds for minimal mechanisms, which only elicit signals and not beliefs from agents. One way to alleviate this problem, without going to non-minimal mechanisms, is to use the reports from agents across multiple tasks to estimate the signal distribution of an agent. This is our goal: to use minimal peer prediction mechanisms with heterogeneous agents, while using reports from agents for both learning and scoring. We also want to provide robustness against coordinated misreports.

As a starting point, we consider the *correlated agreement* (CA) mechanism proposed in Shnayder et al. [2016]. If the agents are homogeneous and the designer has knowledge of their joint signal distribution, then the CA mechanism is *informed truthful*, i.e. no (even coordinated) strategy profile provides more expected payment than truth-telling and the expected payment under an uninformed strategy (where an agent’s report is independent of signal) is strictly less than the expected payment under truth-telling. These two properties remove the incentive for coordinated deviations and promote effort in acquiring and reporting signals, respectively. In a detail-free variation, where the designer learns the signal distribution from reports, approximate incentive alignment is provided (still maintaining the second property as a strict guarantee.) The detail-free CA mechanism can be extended to handle agent heterogeneity, but a naive approach would require learning the joint signal distributions between every pair of agents, and the total number of reports that need to be collected would be prohibitive for many settings.

Our contributions: We design the first minimal and detail-free mechanism for peer prediction with heterogeneous agents, and establish that the sample complexity is linear in the number of agents, while providing an incentive guarantee of approximate informed truthfulness. First inAs with the CA mechanism, this is a multi-task mechanism in that each agent makes reports across multiple tasks. Since our mechanism has a learning component, the task assignments to agents should be such that both the goals of incentive alignment and learning are simultaneously achieved. We consider two assignment schemes under which these goals can be achieved and analyze the sample complexity of our methods for these schemes.

The new mechanism clusters the agents based on their reported behavior and learns the pairwise correlations between these clusters. The clustering introduces one component of the incentive approximation, and could be problematic in the absence of a good clustering such that agents within a cluster behave similarly. Using eight real-world datasets, which contain reports of users on crowdsourcing platforms for multiple labeling tasks, we show that the clustering error is small in practice even when using a relatively small number of clusters. The additional approximation that results from needing to learn the pairwise correlations between clusters can be made arbitrarily small using a sufficient number of signal reports.

Another contribution is to connect, we think for the first time, the peer prediction literature with the extensive and influential literature on latent, confusion matrix models of label aggregation [Dawid and Skene, 1979b]. The Dawid-Skene model assumes that signals are generated independently, conditional on a latent attribute of a task and according to an agent’s confusion matrix. We estimate the cluster average confusion matrices using recent developments in tensor decomposition algorithms [Anandkumar et al., 2014, Zhang et al., 2016]. These matrices are used to design reward schemes and achieve approximate informed truthfulness.

In effect, the mechanism learns how to translate one agent’s signal reports so that they map onto the signal reports of other agents. For example, in the context of a MOOC, an agent in an “extremal” cluster that only uses the ‘A’ and ‘E’ grades may be scored positively when her ‘A’ report is matched with a ‘B’ report from another agent assigned to the same essay and a member of the “accurate” cluster. Similarly, an agent in the “contrarian” cluster may be scored positively when her ‘E’ report is matched with an ‘A’ report from another agent in the “accurate” cluster and assigned to the same essay. In practice, our mechanism would be trained on the data collected during a semester of peer assessment reports, with students then clustered, cluster-pairwise signal distributions estimated, and scores eventually assigned (this scoring done retrospectively).

1.1 Related Work

We focus our discussion on related work about minimal mechanisms, but remark that we are aware of no, non-minimal mechanisms (following from [Prelec, 2004]) that handle agent heterogeneity. Miller et al. [2005] introduced the peer prediction problem, and proposed an incentive-aligned mechanism for a single-task setting. However their mechanism requires knowledge of the joint signal distribution and is vulnerable to coordinated misreports. In regard to coordinated misreports, Jurca et al. [2009] show how to eliminate uninformative, pure-strategy equilibria through a three-peer mechanism, and Kong et al. [2016] provide a method to design robust, single-task, binary signal mechanisms (but need knowledge of the joint signal distribution). Frongillo and Witkowski [2017] provide a characterization of minimal (single task) peer prediction mechanisms.

Witkowski and Parkes [2013] introduced the combination of learning and peer prediction, coupling the estimation of the signal prior together with the shadowing mechanism. Some results make use of reports from a large population. Radanovic and Faltings [2015b], for example, establish robust incentive properties in a large-market limit where both the number of tasks and the number of agents assigned to each task grow without bound. Radanovic et al. [2016] provide complementary theoretical results, giving a mechanism in which truthfulness is the equilibrium with highest payoff in the asymptote of a large population and with a structural property on the signal distribution.

Dasgupta and Ghosh [2013] show that robustness to coordinated misreports can be achieved for binary signals in a small population by using a multi-task mechanism. The idea is to reward agents if they provide the same signal on the same task, but punish them if one agent’s report on one task is the same as another’s on another task. The correlated agreement (CA) mechanism [Shnayder et al., 2016] generalizes this mechanism to handle multiple signals, and uses reports to estimate the correlation structure on pairs of signals without compromising incentives. In related work, Kong and Schoenebeck [2016] show that many peer prediction mechanisms can be derived within a single information-theoretic framework. Their results use different technical tools than Shnayder et al. [2016], and also include a different multi-signal generalization of the Dasgupta-Ghosh mechanism that provides robustness against coordinated misreports in the limit of a large number of tasks. Shnayder et al. [2016] adopt replicator dynamics as a model of population learning in peer prediction, and confirm that these multi-task mechanisms (including Kamble et al. [2015]) are successful at avoiding uninformed equilibria.

There are very few results on handling agent heterogeneity in peer prediction. For binary signals, the method of [Dasgupta and Ghosh \[2013\]](#) is likely to be an effective solution because their assumption on correlation structure will tend to hold for most reasonable models of heterogeneity. But it will break down for more than binary signals, as explained by [Shnayder et al. \[2016\]](#). Moreover, although the CA mechanism can in principle be extended to handle heterogeneity, it is not clear how the required statistical information about signal distributions can be efficiently learned and coupled with an analysis of approximate incentives. For a setting with binary signals and where each task has one of a fixed number of latent types, [Kamble et al. \[2015\]](#) design a mechanism that provides strict incentive compatibility for a suitably large number of heterogeneous agents, and when the number of tasks grows without bound (while allowing each agent to only provide reports on a bounded number of tasks). Their result is restricted to binary signals, and requires a strong regularity assumption on the generative model of signals. Finally, we consider only binary effort of a user, i.e. the agent either invests effort and receives a signal or does not invest effort and receives an uninformed signal. Other work that has considered the importance of motivating effort in the context of peer prediction includes [Liu and Chen \[2016\]](#) and [Witkowski et al. \[2013\]](#).¹ See [Mandal et al. \[2016\]](#) for a setting with heterogeneous tasks but homogeneous agents.

2 MODEL

Let notation $[t]$ denote $\{1, \dots, t\}$ for $t \in \mathbb{N}$. We consider a population of agents $P = [\ell]$, and use indices such as p and q to refer to agents from this population. There is a set of tasks $M = [m]$. When an agent performs a task, she receives a signal from $N = [n]$. As mentioned before, we assume that the effort of an agent is binary. We also assume that the tasks are *ex ante* identical, that is, the signals of an agent for different tasks are sampled i.i.d. We use S_p to denote the random variable for the signal of agent p for a task.

We work in the setting where the agents are heterogeneous, i.e., the distribution of signals can be different for different agents. We say they vary by *signal type*. Let $D_{p,q}(i, j)$ denote the joint probability that agent p receives signal i while agent q receives signal j on a random task. Let $D_p(i)$ and $D_q(j)$ denote the corresponding marginal probabilities. We define the *Delta matrix* $\Delta_{p,q}$ between agents p and q as

$$\Delta_{p,q}(i, j) = D_{p,q}(i, j) - D_p(i) \cdot D_q(j), \quad \forall i, j \in [n]. \quad (1)$$

[Shnayder et al. \[2016\]](#) show that it is without loss of generality for the class of mechanisms we study in this paper to assume that an agent's strategy is uniform across different tasks. Given this, let R_p denote the random variable for the report of agent p for a task. The strategy of agent p , denoted F^p , defines the distribution on reports for each possible signal i , with $F_{ir}^p = \Pr(R_p = r | S_p = i)$. The collection of agent strategies, denoted $\{F^p\}_{p \in P}$, is the *strategy profile*. A strategy of agent p is *informed* if there exist distinct $i, j \in [n]$ and $r \in [n]$ such that $F_{ir}^p \neq F_{jr}^p$, i.e., if not all rows of F^p are identical. We say that a strategy is *uninformed* otherwise.

