Behavior Identification in Two-stage Games for Incentivizing Citizen Science Exploration

Yexiang Xue[1], Ian Davies[2], Daniel Fink[2], Christopher Wood[2], and Carla P. Gomes[1]

[1]Computer Science Department, Cornell University
{yexiang,gomes}@cs.cornell.edu
[2]Cornell Lab of Ornithology
{id99,daniel.fink,chris.wood}@cornell.edu

Abstract. We consider two-stage games in which a leader seeks to direct the activities of independent agents by offering incentives. A good leader's strategy requires an understanding of the agents' utilities and the ability to predict agent behavior. Moreover, the optimization of outcomes requires an agent behavior model that can be efficiently incorporated into the leader's model. Here we address the agent behavior modeling problem and show how it can be used to reduce bias in a challenging citizen science application. Adapting ideas from Discrete Choice Modeling in behavioral economics, we develop a probabilistic behavioral model that takes into account variable patterns of human behavior and suboptimal actions. By modeling deviations from baseline behavior we are able to accurately predict future behavior based on limited, sparse data. We provide a novel scheme to fold the agent model into a bi-level optimization as a single Mixed Integer Program, and scale up our approach by adding redundant constraints, based on novel insights of an easy-hard-easy phase transition phenomenon. We apply our methodology to a game called Avicaching, in collaboration with eBird, a well-established citizen science program that collects bird observations for conservation. Field results show that our behavioral model performs well and that the incentives are remarkably effective at steering citizen scientists' efforts to reduce bias by exploring under-sampled areas. Moreover, the data collected from Avicaching improves the performance of species distribution models.

1 Introduction

Many game applications involve a leader, who commits to a strategy before her followers. Thus in order to come up with an optimal strategy, the leader must factor in the reasoning process of her followers. This leads naturally to the following bi-level optimization:

$$\begin{aligned} \textbf{Leader:}\ \ & \underset{\mathbf{a}_1}{\text{maximize}} \quad U_L(\mathbf{a}_1, \mathbf{a}_2), \\ & \text{subject to} \quad \textbf{Followers:}\ \ \mathbf{a}_2 \leftarrow \underset{\mathbf{a}_2}{\text{argmax}}\, U_F(\mathbf{a}_2, \mathbf{a}_1). \end{aligned}$$

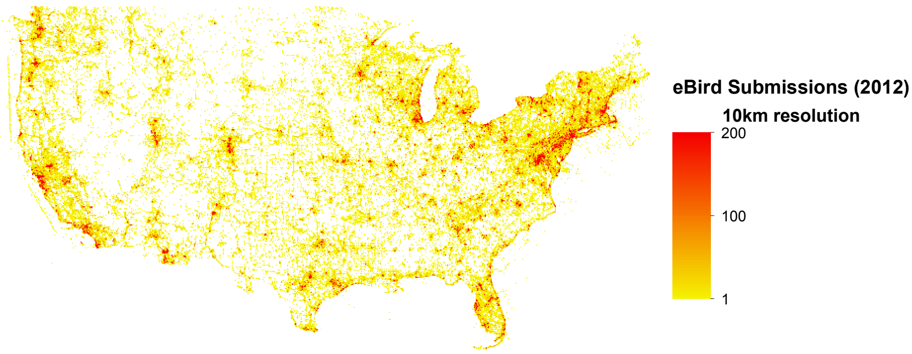

Fig. 1: Number of observations submitted to *eBird* in 2012 in the Continental US. Submissions are biased towards population centers.

Here the leader's utility function U_L is known a priori, but the utilities of the individual followers U_F are unknown by the leader. $\mathbf{a}_1$ and $\mathbf{a}_2$ are the actions of the leader and the followers, respectively.

At the heart of solving this problem lies the challenge of identifying the utility functions that govern the followers' behavior. On one hand, the behavioral model has to be capable of capturing complex, highly variable human behavior and it should be robust to make predictions with limited, sparse data. On the other hand, the behavioral model has to be efficiently incorporated into the overall bi-level optimization problem.

In this paper, we address the behavioral identification problem in two-stage games to reduce data bias in citizen science projects, such as *Zooniverse*, *Coralwatch*, and *eBird* [20,7,29]. These projects use crowdsourcing techniques to engage the public as agents in the data collection process to address scientific questions determined by project leaders. Despite their tremendous success, the data collected often suffer from biases, which arises from fundamental mismatches between the personal motivations that determine how individual agents collect data and the data needs for scientific inquiry. For example, projects that allow participants to choose where and when to make observations tend to collect the most data near areas of human activity, see (see Figure 1). Uneven geographic (and temporal) data density presents a challenge for scientific studies.

Previous work has shown that games are effective in steering citizen scientists towards crucial scientific tasks [31]. Under a two-stage game scenario, individual participants are offered incentives to spend more effort collecting data at sites identified as important by project leaders. In this gamified setting, a key problem is the *optimal reward allocation* problem: how to design a reward scheme which maximizes citizen scientists' overall contribution to science.

The reward allocation problem is closely related to the Principal-Agent Problem, first raised in behavioral economics [26]. More recently, it has also been studied in computer science [1,14,10,3,15]. It is also related to the Stackelberg pricing games [9,8,22,11], in which the leader commits to a strategy before her followers. In crowdsourcing, related work includes mechanisms to improve the crowd performance [23,19,17,28,30,4,27,6,2]. The reward allocation problem is a bi-level optimization that includes as a crucial component the modeling of citizen scientists' behavior.