2.1 Multi-Task Peer Prediction

We consider *multi-task peer prediction* mechanisms. In particular, we adapt the *correlated agreement* (CA) mechanism [[Shnayder et al., 2016](#)] to our setting with heterogeneous agents. For every pair of agents $p, q \in P$, we define a *scoring matrix* $S_{p,q} : [n] \times [n] \rightarrow \mathbb{R}$ as a means of scoring agent

¹[Cai et al. \[2015\]](#) work in a different model, showing how to achieve optimal statistical estimation from data provided by rational agents. These authors focus on the cost of effort, and do not consider misreports (and their mechanism is vulnerable to coordinated misreports).

reports.² We randomly divide (without the knowledge of an agent) the set of tasks performed by each agent into nonempty sets of *bonus tasks* and *penalty tasks*. For agent p , we denote the set of her bonus tasks by M_1^p and the set of her penalty tasks by M_2^p .

To calculate the payment to an agent p for a bonus task $t \in M_1^p$, we do the following:

- (1) Randomly select an agent $q \in P \setminus \{p\}$ such that $t \in M_1^q$, and the set $M_2^p \cup M_2^q$ has at least 2 distinct tasks, and call q the *peer* of p .
- (2) Pick tasks $t' \in M_2^p$ and $t'' \in M_2^q$ randomly such that $t' \neq t''$ (t' and t'' are the penalty tasks for agents p and q respectively)
- (3) Let the reports of agent p on tasks t and t' be r_p^t and $r_p^{t'}$, respectively and the reports of agent q on tasks t and t'' be r_q^t and $r_q^{t''}$ respectively.
- (4) The payment of agent p for task t is then $S_{p,q}(r_p^t, r_q^t) - S_{p,q}(r_p^{t'}, r_q^{t''})$.

The total payment to an agent is the sum of payments for the agent's bonus tasks. We assume that from an agent's perspective, every other agent is equally likely to be her peer. This requires agents not to know each other's task assignments. In Section 4 we give two task assignment schemes such that from an agent's perspective all peers are equally likely.

The expected payment to agent p for any bonus task performed by her—equal across all bonus tasks as the tasks are *ex ante* identical—is given by

$$u_p(F^p, \{F^q\}_{q \neq p}) = \frac{1}{\ell - 1} \sum_{q \neq p} \sum_{i,j} \Delta_{p,q}(i,j) \sum_{r_p, r_q} S_{p,q}(r_p, r_q) \cdot F_{ir_p}^p \cdot F_{jr_q}^q. \quad (2)$$

Using the fact that there exist optimal solutions to linear functions that are extremal, it is easy to show that there always exists an optimal strategy for agent p that is deterministic (see also Shnayder et al. [2016]).

LEMMA 2.1. *For every player p , and any strategies of others, there always exists an optimal strategy F^p maximizing u_p that is deterministic.*

Hereafter, we assume without loss of generality that agent strategies are deterministic. For a deterministic strategy F^p of agent p , we will slightly abuse notation and write F_i^p to denote the signal reported by agent p when she observes signal i . For a deterministic strategy profile $\{F^q\}_{q \in P}$, the expected payment to agent p is

$$u_p(F^p, \{F^q\}_{q \neq p}) = \frac{1}{\ell - 1} \sum_{q \neq p} \sum_{i,j} \Delta_{p,q}(i,j) \cdot S_{p,q}(F_i^p, F_j^q). \quad (3)$$

2.2 Informed Truthfulness

Following Shnayder et al. [2016], we define the notion of approximate informed truthfulness for a multi-task peer prediction mechanism.

Definition 2.2. (ε -informed truthfulness) We say that a multi-task peer prediction mechanism is ε -informed truthful, for some $\varepsilon \geq 0$, if and only if for every strategy profile $\{F^q\}_{q \in P}$ and every agent $p \in P$, we have $u_p(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) \geq u_p(F^p, \{F^q\}_{q \neq p}) - \varepsilon$, where $\mathbb{I}$ is the truthful strategy, and $u_p(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) > u_p(F_0^p, \{F^q\}_{q \neq p})$ where F_0^p is an uninformed strategy.

An ε -informed truthful mechanism ensures that every agent prefers (up to ε) the truthful strategy profile over any other strategy profile, and strictly prefers the truthful strategy profile over any uninformed strategy. For a small ε , this is responsive to the main concerns about incentives in peer

²Multi-task peer prediction mechanisms for homogeneous agents need a single Delta and a single scoring matrix.

prediction: a minimal opportunity for coordinated manipulations, and a strict incentive to invest effort in collecting and reporting an informative signal.³

2.3 Agent Clustering

While the natural extension of the detail-free CA mechanism to our setting with heterogeneous agents preserves informed truthfulness, it would require learning the Delta matrix between every pair of agents. In order to learn this from data without additional assumptions, we would require $\Omega(\ell^2)$ samples, which will often be impractical. Rather, the number of reports in a practical mechanism should scale closer to linearly in the number of agents.

In response, we will assume that agents can be (approximately) clustered into a bounded number K of agent signal types, with all agents of the same type having nearly identical signal distributions. Let $G_1, \dots, G_K$ denote a partitioning of agents into K clusters. With a slight abuse of notation, we also use $G(p)$ to denote the cluster to which agent p belongs.

Definition 2.3. We say that clustering $G_1, \dots, G_K$ is ε_1 -accurate, for some $\varepsilon_1 \geq 0$, if for every pair of agents $p, q \in P$,

$$\|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1 \leq \varepsilon_1, \quad (4)$$

where $\Delta_{G(p),G(q)}$ is the *cluster Delta matrix* between clusters $G(p)$ and $G(q)$, defined as the average of the Delta matrices between agents in $G(p)$ and agents in $G(q)$:

$$\Delta_{G_s, G_t} = \frac{1}{|G_s| \times |G_t|} \sum_{p \in G_s, q \in G_t} \Delta_{p,q}.$$

Our mechanism will use $\Delta_{G(p),G(q)}$ in place of $\Delta_{p,q}$ to define the scoring matrix $S_{p,q}$. Thus, the accuracy of the clustering will translate directly to the incentive guarantee of the mechanism and we seek a clustering that is as accurate as possible. For the incentive properties to hold, we will need high between-cluster similarity between agents (i.e., two agents in cluster 1 interact with agents in other clusters in a similar way).

There is an inverse relationship between the number of clusters, K , and the cluster accuracy, ε_1 . On the one extreme, we can let every agent be a separate cluster ($K = \ell$), which results in $\varepsilon_1 = 0$. On the other hand, a small number of clusters is essential for reasonable sample complexity, as we will need to learn K^2 cluster Delta matrices. In Section 4, we give a learning algorithm that can learn all the pairwise cluster Delta matrices with $\tilde{O}(K)$ samples given a clustering of the agents. In Section 5, we show using real-world data that a reasonably small clustering error can be achieved with relatively few clusters.

3 CORRELATED AGREEMENT FOR CLUSTERED, HETEROGENEOUS AGENTS

In this section we define mechanism CAHU, presented as Algorithm 1, which takes as input a clustering as well as estimates of the cluster Delta matrices. Specifically, CAHU takes as input a clustering $G_1, \dots, G_K$, together with estimates of cluster Delta matrices $\{\bar{\Delta}_{G_s, G_t}\}_{s,t \in [K]}$.

Definition 3.1. We say that a clustering $\{G_s\}_{s \in [K]}$ and the estimates $\{\bar{\Delta}_{G_s, G_t}\}_{s,t \in [K]}$ are $(\varepsilon_1, \varepsilon_2)$ -accurate if

- $\|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1 \leq \varepsilon_1$ for all agents $p, q \in P$, i.e., the clustering is ε_1 -accurate, and
- $\|\Delta_{G_s, G_t} - \bar{\Delta}_{G_s, G_t}\|_1 \leq \varepsilon_2$ for all clusters $s, t \in [K]$, i.e., the cluster Delta matrix estimates are ε_2 -accurate.