Here we propose **a novel probabilistic model to capture agents' behavior in two-stage games**, adapting ideas from *Discrete Choice Modeling* in behavior economics [21], as well as **a novel Mixed Integer Programming encoding to solve the reward allocation problem,** in which the proposed probabilistic behavioral model is folded as linear constraints. We also **apply our novel behavioral model into a real citizen science domain**. Our contributions are multi-dimensional:

- On machine learning side, our proposed behavioral model is **(1)** *structural*, meaning that its parameters provide intuitive insights into agents' decision-making process, as well as **(2)** *generative*, meaning that it can generalize to new circumstances with different environmental features and reward treatments. Unlike the knapsack model in previous work [31], our model is **(3)** *probabilistic*. Therefore it is able to account for complex human behavior, as well as suboptimal actions. Instead of directly modeling agents' preferences, which would be difficult to capture, we break the model into **(4)** a *conditional* form, and focus on modeling agents' *deviation* from their baseline behavior under zero reward treatments, alleviating the data sparsity problem by effectively *taking advantage of the relatively abundant historical data before the introduction of the reward game.*
- On the inference side, despite the fact that the reward allocation problem is a bi-level optimization, we are able to **(5)** ***fold the behavioral model into the global problem as a set of linear constraints***, therefore the entire reward allocation problem is solved with a ***single Mixed Integer Program (MIP)***. In addition, we add redundant constraints to trigger pruning, therefore scaling up the MIP encoding to large instances, based on observations of a novel **(6)** ***easy-hard-easy phase transition phenomenon*** **[13]** in the empirical complexity.
- On the application side, we **(7)** apply our behavioral model into a recently launched gamification application called *Avicaching* [31], in the well-established *eBird* citizen science program. Our behavioral model is able to better capture the decision process of the participants than previously proposed models with real field data. Furthermore, the reward designed by the optimal reward allocation algorithm proves to be effective in minimizing the bias in *eBird* data collection process.
- Finally, in terms of addressing *the core scientific goal of ebird*, we show the benefit of having data collected from the *Avicaching* game by demonstrating **(8)** a clear boost in the performance of species distribution modeling when adding data from Avicaching locations.

2 Two Stage Game for Bias Reduction

In our two-stage game setting, citizen scientists visit a set of locations and report their observations of events of interest in those locations. Our model can be generalized to other scientific exploration activities as well [31]. The incentive game involves two self-interested parties: the organizer and the agents. On one side, rational agents (e.g., citizen scientists) select a set of locations to visit that maximizes their own utilities under budgets. On the other side, the organizer (e.g., a citizen science program) uses rewards to encourage agents to visit locations with large scientific value. For example, in *eBird*, bird watchers choose their sites to visit based on a combination of environmental values, personal preference and convenience. The organizer in turn sets external rewards at different locations to promote uniform exploration activities. At a high level, this leads to a bi-level optimization problem:

$$\begin{aligned} \textbf{Organizer:}\ & \underset{\mathbf{r}}{\text{maximize}} && U_o(\mathbf{v}, \mathbf{r}), \\ & \text{subject to} && \textbf{Agents:}\ \mathbf{v} \leftarrow V_a(\mathbf{f}, \mathbf{r}). \end{aligned} \tag{1}$$

In this formulation, $\mathbf{r}$ is the external reward that the organizer uses to steer the agents, and $\mathbf{v}$ are the response from the agents, affected by internal utilities, which is determined by feature vector $\mathbf{f}$, and external rewards $\mathbf{r}$ set by the organizer. $U_o(\mathbf{v}, \mathbf{r})$ is the utility function of the organizer, which depends on agents' response $\mathbf{v}$.

Addressing the Organizer-Agent Problem requires a good behavioral model for agents $V_a(\mathbf{f}, \mathbf{r})$, which involves challenges from two associated problems: one is the *Identification Problem* and the other one is the *Pricing Problem.* For the ***Identification*** Problem, we need to learn an agent model to predict noisy human behavior under different reward treatments. For the ***Pricing*** Problem, we need to incorporate the identified agent model into the bi-level optimization (shown in Equation 1) to solve the overall reward allocation problem.

The organizer's goal is to promote a balanced exploration activity. Let $L = \{l_1, l_2, \ldots, l_n\}$ be the set of locations, and y_i be the amount of effort agents devote to location l_i. We normalize y_i so that $\sum_{i=1}^{n} y_i = 1$. In other words, y_i is proportional to the number of observations submitted at location l_i. Denote by $\mathbf{y}$ the column vector $(y_1, \ldots, y_n)^T$ and by $\overline{\mathbf{y}}$ the constant column vector $(\overline{y}, \ldots, \overline{y})^T$ where $\overline{y} = \frac{1}{n} \sum_{i=1}^{n} y_i = \frac{1}{n}$. To promote a uniform sampling activity, we model the organizer's objective as to minimize the bias in agents' sampling effort: minimize $D_p = \frac{1}{n} ||\mathbf{y} - \overline{\mathbf{y}}||_p^p$. Given this definition, D_1 corresponds to the *mean absolute deviation*, while D_2 corresponds to the *sample variance.* Other objectives could be used, e.g., maximizing the entropy of $\mathbf{y}$ in order to minimize its distance to a uniform distribution.[1]

[1] Uncertainty measures, often used in active learning [25], are typically tied to one particular predictive model. We did not use them because of the need to meet multiple scientific goals in our application.

3 Probabilistic Behavior Model

A key to solving the reward allocation problem is to identify a good behavioral model, which captures agents' preferences to environmental features as well as external rewards. It is challenging, given (1) the *complex and highly variable human behavior*, which cannot be fully captured by environmental variables. Moreover, (2) the data collected with an incentive game in the field is *limited*, since we cannot afford to alienate the community by changing the rewards dramatically. On the other hand, there is much historical data for participants collected without the reward game. How to make full use of this piece of data becomes an interesting question. (3) To efficiently *support decision making*, our behavioral model needs to be able to *fit nicely into the bi-level optimization framework of the pricing problem.* In this paper, we introduce a novel probabilistic model to capture the agents' behavior.

- It takes a **structural** approach, which jointly learns how agents distribute their effort among all locations, rather than predicting the amount of effort spent in each location independently.
- We adopt the idea of **the Discrete Choice Model** in behavioral economics [21], which captures agents' noisy behavior as well as suboptimal actions.
- We alleviate the data sparsity problem by focusing on modeling the conditional probabilities characterizing people's **deviation** from their normal behaviors under no reward treatments, thus effectively taking advantage of relatively abundant historical data without rewards.
- Finally, this structural and generative model allows us to *fold the agents' model as a set of linear constraints* into the reward allocation problem, therefore the entire problem can be solved by **a single MIP**.

During one round of reward treatment, suppose we offer an agent an extra reward r_i for one observation made at location i. Let $\mathbf{r} = (r_1, \ldots, r_n)^T$ be the reward vector. Let $y_{j,i}$ be the amount of effort that agent j devote to location i. We normalize the effort such that $\sum_{i=1}^{n} y_{j,i} = 1$. Let $\mathbf{y}_j = (y_{j,1}, \ldots, y_{j,n})^T$ be vector characterizing the distribution of effort.

Behavioral modeling is to fit a function $\mathbf{y}_j = V_a(\mathbf{f}, \mathbf{r})$ which predicts how agent j distributes his effort $\mathbf{y}_j$ based on environmental features $\mathbf{f}$ and the current reward vector $\mathbf{r}$. One option is to fit V_a as a joint distribution. Unfortunately, this is challenging given the multitude of subtle factors affecting human behavior. Luckily, most participants in our reward game participate heavily in eBird. We have much historical data on them before the reward game, so we hope to use this data to capture their subtle preferences. We therefore break down the agents' behavior into a conditional form, comprising each participant's historical preferences $\mathbf{x}_j$ without external rewards, and the deviation of new behavior $\mathbf{y}_j$ under reward treatment from the baseline behavior $\mathbf{x}_j$. $\mathbf{x}_j$ is summarized based on agents' past behavior during the same time of the year, across previous years. For recently joined participants, we use the population mean as their baseline distribution.

We focus on modeling the *conditional* part, which predicts the *deviation* of people's behavior from $\mathbf{x}_j$ to $\mathbf{y}_j$. Notice that it is a simpler problem than fitting V_a as a joint distribution directly, because *the only main effect that is in the field during the reward treatment period of* $\mathbf{y}_j$*, but not in the baseline treatment period of* $\mathbf{x}_j$*, is the introduction of reward* $\mathbf{r}$. Therefore, the effects of rewards are much stronger in the conditional distribution. We model the transformation matrix P connecting $\mathbf{y}_j$ and $\mathbf{x}_j$, which depends on internal utility features $\mathbf{f}$, and external rewards $\mathbf{r}$:

$$\mathbf{y}_j = P(\mathbf{f}, \mathbf{r})\ \mathbf{x}_j. \tag{2}$$

Many machine learning applications share similar ideas as ours in terms of modeling the conditional part in the joint data distribution [18,12]. Let $p_{u,v}$ be the entry of matrix P at the u-th row and the v-th column. Intuitively, $p_{u,v}$ denotes the proportion of effort that originally was spent in location v, but has been shifted to location u. Motivated by the Discrete Choice Model in behavioral economics [21], we further parameterize the matrix P as:

$$p_{u,v} = \frac{\exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r_u))}{\sum_{u'} \exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u',v}, r_{u'}))}. \tag{3}$$

In this formulation, $\mathbf{f}_{u,v}$ is the environmental feature vector for the transition from location v to location u, which includes features for location u and v individually, such as underlying landscapes, interesting species to see, historical popularities, as well as features that depend on both the two locations, such as the traveling distance, etc. ϕ is a function that maps features to a high dimensional space, which includes singular effect terms as well as cross effect terms. $\mathbf{w}$ is a vector that gives relative weights to different features in the output space of ϕ. The dimensionality of $\mathbf{w}$ is the same as the output of function ϕ.

In previous work [31], agents' behavior are modeled as solving knapsack problems: agents select the best set of locations to visit, which jointly maximizes the reward $\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r_u)$, subject to a cost constraint. Since Equation 3 is a softmax function, our proposed model can be viewed as an extension of the knapsack model to the probabilistic case. Indeed, suppose agents always take the optimal action (as in the knapsack case), their behavior will demonstrate a logit form as shown in Equation 3, if apart from the features in $\phi(\mathbf{f}_{u,v}, r_u)$, their actions are further affected by a set of factors that are only known to agents themselves and with an extreme value distribution [24].

Nevertheless, compared to the knapsack model, our behavioral model is considerably more realistic. Our model is probabilistic, thus it is able to represent variability in agent behavior, as well as uncertainty on the part of the organizer. Suppose one agent chooses to visit either location A or B, but with 70% chance for A, and 30% chance for B. The deterministic knapsack model has to learn a utility function that either predicts that A is a better option than B or vice versa. Our model can come up with an optimal reward scheme in this probabilistic setting. Besides, in the knapsack model, agents' behavior is subject to a strict budget limit. In reality, people occasionally venture beyond their normal travel distance. Our model is able to capture this aspect, by learning a soft penalty on the traveling distance.

3.1 Identification Problem

The identification problem learns the parameters of the agents' behavior model by examining agents' responses to various reward treatments. Specifically, we are given a dataset $\mathcal{D}$ composed of quadruples $(\mathbf{x}_{j,t}, \mathbf{y}_{j,t}, \mathbf{r}_t, \mathbf{f}_t)$, in which $\mathbf{x}_{j,t}$ and $\mathbf{y}_{j,t}$ are the visit densities of one citizen scientist without and with the reward treatment $\mathbf{r}_t$. $\mathbf{f}_t$ is the environmental feature vector during the period of the treatment. We need to identify weights $\mathbf{w}$ that best matches $\mathbf{y}_{j,t}$ with $P(\mathbf{f}_t, \mathbf{r}_t; \mathbf{w})\ \mathbf{x}_{j,t}$. Using the L2 loss, we minimize the following empirical risk function:

$$R(\mathbf{w}) = \sum_{j,t} \left(u_{j,t}(\mathbf{y}_{j,t} - P(\mathbf{f}_t, \mathbf{r}_t; \mathbf{w})\ \mathbf{x}_{j,t})\right)^2 . \tag{4}$$

Here, instances are weighted by $u_{j,t}$, which is the total number of submissions of the corresponding citizen scientist during one reward treatment r_t. We fit a common $\mathbf{w}$ for all citizen scientists, due to limited amount of data.