³We do not model the cost of effort explicitly in this paper, but a binary cost model (effort $\rightarrow$ signal, no-effort $\rightarrow$ no signal) can be handled in a straightforward way. See Shnayder et al. [2016].

When we have a clustering and estimates of the delta matrices which are $(\varepsilon_1, \varepsilon_2)$ -accurate, we prove that the CAHU mechanism is $(\varepsilon_1 + \varepsilon_2)$ -informed truthful. In Section 4, we present algorithms that can learn an ε_1 -accurate clustering and ε_2 -accurate estimates of cluster Delta matrices.

Throughout the rest of this section, we will use ε_1 to denote the clustering error and ε_2 to denote the learning error. We remark that the clustering error ε_1 is determined by the level of similarity present in agent signal-report behavior, as well as the number of clusters K used, whereas the learning error ε_2 depends on how many samples the learning algorithm sees.

Algorithm 1 Mechanism CAHU

Let $\text{Sign}(x) = 1$ if $x > 0$, and 0 otherwise.

Input:

A clustering $G_1, \dots, G_K$ such that $\|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1 \leq \varepsilon_1$ for all $p, q \in P$;
 estimates $\{\bar{\Delta}_{G_s, G_t}\}_{s,t \in [K]}$ such that $\|\bar{\Delta}_{G_s, G_t} - \Delta_{G_s, G_t}\|_1 \leq \varepsilon_2$ for all $s, t \in [K]$; and
 for each agent $p \in P$, her bonus tasks M_1^p , penalty tasks M_2^p , and responses $\{r_b^p\}_{b \in M_1^p \cup M_2^p}$.

Method:

- 1: **for** every agent $p \in P$ **do**
 - 2: **for** every task $b \in M_1^p$ **do** ▷ Reward response r_b^p
 - 3: $q \leftarrow$ uniformly at random conditioned on $b \in M_1^q \cup M_2^q$ and (either $|M_2^q| \geq 2, |M_2^p| \geq 2$
 or $M_2^q \neq M_2^p$) ▷ Peer agent
 - 4: Pick tasks $b' \in M_2^p$ and $b'' \in M_2^q$ randomly such that $b' \neq b''$ ▷ Penalty tasks
 - 5: $S_{G(p),G(q)} \leftarrow \text{Sign}(\bar{\Delta}_{G(p),G(q)})$
 - 6: Reward to agent p for task b is $S_{G(p),G(q)}(r_b^p, r_b^q) - S_{G(p),G(q)}(r_{b'}^p, r_{b''}^q)$
 - 7: **end for**
 - 8: **end for**
-

3.1 Analysis of CAHU

Because mechanism CAHU uses scoring matrix $S_{G(p),G(q)}$ to reward agent p given peer agent q in place of a separate scoring matrix $S_{p,q}$ for every pair of agents (p, q) , the payoff equation (3) transforms to

$$u_p(F^p, \{F^q\}_{q \neq p}) = \frac{1}{\ell - 1} \sum_{q \in P \setminus \{p\}} \sum_{i,j} \Delta_{p,q}(i,j) \cdot S_{G(p),G(q)}(F_i^p, F_j^q). \quad (5)$$

Next, we prove a sequence of results that show that mechanism CAHU is $(\varepsilon_1 + \varepsilon_2)$ -informed truthful when the clustering is ε_1 -accurate and the cluster Delta matrix estimates are ε_2 -accurate. We first show that the mechanism would be (exactly) informed truthful if it used the Delta matrices between agents instead of using estimates of cluster Delta matrices. In the following we use $u_p^*(\cdot)$ (resp. $u_p(\cdot)$) to denote the utility of agent p under the CA (resp. CAHU) mechanism.

LEMMA 3.2. *For a strategy profile $\{F^q\}_{q \in P}$ and an agent $p \in P$, define*

$$u_p^*(F^p, \{F^q\}_{q \neq p}) = \frac{1}{\ell - 1} \sum_{q \in P \setminus \{p\}} \sum_{i,j} \Delta_{p,q}(i,j) \cdot S_{p,q}(F_i^p, F_j^q),$$

where $S_{p,q}(i,j) = \text{Sign}(\Delta_{p,q}(i,j))$ for all $i, j \in [n]$. Then, $u_p^*(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) \geq u_p^*(F^p, \{F^q\}_{q \neq p})$. Moreover, for any uninformed strategy $F^p = r$, $u_p^*(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) > u_p^*(r, \{F^q\}_{q \neq p})$.

The proof of all theorems and lemmas appear in the full version of the paper.⁴ We now use this result to prove one of our main theorems, which is the desired informed truthfulness guarantee about mechanism CAHU.

THEOREM 3.3. *With $(\varepsilon_1, \varepsilon_2)$ -accurate clustering and learning, mechanism CAHU is $(\varepsilon_1 + \varepsilon_2)$ -informed truthful if $\min_p u_p^*(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) > \varepsilon_1$. In particular,*

- (1) *For every profile $\{F^q\}_{q \in P}$ and agent $p \in P$, we have $u_p(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) \geq u_p(F^p, \{F^q\}_{q \neq p}) - \varepsilon_1 - \varepsilon_2$.*
- (2) *For any uninformed strategy F_0^p , $u_p(F_0^p, \{F^q\}_{q \neq p}) < u_p(\mathbb{I}, \{\mathbb{I}\}_{q \neq p})$.*

The CAHU mechanism always ensures that there is no strategy profile which gives an expected utility more than $\varepsilon_1 + \varepsilon_2$ above truthful reporting. The condition $\min_p u_p^*(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) > \varepsilon_1$ is required to ensure that any uninformed strategy gives strictly less than the truth-telling equilibrium. This is important to promote effort in collecting and reporting an informative signal. Writing it out, this condition requires that for each agent p the following holds :

$$\frac{1}{\ell - 1} \sum_{q \neq p} \sum_{i, j: \Delta_{p, q}(i, j) > 0} \Delta_{p, q}(i, j) > \varepsilon_1. \quad (6)$$

In particular, a sufficient condition for this property is that for every pair of agents the expected reward on a bonus task in the CA mechanism when making truthful reports is at least ε_1 , i.e. for every pair of agents p and q ,

$$\sum_{i, j: \Delta_{p, q}(i, j) > 0} \Delta_{p, q}(i, j) > \varepsilon_1. \quad (7)$$

In turn, as pointed out in Shnayder et al. [2016], the LHS in (7) quantity can be interpreted as a measure of how much positive correlation there is in the joint distribution on signals between a pair of agents. Note that it is not important that this is same-signal correlation. For example, this quantity would be large between an accurate and an always-wrong agent in a binary-signal domain, since the positive correlation would be between one agent’s report and the flipped report from the other agent.

The incentive properties of the mechanism are retained when used together with learning the cluster structure and cluster Delta matrices. This follows because the agents can do no better than have the mechanism be designed with the correct scoring matrices (see Shnayder et al. [2016]).

4 LEARNING THE AGENT SIGNAL TYPES

In this section, we provide algorithms for learning a clustering of agent signal types from reports, and further, for learning the cluster pairwise Δ matrices. The estimates of the Δ matrices can then be used to give an approximate-informed truthful mechanism. Along the way, we couple our methods with the latent “confusion matrix” methods of Dawid and Skene [1979b].

Recall that m is the total number of tasks about which reports are collected. Reports on m_1 of these tasks will also be used for clustering, and reports on a further m_2 of these tasks will be used for learning the cluster pairwise Δ matrices. We consider two different schemes for assigning agents to tasks for the purpose of clustering and learning (see Figures 1 and 2):

- (1) **Fixed Task Assignment:** Each agent is assigned to the same, random subset of tasks of size $m_1 + m_2$ of the given m tasks.
- (2) **Uniform Task Assignment:** For clustering, we select two agents r_1 and r_2 , uniformly at random, to be *reference agents*. These agents are assigned to a subset of tasks of size

⁴The full version is available on the authors’ website.