We specify regularizers to prevent overfitting. It is a common practice to penalize the norm of $\mathbf{w}$ in regularizers. However, when the data is uninformative, a baseline model should always make predictions based on baseline density, i.e., predict $\mathbf{y} = \mathbf{x}$. This suggests that matrix P should be close to the identity matrix in such uninformative case. However, setting $\mathbf{w} = \mathbf{0}$ will make all $p_{u,v} = \frac{1}{n}$ according to Equation 3, which renders P away from the identity matrix. In this case, we add an indicator variable $\mathtt{1}_{u,v}$ as a special feature. $\mathtt{1}_{u,v} = 1$ if and only if $u = v$, and the entries in matrix P becomes:

$$p_{u,v} = \frac{\exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r_u) + \eta \cdot \mathtt{1}_{u,v})}{\sum_{u'} \exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u',v}, r_{u'}) + \eta \cdot \mathtt{1}_{u',v})}. \tag{5}$$

P now becomes close to an identity matrix if $\mathbf{w}$ is close to $\mathbf{0}$ and η is positive. We minimize the following augmented risk function:

$$R(\mathbf{w}) = \sum_{j,t} \left(u_{j,t}(\mathbf{y}_{j,t} - P(\mathbf{f}_t, \mathbf{r}_t; \mathbf{w})\ \mathbf{x}_{j,t})\right)^2 + \lambda \cdot |\mathbf{w}|_1. \tag{6}$$

Here, the classical L1 regularizer $\lambda \cdot |\mathbf{w}|_1$ helps identify important factors by learning a sparse $\mathbf{w}$ vector. Apart from tuning λ, we also tune η to control how closely the predicted $\mathbf{y}$ should match historical densities $\mathbf{x}$. The minimization problem in Equation 6 can be solved by gradient descent. We use BFGS algorithm [5], which further accelerates descent using second order information.

3.2 Pricing Problem

Given a learned behavioral model, the pricing problem is to minimize the spatial bias D_p, subject to the behavioral model:

$$\begin{aligned} \underset{\mathbf{r}}{\text{minimize}} \quad & D_p = \frac{1}{n}||\mathbf{y} - \overline{\mathbf{y}}||_p^p, \\ \text{subject to} \quad & \mathbf{y} = P(\mathbf{f}, \mathbf{r}; \mathbf{w})\ \mathbf{x}, \\ & r_i \in R. \end{aligned} \tag{7}$$

In this formulation, $\mathbf{x} = (x_1, \ldots, x_n)^T$ is the normalized distribution of effort among all agents. Matrix P is learned from the approach given in the previous section. In practice, because people are more accustomed to only a few distinct reward levels, we further restrict r_i to take a set of discrete values in set R.

The main challenge to solve the pricing problem is the sum-exponential form of the entires of matrix P (Equation 5). Nevertheless, in this paper we are able to show that the sum-exponential form can be captured by a set of linear constraints. Therefore *the pricing problem can be formalized as a single Mixed Integer Program (MIP).*

Suppose R has K different reward levels: $R = \{R_1, \ldots, R_K\}$. Introduce indicator variables $dr_{i,k}$ for $i \in \{1, \ldots, n\}$ and $k \in \{1, \ldots, K\}$. $dr_{i,k} = 1$ if and only if r_i, the reward for location i, is R_k. r_i can take only one value in R, so $dr_{i,k}$ should satisfy:

$$\sum_{k=1}^{K} dr_{i,k} = 1, \quad \forall i \in \{1, \ldots, n\}. \tag{8}$$

The challenge is the sum-exponential operator in Equation 5. To overcome this difficulty, we introduce extra variables α_v ($\alpha_v \geq 0$) for $v \in \{1, \ldots, n\}$, and we use linear constraints to enforce

$$\alpha_v = \frac{1}{Z_v} = \frac{1}{\sum_{u'} \exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u',v}, r_{u'}) + \eta \cdot \mathtt{1}_{u',v})}. \tag{9}$$

Here Z_v is the partition function in Equation 5. We first substitute α_v into Equation 5, and get:

$$p_{u,v} = \exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r_u) + \eta \cdot \mathtt{1}_{u,v}) \cdot \alpha_v. \tag{10}$$

However, in this case both r_u and α_v are variables, so Equation 10 is not linear. To linearize it, we rewrite this equation in the following conditional form:

$$\begin{aligned} dr_{u,k} = 1 \ \Rightarrow \ p_{u,v} = \alpha_v \exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, R_k) + \eta \cdot \mathtt{1}_{u,v}), \\ \forall k \in \{1 \ldots K\}, \forall u, v \in \{1 \ldots n\}. \end{aligned} \tag{11}$$

Here, $\mathbf{w}$ is learned from the identification problem, so it is a constant in the pricing problem. When r_u is fixed to R_k ($dr_{u,k} = 1$), $\exp(\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r_u) + \eta \cdot \mathtt{1}_{u,v})$ becomes a constant, so the right-hand side of Equation 11 is indeed a linear equation over α_v. We can enforce the conditional constraints using the big-M formulation. Next, we require the columns of P sum to 1:

$$\sum_{u=1}^{n} p_{u,v} = 1, \quad \forall v \in \{1, \dots, n\}. \tag{12}$$

It can be shown in the following Theorem that Equations 11 and 12 guarantee that $\alpha_v = 1/Z_v$. Further because of Equation 10, we must have the fact that $p_{u,v}$ satisfies the sum-exponential form in Equation 5.

Theorem 1. *Equation 11 and Equation 12 guarantee that* $\alpha_v = 1/Z_v$, $\forall v \in \{1, \dots, n\}$.

Proof. Equation 11 forces α_v to be proportional to $1/Z_v$ and Equation 12 constrains the sum of $p_{u,v}$ to be 1.