Fig. 1. Fixed Task Assignment

Fig. 2. Uniform Task Assignment

Table 1. Sample complexity for the CAHU mechanism. The rows indicate the assignment scheme and the columns indicate the modeling assumption. Here ℓ is the number of agents, n is the number of signals, ϵ' is a parameter that controls learning accuracy $\frac{\epsilon}{\epsilon'}$, γ is a clustering parameter, K is the number of clusters, and m_1 (resp. m_2) is the size of the set of tasks from which the tasks used for clustering (resp. learning) are sampled.

	No Assumption	Dawid-Skene
Fixed Assignment	Clustering: $\tilde{O}\left(\frac{\ell n^2}{\gamma^2}\right)$ Learning: $\tilde{O}\left(\frac{K n^2}{(\epsilon')^2}\right)$	Clustering: $\tilde{O}\left(\frac{\ell n^2}{\gamma^2}\right)$ Learning: $\tilde{O}\left(\frac{\ell n^7}{(\epsilon')^2}\right)$
Uniform Assignment	Clustering: $\tilde{O}\left(\frac{\ell n^2}{\gamma^2} + m_1\right)$ Learning: $\tilde{O}\left(K m_2^{7/8} \sqrt{\frac{n^2}{(\epsilon')^2}}\right)$	Clustering: $\tilde{O}\left(\frac{\ell n^2}{\gamma^2} + m_1\right)$ Learning: $\tilde{O}\left(\frac{K n^7}{(\epsilon')^2}\right)^\dagger$

$m_1(< m)$. For all other agents, we then assign a required number of tasks, s_1 , uniformly at random from the set of m_1 tasks. For learning the cluster pairwise Δ -matrices, we also assign one agent from each cluster to some subset of tasks of size s_2 , selected uniformly at random from a second set of $m_2(< m - m_1)$ tasks.

For each assignment scheme, the analysis establishes that there are enough agents who have done a sufficient number of joint tasks. Table 1 summarizes the sample complexity results, stating them under two different assumptions about the way in which signals are generated.

4.1 Clustering

We proceed by presenting and analyzing a simple clustering algorithm.

Definition 4.1. A clustering $G_1, \dots, G_K$ is ϵ -good if for some $\gamma > 0$

$$G(q) = G(r) \Rightarrow \|\Delta_{pq} - \Delta_{pr}\|_1 \leq \epsilon - 4\gamma \quad \forall p \in [\ell] \setminus \{q, r\} \quad (8)$$

$$G(q) \neq G(r) \Rightarrow \|\Delta_{pq} - \Delta_{pr}\|_1 > \epsilon \quad \forall p \in [\ell] \setminus \{q, r\} \quad (9)$$

We first show that an ϵ -good clustering, if exists, must be unique.

THEOREM 4.2. *Suppose there exist two clustering $\{G_j\}_{j \in [K]}$ and $\{T_i\}_{i \in [K']}$ that are ϵ -good. Then $K' = K$ and $G_j = T_{\pi(j)}$ for some permutation π over $[K]$.*

[†]For an arbitrary m_2 , this bound is $K m_2$ as long as m_2 is $\Omega\left(n^7/(\epsilon')^2\right)$

^{*}In the no assumption approach (resp. Dawid-Skene Model), ϵ' is the error in the estimation of the joint probability distribution (resp. aggregate confusion matrix).

Algorithm 2 Clustering

Input: ε, γ such that there exists an ε -good clustering with parameter γ .

Output: A clustering $\{\hat{G}_t\}_{t=1}^{\hat{K}}$

```

1:  $\hat{G} \leftarrow \emptyset, \hat{K} \leftarrow 0$  ▷  $\hat{G}$  is the list of clusters,  $\hat{K} = |\hat{G}|$ 
2: Make a new cluster  $\hat{G}_1$  and add agent 1
3: Add  $\hat{G}_1$  to  $\hat{G}, \hat{K} \leftarrow \hat{K} + 1$ 
4: for  $i = 2, \dots, \ell$  do
5:   for  $t \in [\hat{K}]$  do
6:     Pick an arbitrary agent  $q_t \in \hat{G}_t$ 
7:     Pick  $p_t \in [I] \setminus \{i, q_t\}$  (Fixed) or  $p_t \in \{r_1, r_2\} \setminus \{i, q_t\}$  (Uniform), such that  $p_t$  has at
       least  $\Omega(\frac{n^2 \log(K\ell/\delta)}{\gamma^2})$  tasks in common with both  $q_t$  and  $i$ 
8:     Let  $\bar{\Delta}_{p_t, q_t}$  be the empirical Delta matrix from reports of agents  $p_t$  and  $q_t$ 
9:     Let  $\bar{\Delta}_{p_t, i}$  be the empirical Delta matrix from reports of agents  $p_t$  and  $i$ 
10:   end for
11:   if  $\exists t \in [\hat{K}] : \|\bar{\Delta}_{p_t, q_t} - \bar{\Delta}_{p_t, i}\|_1 \leq \varepsilon - 2\gamma$  then
12:     add  $i$  to  $\hat{G}_t$  (with ties broken arbitrarily for  $t$ )
13:   else
14:     Make a new cluster  $\hat{G}_{\hat{K}+1}$  and add agent  $i$  to it
15:     Add  $\hat{G}_{\hat{K}+1}$  to  $\hat{G}, \hat{K} \leftarrow \hat{K} + 1$ 
16:   end if
17: end for
    
```

Since there is a unique ε -good clustering (up to a permutation), we will refer to this clustering as the *correct clustering*. The assumption that there exists an ε -good clustering is stronger than Equation (4) introduced earlier.

In particular, identifying the correct clustering requires need eq. (9), i.e. the Δ -matrices of two agents belonging to two different clusters are different with respect to every other agent. So, we need low inter-cluster similarities in addition to high intra-cluster similarities. The pseudo-code for the clustering algorithm is presented in Algorithm 2. This algorithm iterates over the agents, and forms clusters in a greedy manner. First, we prove that as long as we can find an agent p_t that has $\Omega(\frac{n^2 \log(\ell/\delta)}{\gamma^2})$ tasks in common with both q_t and i , then the clustering produced by Algorithm 2 is correct with probability at least $1 - \delta$.

THEOREM 4.3. *If for all $i \in P$ and $q_t \in G(i)$, there exists p_t which has $\Omega(\frac{n^2 \log(\ell/\delta)}{\gamma^2})$ tasks in common with both q_t and i , then Algorithm 2 recovers the correct clustering i.e. $\hat{G}_t = G_t$ for $t = 1, \dots, K$ with probability at least $1 - \delta$.*

Next we show how the assumption in regard to task overlap is satisfied under each assignment scheme, and characterize the sample complexity of learning the clusterings under each scheme. In the fixed assignment scheme, all the agents are assigned to the same set of $m_1 = \Omega(\frac{n^2}{\gamma^2} \log(K\ell/\delta))$

Fig. 3. Algorithm 2 checks whether i and q_t are in the same cluster by estimating Δ_{p_t, q_t} and $\Delta_{p_t, i}$.

tasks. Thus, for each agent pair q_t and i , any other agent in the population can act as p_t . The total number of tasks performed is $O\left(\frac{\ell n^2}{\gamma^2} \log(K\ell/\delta)\right)$.

In the uniform assignment scheme, we select two agents r_1 and r_2 uniformly at random to be reference agents, and assign these agents to each of $m_1 = \Omega\left(\frac{n^2}{\gamma^2} \log(K\ell/\delta)\right)$ tasks. For all other agents we then assign $s_1 = \Omega\left(\frac{n^2}{\gamma^2} \log(K\ell/\delta)\right)$ tasks uniformly at random from this set of m_1 tasks. If $m_1 = s_1$, then the uniform task assignment is the same as fixed task assignment. However, in applications, for example [Karger et al., 2011], where one wants the task assignments to be more uniform across tasks, it will make sense to use a larger value of m_1 . The reference agent r_1 can act as p_t for all agent pairs q_t and i other than r_1 . Similarly, reference r_2 can act as p_t for all agent pairs q_t and i other than r_2 . If $q_t = r_1$ and $i = r_2$ or $q_t = r_2$ and $i = r_1$, then any other agent can act as p_t . The total number of tasks performed is $\Omega\left(\frac{\ell n^2}{\gamma^2} \log(K\ell/\delta) + m_1\right)$, which is sufficient for the high probability result.

4.2 Learning the Cluster Pairwise Δ Matrices

We proceed now under the assumption that the agents are clustered into K groups, $G_1, \dots, G_K$. Our goal is to estimate the cluster-pairwise delta matrices Δ_{G_s, G_t} as required by Algorithm 1. We estimate the Δ_{G_s, G_t} under two different settings: when we have no model of the signal distribution, and in the Dawid-Skene latent attribute model.

Algorithm 3 Learning- Δ -No-Assumption

- 1: **for** $t = 1, \dots, K$ **do**
 - 2: Chose agent $q_t \in G_t$ arbitrarily.
 - 3: **end for**
 - 4: **for** each pair of clusters G_s, G_t **do**
 - 5: Let q_s and q_t be the chosen agents for G_s and G_t , respectively.
 - 6: Let $\bar{D}_{q_s, q_t}$ be the empirical estimate of D_{q_s, q_t} such that $\|\bar{D}_{q_s, q_t} - D_{q_s, q_t}\|_1 \leq \varepsilon'$ with probability at least $1 - \delta/K^2$
 - 7: Let $\bar{\Delta}_{q_s, q_t}$ be the empirical Delta matrix computed using $\bar{D}_{q_s, q_t}$
 - 8: Set $\bar{\Delta}_{G_s, G_t} = \bar{\Delta}_{q_s, q_t}$
 - 9: **end for**
-

4.2.1 Learning the Δ -Matrices with No Assumption. We first characterize the sample complexity of learning the Δ -matrices in the absence of any modeling assumptions. In order to estimate $\bar{\Delta}_{G_s, G_t}$, Algorithm 3 first picks agent q_s from cluster G_s , estimates $\bar{\Delta}_{q_s, q_t}$ and use this estimate in place of $\bar{\Delta}_{G_s, G_t}$. For the fixed assignment scheme, we assign the agents q_s to the same set of tasks of size $O\left(\frac{n^2}{(\varepsilon')^2} \log(K/\delta)\right)$. For the uniform assignment scheme, we assign the agents to subsets of tasks of an appropriate size among the pool of m_2 tasks.

THEOREM 4.4. *Given an ε -good clustering $\{G_s\}_{s=1}^K$, if the number of shared tasks between any pair of agents q_s, q_t is $O\left(\frac{n^2}{(\varepsilon')^2} \log(K/\delta)\right)$, then algorithm 3 guarantees that for all s, t , $\|\bar{\Delta}_{G_s, G_t} - \Delta_{G_s, G_t}\|_1 \leq 3\varepsilon' + 2\varepsilon$ with probability at least $1 - \delta$. The total number of samples collected by the algorithm is $O\left(\frac{Kn^2}{(\varepsilon')^2} \log(K/\delta)\right)$ (resp. $O\left(Km_2^{7/8} \sqrt{\frac{n^2}{(\varepsilon')^2} \log(K/\delta)}\right)$ w.h.p.) under the fixed (resp. uniform) assignment scheme.*

4.2.2 Learning the Δ -matrices Under the Dawid-Skene Model. In this section, we assume that the agents receive signals according to the [Dawid and Skene \[1979a\]](#) model. In this model, each task has a latent attribute and each agent has a confusion matrix to parameterize its signal distribution conditioned on this latent value. More formally:

- Let $\{\pi_k\}_{k=1}^n$ denote the prior probability over n latent values.
- Agent p has *confusion matrix* $C^p \in \mathbb{R}^{n \times n}$, such that $C_{ij}^p = D_p(S_p = j | T = i)$ where T is the latent value. Given this, the joint signal distribution for a pair of agents p and q is

$$D_{p,q}(S_p = i, S_q = j) = \sum_{k=1}^n \pi_k C_{ki}^p C_{kj}^q, \quad (10)$$

and the marginal signal distribution for agent p is

$$D_p(S_p = i) = \sum_{k=1}^n \pi_k C_{ki}^p. \quad (11)$$

For cluster G_t , we write $C^t = \frac{1}{|G_t|} \sum_{p \in G_t} C^p$ to denote the aggregate confusion matrix of G_t . As before, we assume that we are given an ε -good clustering, $G_1, \dots, G_K$, of the agents. Our goal is to provide an estimate of the Δ_{G_s, G_t} -matrices.

Lemma 4.5 proves that in order to estimate Δ_{G_s, G_t} within an L1 distance of ε' , it is enough to estimate the aggregate confusion matrices within an L1 distance of $\varepsilon'/4$. So in order to learn the pairwise delta matrices between clusters, we first ensure that for each cluster G_t , we have $\|\bar{C}^t - C^t\|_1 \leq \varepsilon'/4$ with probability at least $1 - \delta/K$, and then use the following formula to compute the delta matrices:

$$\Delta_{G_s, G_t}(i, j) = \sum_{k=1}^n \pi_k \bar{C}_{ki}^s \bar{C}_{kj}^t - \sum_{k=1}^n \pi_k \bar{C}_{ki}^s \sum_{k=1}^n \pi_k \bar{C}_{kj}^t \quad (12)$$

LEMMA 4.5. *For all G_a, G_b , $\|\bar{C}^a - C^a\|_1 \leq \varepsilon'/4$ and $\|\bar{C}^b - C^b\|_1 \leq \varepsilon'/4 \Rightarrow \|\bar{\Delta}_{G_a, G_b} - \Delta_{G_a, G_b}\| \leq \varepsilon'$.*

We turn now to estimating the aggregate confusion matrix of each cluster. Let us assume for now that the agents are assigned to the tasks according to the uniform assignment scheme, i.e. agent p belonging to cluster G_a is assigned to a subset of B_a tasks selected uniformly at random from a pool of m_2 tasks. For cluster G_a , we choose $B_a = \frac{m_2}{|G_a|} \ln(\frac{m_2 K}{\beta})$. This implies:

- (1) For each $j \in [m_2]$, $\Pr[\text{agent } p \in G_a \text{ completes task } j] = \frac{\log(m_2 K / \beta)}{|G_a|}$, i.e. each agent p in G_a is equally likely to complete every task j .
- (2) $\Pr[\text{task } j \text{ is unlabeled by } G_a] = \left(1 - \frac{\log(m_2 K / \beta)}{|G_a|}\right)^{|G_a|} \leq \frac{\beta}{m_2 K}$. Taking a union bound over the m_2 tasks and K clusters, we get the probability that any task is unlabeled is at most β . Now if we choose $\beta = 1/\text{poly}(m_2)$, we observe that with probability at least $1 - 1/\text{poly}(m_2)$, each task j is labeled by some agent in each cluster when $B_a = \tilde{O}(\frac{m_2}{|G_a|})$.

Left to do is to provide an algorithm and sample complexity for learning the aggregate confusion matrices. For this, we will use n dimensional unit vectors to denote the reports of the agents (recall that there are n possible signals). In particular agent p 's report on task j , $r_{pj} \in \{0, 1\}^n$. If p 's report on task j is c , then the c -th coordinate of r_{pj} is 1 and all the other coordinates are 0. The expected value of agent p 's report on j th task is $\mathbb{E}[r_{pj}] = \sum_{k=1}^n \pi_k C_{kj}^p$. The aggregated report for a cluster G_t is given as $R_{tj} = \frac{1}{|G_t|} \sum_{p \in G_t} r_{pj}$.

Suppose we want to estimate the aggregate confusion matrix C^1 of some cluster G_1 . To do so, we first pick three clusters G_1, G_2 and G_3 and write down the corresponding cross moments. Let

(a, b, c) be a permutation of the set $\{1, 2, 3\}$. We have:

$$M_1 = E[R_{aj}] = \sum_k \pi_k C_k^a \quad (13)$$

$$M_2 = E[R_{aj} \otimes R_{bj}] = \sum_k \pi_k C_k^a \otimes C_k^b \quad (14)$$

$$M_3 = E[R_{aj} \otimes R_{bj} \otimes R_{cj}] = \sum_k \pi_k C_k^a \otimes C_k^b \otimes C_k^c \quad (15)$$

The cross moments are asymmetric, however using Theorem 3.6 of [Anandkumar et al. \[2014\]](#), we can write the cross-moments in a symmetric form.