Next we model the objective function D_p. Here we provide a formulation for D_1.[2] The key is to model the absolute difference $|y_i - \overline{y}|$. Introduce variable t_i for $|y_i - \overline{y}|$, $i \in \{1, \dots, n\}$, and constraints $t_i \geq y_i - \overline{y}$ and $t_i \geq \overline{y} - y_i$ to guarantee that $t_i \geq |y_i - \overline{y}|$. Then we can modify the objective so as to minimize $\sum_{i=1}^{n} t_i$.

In practice, we find the MIP encoding with the constraints in Equations 8-12 does not scale well with small external rewards (see section 4.3). In this case, we add redundant constraints to facilitate constraint propagation and pruning. When $r_u = R_k$, we add these redundant constraints:

$$p_{u,v} \leq \frac{\exp(g_{u,v}(R_k))}{\exp(g_{u,v}(R_k)) + \sum_{u' \neq u} \exp(\min_{r \in R} g_{u',v}(r))}, \tag{13}$$

and

$$p_{u,v} \geq \frac{\exp(g_{u,v}(R_k))}{\exp(g_{u,v}(R_k)) + \sum_{u' \neq u} \exp(\max_{r \in R} g_{u',v}(r))}. \tag{14}$$

Here, $g_{u,v}(r)$ is an abbreviation for $\mathbf{w} \cdot \phi(\mathbf{f}_{u,v}, r) + \eta \cdot \mathtt{1}_{u,v}$. The right hand side of these two inequalities are clearly the upper and lower bound of $p_{u,v}$, because all free variables are fixed to their most extreme values.

4 Experiments

4.1 Applying the Behavioral Model to *Avicaching*

We apply our behavioral model into ***Avicaching*** [31], a recently launched gamified application to reduce the data bias problem within *eBird*, a well-established citizen science program. *Avicaching* is created in the spirit of promoting "friendly competition and cooperation" among *eBird* participants. *Avicaching* started in March 2015 as a pilot study in Tompkins and Cortland counties, New York. A set of publicly accessible locations with no prior eBird observations were defined

[2] One needs solve a Mixed Quadratic Program if he uses objective function D_2.

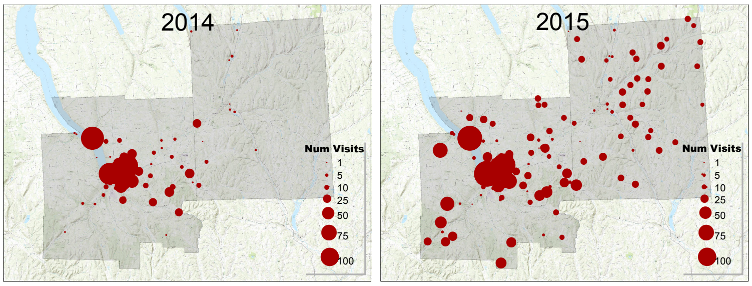

Fig. 2: The comparison of the locations of submissions in *eBird* in Tompkins and Cortland County in New York State. The size of the circles represent the number of submissions. (Left) from Mar 28 to Aug 31, 2014, before *Avicaching*. (Right) from Mar 28 to Aug 31, 2015, after *Avicaching* is introduced. Effort is shifted towards under-sampled locations significantly. Study area is shaded.

as *Avicaching* locations: bird watchers received extra *avicaching* points for every checklist they submitted in those locations. These locations were selected around under-covered regions from the current *eBird* dataset, emphasizing important yet under-sampled land types, such as agricultural land and forest. *Avicaching* points have intrinsic value to bird watchers, because they mark their scientific contribution to *eBird*. In addition, other rewards, such as binoculars, were also provided in the form of a lottery, which is based on the total *avicaching* points earned by each participant. The *Avicaching* points were updated every week. The probabilistic behavioral model was used in the bi-level optimization problem, which allocates optimal rewards to locations to minimize the spatial bias. We used the participants' response in the first few weeks to train our behavioral model.

Encouraged by *Avicaching*, bird watchers shifted their effort towards under-sampled locations. As visually demonstrated in Figure 2, 482 *eBird* observations were submitted from *Avicaching* locations, out of the 2,522 observations in total for Tompkins and Cortland County during summer months from June 15 to Sep 15, 2015. 19.1% of birding effort has shifted from oversampled locations to under-sampled *Avicaching* locations, which received zero submissions before. Cortland, as an under-sampled county, received 202 observations during these three summer months in 2015, when *Avicaching* is in the field, which is 2.3 times the number of visits of the previous two years combined (there are in total 87 submissions from Cortland during the same period of time in 2013 and 2014). In terms of uniformity, the normalized D_2 score ($\frac{1}{n}||\mathbf{Y} - \overline{\mathbf{Y}}||_2^2/\overline{Y}$), dropped from 0.017 in 2013, 0.017 in 2014 to 0.013 in 2015 during the period of time.

4.2 Evaluation of the Probabilistic Behavioral Model for the Identification Problem

The behavioral model used in the reward allocation problem of one week is fit using the data since the beginning of *Avicaching* and up to that week. The data

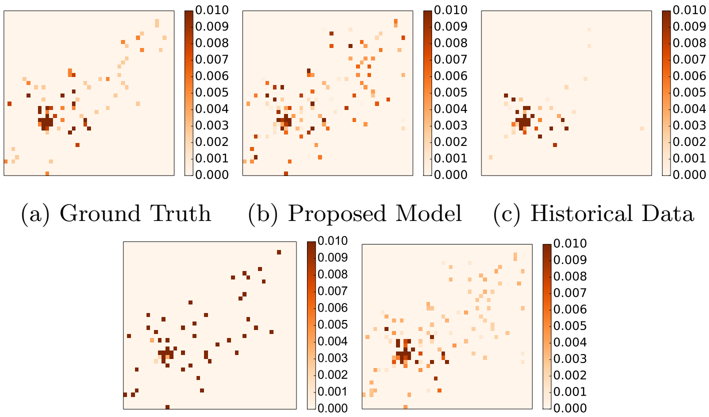

(a) Ground Truth (b) Proposed Model (c) Historical Data

(d) Structural SVM (e) Random Forest

Fig. 3: The comparison of probabilities of visiting each location predicted by various behavioral model on one test set. The range was selected to highlight locations with small probabilities. The proposed model matches closest to the ground truth (note the color scale).