LEMMA 4.6. Assume that the vectors $\{C_1^t, \dots, C_n^t\}$ are linearly independent for each $t \in \{1, 2, 3\}$. For any permutation (a, b, c) of the set $\{1, 2, 3\}$ define

$$\begin{aligned} R'_{aj} &= E[R_{cj} \otimes R_{bj}] (E[R_{aj} \otimes R_{bj}])^{-1} R_{aj} \\ R'_{bj} &= E[R_{cj} \otimes R_{aj}] (E[R_{bj} \otimes R_{aj}])^{-1} R_{bj} \\ M_2 &= E[R'_{aj} \otimes R'_{bj}] \text{ and } M_3 = E[R'_{aj} \otimes R'_{bj} \otimes R_{cj}] \end{aligned}$$

$$\text{Then } M_2 = \sum_{k=1}^n \pi_k C_k^c \otimes C_k^c \text{ and } M_3 = \sum_{k=1}^n \pi_k C_k^c \otimes C_k^c \otimes C_k^c$$

We cannot compute the moments exactly, but rather estimate the moments from samples observed from different tasks. Furthermore, for a given task j , instead of exactly computing the aggregate label R_{gj} , we select one agent p uniformly at random from G_g and use agent p 's report on task j as a proxy for R_{gj} . We will denote the corresponding report as $\tilde{R}_{gj}$. The next lemma proves that the cross-moments of $\{\tilde{R}_{gj}\}_{g=1}^K$ and $\{R_{gj}\}_{g=1}^K$ are the same.

- LEMMA 4.7. (1) For any group G_a , $E[\tilde{R}_{aj}] = E[R_{aj}]$
 (2) For any pair of groups G_a and G_b , $E[\tilde{R}_{aj} \otimes \tilde{R}_{bj}] = E[R_{aj} \otimes R_{bj}]$
 (3) For any three groups G_a, G_b and G_c , $E[\tilde{R}_{aj} \otimes \tilde{R}_{bj} \otimes \tilde{R}_{cj}] = E[R_{aj} \otimes R_{bj} \otimes R_{cj}]$

The next set of equations show how to approximate the moments M_2 and M_3 :

$$\hat{R}'_{aj} = \left(\frac{1}{m_2} \sum_{j'=1}^{m_2} \tilde{R}_{cj'} \otimes \tilde{R}_{bj'} \right) \left(\frac{1}{m_2} \sum_{j'=1}^{m_2} \tilde{R}_{aj'} \otimes \tilde{R}_{bj'} \right)^{-1} \tilde{R}_{aj} \quad (16)$$

$$\hat{R}'_{bj} = \left(\frac{1}{m_2} \sum_{j'=1}^{m_2} \tilde{R}_{cj'} \otimes \tilde{R}_{aj'} \right) \left(\frac{1}{m_2} \sum_{j'=1}^{m_2} \tilde{R}_{bj'} \otimes \tilde{R}_{aj'} \right)^{-1} \tilde{R}_{bj} \quad (17)$$

$$\hat{M}_2 = \frac{1}{m_2} \sum_{j'=1}^{m_2} \hat{R}'_{aj'} \otimes \hat{R}'_{bj'} \text{ and } \hat{M}_3 = \frac{1}{m_2} \sum_{j'=1}^{m_2} \hat{R}'_{aj'} \otimes \hat{R}'_{bj'} \otimes \tilde{R}_{cj'} \quad (18)$$

We use the tensor decomposition algorithm (4) on $\hat{M}_2$ and $\hat{M}_3$ to recover the aggregate confusion matrix $\hat{C}^c$ and $\hat{\Pi}$, where $\hat{\Pi}$ is a diagonal matrix whose k -th component is $\hat{\pi}_k$, an estimate of π_k . In order to analyze the sample complexity of Algorithm 4, we need to make some mild assumptions about the problem instance. For any two clusters G_a and G_b , define $S_{ab} = E[R_{aj} \otimes R_{bj}] = \sum_{k=1}^n \pi_k C_k^a \otimes C_k^b$. We make the following assumptions:

Algorithm 4 Estimating Aggregate Confusion Matrix

Input: K clusters of agents $G_1, G_2, \dots, G_K$ and the reports $\tilde{R}_{gj} \in \{0, 1\}^n$ for $j \in [m]$ and $g \in [K]$
Output: Estimate of the aggregate confusion matrices $\tilde{C}^g$ for all $g \in [K]$

- 1: Partition the K clusters into groups of three
- 2: **for** Each group of three clusters $\{g_a, g_b, g_c\}$ **do**
- 3: **for** $(a, b, c) \in \{(g_b, g_c, g_a), (g_c, g_a, g_b), (g_a, g_b, g_c)\}$ **do**
- 4: Compute the second and the third order moments $\hat{M}_2 \in \mathbb{R}^{n \times n}$, $\hat{M}_3 \in \mathbb{R}^{n \times n \times n}$. ▶
- Compute $\tilde{C}^g$ and $\tilde{\Pi}$ by tensor decomposition
- 5: Compute whitening matrix $\hat{Q} \in \mathbb{R}^{n \times n}$ such that $\hat{Q}^T \hat{M}_2 \hat{Q} = I$
- 6: Compute eigenvalue-eigenvector pairs $(\hat{\alpha}_k, \hat{v}_k)_{k=1}^n$ of the whitened tensor $\hat{M}_3(\hat{Q}, \hat{Q}, \hat{Q})$ by using the robust tensor power method
- 7: Compute $\hat{w}_k = \hat{\alpha}_k^{-2}$ and $\hat{\mu}_k = (\hat{Q}^T)^{-1} \hat{\alpha} \hat{v}_k$
- 8: For $k = 1, \dots, n$ set the k -th column of $\tilde{C}^c$ by some $\hat{\mu}_k$ whose k -th coordinate has the greatest component, then set the k -th diagonal entry of $\tilde{\Pi}$ by $\hat{w}_k$
- 9: **end for**
- 10: **end for**

- (1) There exists $\sigma_L > 0$ such that $\sigma_n(S_{ab}) \geq \sigma_L$ for each pair of clusters a and b , where $\sigma_n(M)$ is the n -th smallest eigenvalue of M .
- (2) $\kappa = \min_{t \in [k]} \min_{s \in [n]} \min_{r \neq s} \{C_{rr}^t - C_{rs}^t\} > 0$

The first assumption implies that the matrices S_{ab} are non-singular. The second assumption implies that within a group, the probability of assigning the correct label is always higher than the probability of assigning any incorrect label. The following theorem gives the number of tasks each agent needs to complete to get an ε' -estimate of the aggregate confusion matrices.

THEOREM 4.8. *For any $\varepsilon' \leq \min \left\{ \frac{31}{\sigma_L^2}, \frac{\kappa}{2} \right\} n^2$ and $\delta > 0$, if the size of the universe of shared tasks m_2 is at least $O \left(\frac{n^7}{(\varepsilon')^2 \sigma_L^{11}} \log \left(\frac{nK}{\delta} \right) \right)$, then we have $\|\tilde{C}^t - C^t\|_1 \leq \varepsilon'$ for each cluster G_t . The total number of samples collected by Algorithm 4 is $\tilde{O}(Km_2)$ under the uniform assignment scheme.*

Discussion. If the algorithm chooses $m_2 = \tilde{O} \left(\frac{n^7}{(\varepsilon')^2 \sigma_L^{11}} \right)$, then the total number of samples collected under the uniform assignment scheme is at most $\tilde{O} \left(\frac{n^7}{(\varepsilon')^2 \sigma_L^{11}} \right)$. So far we have analyzed the Dawid-Skene model under the uniform assignment scheme. When the assignment scheme is fixed, the moments of R_{aj} and $\tilde{R}_{aj}$ need not be the same. In this case we will have to run Algorithm 4 with respect to the actual aggregate labels $\{R_{gj}\}_{g=1}^K$. This requires collecting samples from every member of a cluster, leading to a sample complexity of $O \left(\frac{\ell n^7}{(\varepsilon')^2 \sigma_L^{11}} \log \left(\frac{nK}{\delta} \right) \right)$

In order to estimate the confusion matrices, Zhang et al. [2016] require each agent to provide at least $O(n^5 \log((\ell + n)/\delta)/(\varepsilon')^2)$ samples. Our algorithm requires $O(n^7 \log(nK/\delta)/(\varepsilon')^2)$ samples from each cluster. The increase of n^2 in the sample complexity comes about because we are estimating the aggregate confusion matrices in L1 norm instead of the infinity norm. Moreover when the number of clusters is small ($K \ll \ell$), the number of samples required from each cluster does not grow with ℓ . This improvement is due to the fact that, unlike Zhang et al. [2016], we do not have to recover individual confusion matrices from the aggregate confusion matrices.

Note that the Dawid and Skene [1979b] based approach, for the uniform assignment scheme, does not require all agents to provide reports on the same set of shared tasks. Rather, we need

that for each group of three clusters (as partitioned by Algorithm 4 on line 1) and each task, there exists one agent from the three clusters who completes the same task. In particular the reports for different tasks can be acquired from different agents within the same cluster. The assignment scheme makes sure that this property holds with high probability.

We now briefly compare the learning algorithms under the no-assumptions and model-based approach. When it is difficult to assign agents to the same tasks, and when the number of signals is small (which is often true in practice), the Dawid-Skene method has a strong advantage. Another advantage of the Dawid-Skene method is that the learning error ε' can be made arbitrarily small since each aggregate confusion matrix can be learned with arbitrary accuracy, whereas the true learning error of the no-assumption approach is at least 2ε (see Theorem 4.4), and depends on the problem instance.

5 CLUSTERING EXPERIMENTS

In this section, we study the clustering error on eight real-world, crowdsourcing datasets. Six of these datasets are from the *SQUARE* benchmark [Sheshadri and Lease, 2013], selected to ensure a sufficient density of worker labels across different latent attributes as well as the availability of latent attributes for sufficiently many tasks. In addition, we also use the Stanford *Dogs* dataset [Khosla et al., 2011] and the *Expressions* dataset [Mozafari et al., 2012, 2014]. Below, we briefly describe the format of task, the number of agents ℓ , and the number of signals n for each dataset.⁵ The description of the dataset is available in the full version of the paper.

Recall that the maximum incentive an agent has to use a non-truthful strategy in the CAHU mechanism can be upper-bounded in terms of two sources of error:

- The clustering error. This represents how “clusterable” the agents are. From theory, we have the upper bound $\varepsilon_1 = \max_{p,q \in [\ell]} \|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1$.
- The learning error. This represents how accurate our estimates for the cluster Delta matrices are. From theory, we have the upper bound $\varepsilon_2 = \max_{i,j \in [K]} \|\Delta_{G_i,G_j} - \bar{\Delta}_{G_i,G_j}\|_1$.

Based on this, the CAHU mechanism is $(\varepsilon_1 + \varepsilon_2)$ -informed truthful (Theorem 3.3) where ε_1 is the clustering error and ε_2 is the learning error. Even with the best clustering, the clustering error ε_1 cannot be made arbitrarily small because it depends on how close the signal distributions of the agents are as well as the number of clusters. In contrast, the learning error ε_2 of the no-assumption approach is $3\varepsilon' + 2\varepsilon$, (theorem 4.4) where the error due to ε' can indeed be made arbitrarily small by simply acquiring more data about agents’ behavior. Similarly, the learning error ε_2 in the Dawid-Skene approach can be made arbitrarily small by acquiring more agent reports (theorem 4.8). Hence the total error is dominated by the clustering error ε_1 .

For this reason, we focus on the clustering error, and show that it can be small even with a relatively small number of clusters. Rather than the use the weak bound $\max_{p,q \in [\ell]} \|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1$ on the clustering error (which is nevertheless helpful for our theoretical results), we use the following tighter bound from the proof of Theorem 3.3.

$$|u_p^*(\mathbb{I}, \{\mathbb{I}\}_{q \neq p}) - u_p(\mathbb{I}, \{\mathbb{I}\}_{q \neq p})| = \left| \frac{1}{(\ell - 1)} \sum_{q \in P \setminus \{p\}} \sum_{i,j} \Delta_{p,q}(i,j) (\text{Sign}(\Delta_{p,q})_{i,j} - \text{Sign}(\bar{\Delta}_{G(p),G(q)})_{i,j}) \right| \quad (19)$$

⁵We filter each dataset to remove tasks for which the latent attribute is unknown, and remove workers who only perform such tasks. ℓ is the number of agents that remain after filtering.

The datasets specify the latent value of each task. Because of this, we can adopt the Dawid-Skene model, and estimate the confusion matrices from the frequency which with each agent p reports each label j in the case of each latent attribute i .

Typical clustering algorithms take a distance metric over the space of data points and attempt to minimize the maximum cluster diameter, which is the maximum distance between any two points within a cluster. In contrast, our objective (the tighter bound on the incentive in Equation (19)) is a complex function of the underlying confusion matrices. We therefore compare two approaches:

- 1) We cluster the confusion matrices using the standard k -means++ algorithm with the $L2$ norm distance (available in Matlab), and hope that resulting clustering leads to a small incentive bound.⁶
- 2) In the following lemma, we derive a distance metric between confusion matrices for which the maximum cluster diameter is provably an upper bound on the incentive, and use k -means++ with this metric (implemented in Matlab).⁷ Note that computing this metric requires knowledge of the prior over the latent attribute.

LEMMA 5.1. *For all agents p, q, r , we have $\|\Delta_{p,q} - \Delta_{p,r}\|_1 \leq 2 \cdot \sum_k \pi_k \sum_j |C_{kj}^q - C_{kj}^r|$.*

LEMMA 5.2. *For every pair of agents p, q , we have*

$$\|\Delta_{p,q} - \Delta_{G(p),G(q)}\|_1 \leq 2 \cdot \max_{a,b,c \in P: G(a)=G(b)} \|\Delta_{a,c} - \Delta_{b,c}\|_1.$$

Note that $\sum_k \pi_k \sum_j |C_{kj}^q - C_{kj}^r| \leq \|C^q - C^r\|_1$ because $\sum_j |C_{lj}^q - C_{lj}^r| \leq \|C^q - C^r\|_1$. Lemma 5.1, along with Lemma 5.2, shows that the incentive is upper bounded by four times the maximum cluster diameter under our metric. For each dataset, we vary the number of clusters K from 5% to 15% of the number of agents in the dataset. We repeat the experiment 20 times, and select the clustering that produces the smallest incentive bound. Figures 4 and 5 show the incentive bound achieved using the standard $L2$ metric and using our custom metric, respectively. We see that the incentive bound is small compared to the maximum payment of 1 by CAHU, even with the number of clusters K as small as 15% of the number of workers. The number of agents does not seem to affect this bound as long as the number of clusters is small relative to the number of agents. Using our custom metric leads to a clustering with a noticeably better incentive bound.

6 CONCLUSION

We have provided the first, general solution to the problem of peer prediction with heterogeneous agents. This is a compelling research direction, where new theory and algorithms can help to guide practice. In particular, heterogeneity is likely to be quite ubiquitous because of differences in taste, context, judgment, and reliability across users. Beyond testing these methods in a real-world application such as marketing surveys, there remain interesting directions for ongoing research. For example, is it possible to solve this problem with similar sample complexity but without a clustering approach, and to couple methods of peer prediction with optimal methods for inference in crowdsourced classification [Ok et al., 2016], and with methods for task assignment in budgeted

⁶We use $L2$ norm rather than $L1$ norm because the standard k -means++ implementation uses as the centroid of a cluster the confusion matrix that minimizes the sum of distances from the confusion matrices of the agents in the cluster. For $L2$ norm, this amounts to averaging over the confusion matrices, which is precisely what we want. For $L1$ norm, this amounts to taking a pointwise median, which does not even result in a valid confusion matrix. Perhaps for this reason, we observe that using the $L1$ norm performs worse.

⁷For computing the centroid of a cluster, we still average over the confusion matrices of the agents in the cluster. Also, since the algorithm is no longer guaranteed to converge (indeed, we observe cycles), we restart the algorithm when a cycle is detected, at most 10 times.