Method	Normalized MSE
Proposed	**0.26**
Historical	0.36
Structural SVM	0.93
Random Forest	0.37

Table 1: Comparison of predicted performance on the test set. The table shows the normalized mean squared error (MSE). Our proposed model outperforms the other 3 baseline models.

are composed of $(\mathbf{x}_{j,t}, \mathbf{y}_{j,t}, \mathbf{r}_t, \mathbf{f}_t)$ tuples, each of which represents the density of locations a bird watcher visited during one week's reward treatment. There are in total 116 locations in total in this two counties (the length of $\mathbf{x}_{j,t}$ and $\mathbf{y}_{j,t}$), out of which 50 are *Avicaching* locations. We split the dataset into 75% for training, 5% for validation, and the remaining 20% for testing. The data for validation is used to select the values of regularizers. We found the model is not sensitive to the values of regularizers, as long as they are in a proper range. The reported performance is averaged over 3 random splits. The location features we consider for the behavioral model are: the number of visits in each month (popularity), the expected number of species to see (interestingness), the NLCD covariates for the landscape [16], housing density (population center), elevation, distances to rivers, roads, etc, latitude and longitude (geographical regions), convenience factor (distance to reach), and *Avicaching* points (rewards). We also include non-linear transformation of these features and cross terms.

We compare our proposed model with three baseline models. The first model always uses historical density to make predictions, i.e., always predict $\mathbf{y}_{j,t} = \mathbf{x}_{j,t}$. The second model is the structural SVM model from [31], a powerful nonparametric machine learning model optimized for solving knapsack problems. The third benchmark is a continuous-response random forest, which predicts the density $y_{j,t}$ at each location independently with 1,000 trees of depth 10. Random forests are expected to set the benchmark for very good predictive performance. However, the lack of interpretable structures precludes them from being folded into the MIP formulation of the pricing problem. Both the Structural SVM and the random forest model share the same environmental features as our proposed model. We include the baseline density $\mathbf{x}_{j,t}$ in the two models as a separate feature.

Table 1 shows the comparison on normalized mean squared error (MSE), which is $\frac{\sum_{j,t}\sum_{i=1}^{n}(u_{j,t}(y_{j,t,i}^{truth}-y_{j,t,i}^{pred}))^2}{\sum_{j,t}\sum_{i=1}^{n}(u_{j,t}(y_{j,t,i}^{truth}-\overline{y_{j,t,i}^{truth}}))^2}$. Here $y_{j,t,i}^{truth}$ is the true density for agent j in time t, and $y_{j,t,i}^{pred}$ is the predicted value at location i. The squared error is further weighted by $u_{j,t}$ – the number of submissions during the reward period. **Our proposed model clearly outperforms the other 3 models.** To further visualize the difference, the predicted probabilities to visit each location, averaged over all test cases in one test set, are compared with the ground truth in Figure 3. The locations with high probabilities (shown with dark red cells) are historically popular sites. The model based on historical density predicted very well on these sites, because we have rich data on people's birding history, and bird watchers' behavior is relatively stable across different years. Those sites with relative low probabilities (light orange cells) are often under-sampled sites with *Avicaching* rewards. In this case, the historical model missed completely. Structural SVM performed the worst. While random forest performed well qualitatively, it was out-performed by our proposed model (Table 1). Even if the random forest model had comparable performance, it cannot be folded into the MIP to solve the bi-level optimization problem.

4.3 Phase Transition on the Pricing Problem

The scalability of the Mixed Integer Programming encoding proposed for the pricing problem is also important. To evaluate the solver, we generated 5 sets of synthetic instances, with numbers of locations n ranging from 15 to 35. Each set had 30 instances with the same n, generated in a way to best mimic people's behavior. To make it easy for plotting, reward set R contains 2 levels of rewards for these instances: one was 0, and the other was a non-zero reward shown in the horizontal axis of Figure 4 (all 30 instances in one test set shared a common non-zero reward). We kept all other parameters the same, and only varied the non-zero rewards. The curves in Figure 4 report the median time to solve these instances with MIP encoded in CPLEX 12.6, with a single Intel x5690 core and 8GB of memory. Each dot in one curve represents the median time of solving 30 instances in one test set. Two points on a given curve only differ in the reward level.

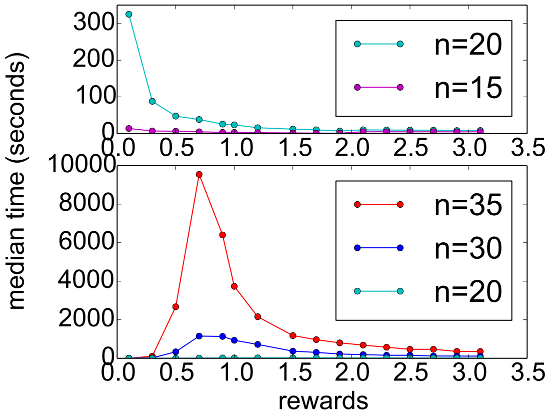

Fig. 4: The easy-hard-easy phase transition for the pricing problem; n is the number of locations. (Upper) The median time to solve instances with various non-zero rewards without the redundant constraints in Eq. 13 and Eq. 14. The time is long for instances with small rewards. (Lower) The median time when redundant constraints are introduced. The easy-hard-easy pattern emerges.