Fig. 4. Incentive bound on each of the 8 different data sets when using k -means++ with the L_2 norm

Fig. 5. Incentive bound on each of the 8 different data sets when using k -means++ with our custom metric

settings [Karger et al., 2014]? This should include attention to adaptive assignment schemes [Khetan and Oh, 2016] that leverage generalized Dawid-Skene models [Zhou et al., 2015], and could connect with recent progress on task heterogeneity within peer prediction [Mandal et al., 2016].

REFERENCES

- Animashree Anandkumar, Rong Ge, Daniel J Hsu, Sham M Kakade, and Matus Telgarsky. 2014. Tensor Decompositions for Learning Latent Variable Models. *Journal of Machine Learning Research* 15, 1 (2014), 2773–2832.
- Yang Cai, Constantinos Daskalakis, and Christos Papadimitriou. 2015. Optimum Statistical Estimation with Strategic Data Sources. In *Proceedings of The 28th Conference on Learning Theory*. 280–296.
- Anirban Dasgupta and Arpita Ghosh. 2013. Crowdsourced Judgement Elicitation with Endogenous Proficiency. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 319–330.
- Alexander Philip Dawid and Allan M Skene. 1979a. Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm. *Applied statistics* (1979), 20–28.
- Philip A Dawid and Allan M Skene. 1979b. Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm. *Applied statistics* 28 (1979), 20–28.
- Jennifer DeBoer, Glenda S Stump, Daniel Seaton, and Lori Breslow. 2013. Diversity in MOOC Students’ Backgrounds and Behaviors in Relationship to Performance in 6.002 x. In *Proceedings of the Sixth Learning International Networks Consortium Conference*, Vol. 4.
- Rafael Frongillo and Jens Witkowski. 2017. A Geometric Perspective on Minimal Peer Prediction. *ACM Transactions on Economics and Computation (TEAC)* (2017).
- Alice Gao, James R Wright, and Kevin Leyton-Brown. 2016. Incentivizing Evaluation via Limited Access to Ground Truth: Peer-Prediction Makes Things Worse. *EC 2016 Workshop on Algorithmic Game Theory and Data Science* (2016).
- Radu Jurca, Boi Faltings, et al. 2009. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research* 34, 1 (2009), 209.
- Vijay Kamble, David Marn, Nihar Shah, Abhay Parekh, and Kannan Ramachandran. 2015. Truth Serums for Massively Crowdsourced Evaluation Tasks. *The 5th Workshop on Social Computing and User-Generated Content* (2015).
- David R. Karger, Sewoong Oh, and Devavrat Shah. 2011. Budget-Optimal Task Allocation for Reliable Crowdsourcing Systems. *CoRR* abs/1110.3564 (2011). <http://arxiv.org/abs/1110.3564>
- David R. Karger, Sewoong Oh, and Devavrat Shah. 2014. Budget-Optimal Task Allocation for Reliable Crowdsourcing Systems. *Operations Research* 62, 1 (2014), 1–24. <https://doi.org/10.1287/opre.2013.1235>
- Ashish Khetan and Sewoong Oh. 2016. Achieving budget-optimality with adaptive schemes in crowdsourcing. In *Annual Conference on Neural Information Processing Systems*. 4844–4852.
- Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. 2011. Novel Dataset for Fine-Grained Image Categorization. In *First CVPR Workshop on Fine-Grained Visual Categorization*.
- Yueying Kong, Katrina Ligett, and Grant Schoenebeck. 2016. Putting Peer Prediction Under the Micro (economic) scope and Making Truth-telling Focal. In *International Conference on Web and Internet Economics*. Springer, 251–264.
- Yueying Kong and Grant Schoenebeck. 2016. A Framework For Designing Information Elicitation Mechanism That Rewards Truth-telling. (2016). <http://arxiv.org/abs/1605.01021>

- Chinmay Kulkarni, Koh Pang Wei, Huy Le, Daniel Chia, Kathryn Papadopoulos, Justin Cheng, Daphne Koller, and Scott R Klemmer. 2015. Peer and Self Assessment in Massive Online Classes. In *Design thinking research*. Springer, 131–168.
- Yang Liu and Yiling Chen. 2016. Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices. *arXiv preprint arXiv:1611.09219* (2016).
- Debmalya Mandal, Matthew Leifer, David C Parkes, Galen Pickard, and Victor Shnayder. 2016. Peer Prediction with Heterogeneous Tasks. *NIPS 2016 Workshop on Crowdsourcing and Machine Learning* (2016).
- Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51 (2005), 1359–1373.
- Barzan Mozafari, Purnamrita Sarkar, Michael J. Franklin, Michael I. Jordan, and Samuel Madden. 2012. Active Learning for Crowd-Sourced Databases. *CoRR* abs/1209.3686 (2012).
- Barzan Mozafari, Purnamrita Sarkar, Michael J. Franklin, Michael I. Jordan, and Samuel Madden. 2014. Scaling Up Crowdsourcing to Very Large Datasets: A Case for Active Learning. *PVLDB* 8, 2 (2014), 125–136.
- Jungseul Ok, Sewoong Oh, Jinwoo Shin, and Yung Yi. 2016. Optimality of Belief Propagation for Crowdsourced Classification. In *Proc. 33rd Int. Conf. on Machine Learning (ICML)*. 535–544.
- Drazen Prelec. 2004. A Bayesian Truth Serum For Subjective Data. *Science* 306, 5695 (2004), 462.
- Goran Radanovic and Boi Faltings. 2015a. Incentive Schemes for Participatory Sensing. In *Proc. Int. Conf. on Autonomous Agents and Multiagent Systems, AAMAS*. 1081–1089.
- Goran Radanovic and Boi Faltings. 2015b. Incentive Schemes for Participatory Sensing. In *AAMAS 2015*.
- Goran Radanovic and Boi Faltings. 2015c. Incentives for Subjective Evaluations with Private Beliefs. In *Proc. 29th AAAI Conf. on Art. Intell. (AAAI’15)*. 1014–1020.
- Goran Radanovic, Boi Faltings, and Radu Jurca. 2016. Incentives for effort in crowdsourcing using the peer truth serum. *ACM Transactions on Intelligent Systems and Technology (TIST)* 7, 4 (2016), 48.
- Aashish Sheshadri and Matthew Lease. 2013. SQUARE: A Benchmark for Research on Computing Crowd Consensus. In *Proc. 1st AAAI Conf. on Human Computation (HCOMP)*. 156–164.
- Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C Parkes. 2016. Informed Truthfulness in Multi-Task Peer Prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 179–196.
- Victor Shnayder, Rafael Frongillo, and David C. Parkes. 2016. Measuring Performance Of Peer Prediction Mechanisms Using Replicator Dynamics. In *Proc. 25th Int. Joint Conf. on Art. Intell. (IJCAI’16)*. 2611–2617.
- Victor Shnayder and David C Parkes. 2016. Practical Peer Prediction for Peer Assessment. In *Fourth AAAI Conference on Human Computation and Crowdsourcing*.
- Luis von Ahn and Laura Dabbish. 2004. Labeling Images with a Computer Game. In *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI’04)*. 319–326.
- Julia Wilkowsky, Amit Deutsch, and Daniel M Russell. 2014. Student Skill and Goal Achievement in the Mapping with Google MOOC. In *Proceedings of the first ACM conference on Learning@ scale conference*. ACM, 3–10.
- Jens Witkowski, Yoram Bachrach, Peter Key, and David Parkes. 2013. Dwelling on the negative: Incentivizing effort in peer prediction. In *First AAAI Conference on Human Computation and Crowdsourcing*.
- Jens Witkowski and David C Parkes. 2013. Learning the Prior in Minimal Peer Prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*. 14.
- Muhammad Bilal Zafar, Krishna P Gummadi, and Cristian Danescu-Niculescu-Mizil. 2016. Message Impartiality in Social Media Discussions. In *ICWSM*. 466–475.
- Yuchen Zhang, Xi Chen, Dengyong Zhou, and Michael I Jordan. 2016. Spectral Methods Meet EM: A Provably Optimal Algorithm for Crowdsourcing. *Journal of Machine Learning Research* 17, 102 (2016), 1–44.
- Dengyong Zhou, Qiang Liu, John C. Platt, Christopher Meek, and Nihar B. Shah. 2015. Regularized Minimax Conditional Entropy for Crowdsourcing. *CoRR* abs/1503.07240 (2015). <http://arxiv.org/abs/1503.07240>