Intuitively, there should be an easy-hard-easy pattern in the empirical complexity of the pricing problem. If the external rewards are too small, then it makes little difference in terms of changing agents' behavior whether one reward is assigned to one location or not. On the other hand, if the rewards are too large, then agents' behavior is completely dominated by these external rewards. It is when the external rewards match agents' internal utilities that the problem becomes hard, and the algorithm needs to plan wisely in allocating rewards. Nevertheless, when the redundant constraints in Equations 13 and 14 were not introduced (Upper Panel of Figure 4), we did not see the easy-hard-easy pattern. Problem instances with small non-zero external rewards were significantly harder than other ones.

The unexpected long runtimes for instances with small rewards were due to the difficulty in propagating constraints. The solver could not automatically discover the fact that the reward was too small to have any substantial impact, so it spends much time on many meaningless branches. This prevented the solver from early pruning, which was often the key to efficient problem solving. Noticing this aspect, we added redundant constraints (Equation 13 and Equation 14) into the MIP formulation. These two equations were obvious necessary conditions for $p_{u,v}$. Adding these two equations helped the solver find bounds on $p_{u,v}$, so it could prove tighter bounds for the objective function, and trigger early pruning more often. After adding these two constraints, the easy-hard-easy phenomenon emerged. We were also able to scale up to larger instances due to better constraint propagation (It takes too long for the solver to run for $n = 30$ and $n = 35$ without additional constraints, so they are not plotted in the upper panel of Figure 4).

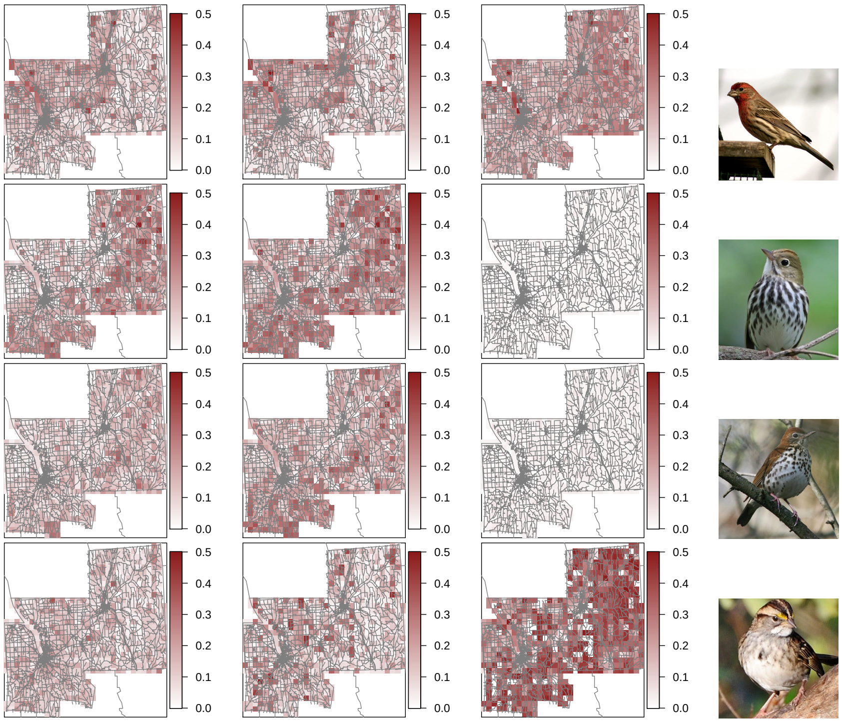

(a) Ground Truth (b) with Avicache (c) without Avicache (d) bird

Fig. 5: The benefit of having observations from avicaching sites. (1st Row) Model for House Finch; (2nd Row) Ovenbird; (3rd Row) Wood Thrush; (4th Row) White-throated Sparrow. Predictive model fit with 2015 data including that from Avicaching sites in Cortland (2rd column) better matches a model close to the ground truth (1st column, fit with all available data, best effort and validated by experts), compared with the model fit without Avicaching data (3rd column).

4.4 Benefit of Avicaching on Species Modeling

We are able to see the benefit of having data collected from avicaching locations on species distribution modeling – the main scientific application of *eBird* data. To fit the species distribution models, we use the data from April to June (the spring migration period), in both Tompkins and Cortland counties, including avicaching and non-avicaching locations. We predict the occupancy of a species based on environmental variables. For each species, we fit random forest models with 1,000 trees, with each tree at the depth of 10.

Figure 5 shows the predicted probabilities of occurrence in heatmaps for random forest models fit with different datasets, for four species in the two counties. The first column shows the distribution models fit with the most comprehensive dataset, which consists of data from both counties, during April to June across several years. Because Tompkins county is the best covered area in *eBird*, the learned model is close to the ground truth, according to bird experts at the Cornell Lab of Ornithology. In the second column, we fit the models using the data only from Cortland County in 2015, including that from avicaching locations. We use Cortland County as an example to represent a large number of counties in the United States, where there are few *eBird* submissions. Then in the third column, we further exclude the data collected from Avicaching locations.

As we can see from Figure 5, the species distribution models in the second column match pretty well in terms of the predicted probabilities with the models in the first column, although they are fitted using much less data. On the contrary, the models in the third column are much worse. Indeed, the log losses improve from 0.44 to 0.30 for Ovenbird, from 0.47 to 0.46 for House Finch, from 0.51 to 0.38 for Wood Thrush and from 0.48 to 0.41 for White-throated sparrow when Avicaching observations are added.

Since the only difference between the models in the second and the third columns is whether the models are learned using the dataset containing observations from *Avicaching* locations, the clear difference in the predictive performance demonstrates the benefit of having data from *Avicaching* locations. From this experiment, we see that *Avicaching* game really helps *eBird* in addressing its ultimate scientific goal.

5 Conclusion

We address the behavior identification problem in two-stage games to reduce the data bias problem in citizen science. We introduce a novel probabilistic behavioral model and show that it is better at capturing noisy human behavior compared to the knapsack model previously used in *Avicaching*, a recently launched gamified application in *eBird*. In addition, the behavioral model can be folded as a set of linear constraints into the bi-level optimization problem for bias reduction, so the whole two-stage game can be solved with a single Mixed Integer Program. We further scale up the encoding to large instances by adding redundant constraints, based on a novel easy-hard-easy phase transition phenomenon. Finally, we also show that the data collected from the *Avicaching* game improves species distribution modeling, therefore it better serves the core scientific goal of citizen science.

Acknowledgements

We are thankful to the anonymous reviewers for comments, thousands of *eBird* participants, and the Cornell Lab of Ornithology for managing the database. This research was supported by National Science Foundation (0832782, 1522054, 1059284, 1356308), ARO grant W911-NF-14-1-0498, the Leon Levy Foundation and the Wolf Creek Foundation.

References

1. Aggarwal, G., Feder, T., Motwani, R., Zhu, A.: Algorithms for multi-product pricing. In: ICALP (2004)
2. Anderson, A., Huttenlocher, D.P., Kleinberg, J.M., Leskovec, J.: Steering user behavior with badges. In: 22nd Int'l World Wide Web Conference, WWW (2013)
3. Bacon, D.F., Parkes, D.C., Chen, Y., Rao, M., Kash, I., Sridharan, M.: Predicting your own effort. In: Proc. of the 11th Int'l Conf. on Autonomous Agents and Multiagent Systems-Volume 2. pp. 695–702 (2012)
4. Bragg, J., Mausam, Weld, D.S.: Crowdsourcing multi-label classification for taxonomy creation. In: HCOMP (2013)
5. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. SIAM J. Sci. Comput. (1995)
6. Chen, X., Lin, Q., Zhou, D.: Optimistic knowledge gradient policy for optimal budget allocation in crowdsourcing. In: ICML (2013)
7. Chiappone, M.: Coral watch program summary. a report on volunteer and scientific efforts to document the status of reefs in the florida keys national marine sanctuary. The Nature Conservancy, Summerland Key, Florida (1996)
8. Conitzer, V., Garera, N.: Learning algorithms for online principal-agent problems (and selling goods online). In: Proc. of the 23rd ICML (2006)
9. Conitzer, V., Sandholm, T.: Computing the optimal strategy to commit to. In: Proc. 7th ACM Conference on Electronic Commerce (EC). pp. 82–90 (2006)
10. Endriss, U., Kraus, S., Lang, J., Wooldridge, M.: Incentive engineering for boolean games. IJCAI Proceedings-International Joint Conference on Artificial Intelligence 22(3), 2602 (2011)
11. Fang, F., Stone, P., Tambe, M.: When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In: IJCAI (2015)
12. Gens, R., Domingos, P.M.: Discriminative learning of sum-product networks. In: Advances in Neural Information Processing Systems. pp. 3248–3256 (2012)
13. Gomes, C.P., Selman, B.: Satisfied with physics. Science 297(5582), 784–5 (2002)
14. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: SODA. pp. 1164–1173 (2005)
15. Hartline, J.D., Koltun, V.: Near-optimal pricing in near-linear time. In: Algorithms and Data Structures, 9th Int'l Workshop, WADS. pp. 422–431 (2005)
16. Homer, C., Dewitz, J., Fry, J., Coan, M., Hossain, N., Larson, C., Herold, N., Mckerrow, A., Vandriel, J.N., Wickham, J.: Completion of the 2001 National Land Cover Database for the conterminous United States. Photogrammetric Engineering and Remote Sensing 73(4), 337–341 (2007), http://www.asprs.org/publications/pers/2007journal/april/highlight.pdf
17. Kawajiri, R., Shimosaka, M., Kashima, H.: Steered crowdsensing: Incentive design towards quality-oriented place-centric crowdsensing. In: UbiComp (2014)
18. Lafferty, J.D., McCallum, A., Pereira, F.C.N.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning. ICML (2001)
19. Li, H., Tian, F., Chen, W., Qin, T., Ma, Z., Liu, T.: Generalization analysis for game-theoretic machine learning. In: AAAI (2015)
20. Lintott, C.J., Schawinski, K., Slosar, A., et al.: Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. Monthly Notices of the Royal Astronomical Society 389(3), 1179–1189 (Sep 2008), http://dx.doi.org/10.1111/j.1365-2966.2008.13689.x

21. McFadden, D.: Modeling the choice of residential location. Spatial Interaction Theory and Residential Location pp. 75–96 (1978)
22. Paruchuri, P., Pearce, J.P., Marecki, J., Tambe, M., Ordóñez, F., Kraus, S.: Playing games for security: an efficient exact algorithm for solving bayesian stackelberg games. In: AAMAS. pp. 895–902 (2008)
23. Radanovic, G., Faltings, B.: Incentive schemes for participatory sensing. In: AAMAS (2015)
24. Rust, J.: Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica 55(5), 999–1033 (1987)
25. Settles, B.: Active learning literature survey. University of Wisconsin, Madison 52(55-66), 11 (2010)
26. Shavell, S.: Risk sharing and incentives in the principal and agent relationship. The Bell Journal of Economics (1979)
27. Singer, Y., Mittal, M.: Pricing mechanisms for crowdsourcing markets. In: Proceedings of the 22nd Internat. Conf. on World Wide Web (WWW) (2013)
28. Singla, A., Santoni, M., Bartók, G., Mukerji, P., Meenen, M., Krause, A.: Incentivizing users for balancing bike sharing systems. In: AAAI (2015)
29. Sullivan, B.L., Aycrigg, J.L., Barry, J.H., others: The ebird enterprise: An integrated approach to development and application of citizen science. Biological Conservation 169(0), 31 – 40 (2014), http://www.sciencedirect.com/science/article/pii/S0006320713003820
30. Tran-Thanh, L., Huynh, T.D., Rosenfeld, A., Ramchurn, S.D., Jennings, N.R.: Crowdsourcing complex workflows under budget constraints. In: Proceedings of the AAAI Conference. AAAI (2015)
31. Xue, Y., Davies, I., Fink, D., Wood, C., Gomes, C.P.: Avicaching: A two stage game for bias reduction in citizen science. In: Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems. AAMAS (2016)