Census Worker Allocation under Uncertainty

Cem Gokmen
Department of Computer Science
Stanford University
cgokmen@stanford.edu

Lynn Dezhen Kong
Department of Computer Science
Stanford University
ldkong@stanford.edu

Connor Toups
Department of Computer Science
Stanford University
ctoups22@stanford.edu

Abstract

The decennial United States census is an undertaking of massive importance and immense difficulty. The census bureau must efficiently allocate workers to count households that have not self-responded to the census. Because households can continue to self respond as workers are being allocated, there is uncertainty not only in deciding which census areas to allocate workers to but when to allocate these workers so as to optimize the total response rate of the census by the end of census-taking period. We formulate a Markov Decision Process to demonstrate how MDPs can effectively allocate workers under this uncertainty, and we incorporate both real data from the 2020 census as well as synthetic data (where real data does not exist publicly) to model our environment. Through this model, we demonstrate how real data privately available to the census bureau can easily replace our synthetic data and how policy-makers can leverage this model to allocate census resources for fairer representation and counting of historically under-counted groups.

1 Introduction

In this paper, we explore the modeling of the United States census as a Markov Decision Process. We model the census as a Markov Decision Process where districts (e.g. finite geographic regions with a population, such as a state, county, or census tract) each accrue response rates (*states*) due to citizens' self-responses with some uncertainty. Additionally, in-person follow-ups to districts (*actions*) can also increase the response rate with some uncertainty. To maximize total response rate across all districts, the census bureau must decide where and when census workers should make in-person follow-ups. A policy can thus be produced that optimizes some reward function using census worker follow-ups throughout the census duration.

The primary goal of this project is to investigate the effects of different approaches of computing the value of different census visit policies. We specifically identify reward schemes that can encode the value of fair representation and counting of historically under-counted minority groups – a significant issue given the Census's role in allocating federal resources to communities.

2 Background

The results of each decennial U.S. census are used to apportion the 435 voting representatives in the House of Representatives and distribute more than $675 billion per year in federal funding to communities [2]. The Census's role in the apportionment of political power and federal aid makes the accuracy of the Census of critical importance. Congress allocated $15 billion to conduct the 2020 census [6][7]. Career workers at the Census Bureau are forced to operate under spending constraints, and we motivate our work by imagining the position they are in as agents and policy-makers.

The Census's goal is to count every person in the United States, exactly once, and count them in the correct location. To count households, the Census relies heavily on self-response wherein households respond to census questionnaires by mail, phone, or – new to the 2020 census – online. Self-response rates, however, don't come anywhere near 100%, and they have continuously dropped over the past half-century [8][4]. Furthermore, these self-response rates vary across demographic groups. Historically, census tracts with high populations of Latinx, Black, and Indigenous people have had low response rates – and therefore were potentially under-counted in the census [5]. In order to count people who do not self-respond to the census, the Census Bureau sends workers to make in-person contact with households. These workers are distributed to census tracts across the country with the goal of getting the total census response-rate to 100% [1].

While the problem of census worker allocation to maximize response rate can be framed as a Markov Decision Process, this framing is rarely explored in research literature. Thus, we built the MDP from scratch based on our understanding of the census response process. Furthermore, we were inspired by the paper by *Cai et al.* on fairness in resource allocation thorough selective information acquisition. They looked at this problem in the context of loans using linear programming; we chose to explore fair representation in the context of the census using MDPs [3].

3 Approach: The Markov Decision Process

3.1 Environment Definitions

Given n districts, we denote a district as $i \in \{1, ..., n\}$. Allowing up to t_{end} time steps, we denote $T = \{1, ..., t_{end}\}$ as the time step space. We also define $P = [0.0, 1.0] \subset \mathbb{R}$ as the response rate space for each district.

3.2 State Space

The MDP state $\mathbf{s} = (t, \mathbf{p}) \in (T \times P^n)$ contains the current time in t and the response rate $p_i \in P$ of each district in vector $\mathbf{p}$. Note that the response rates, and thus the state, is continuous.

3.3 Action Space

An action $a \in \{0, ..., n\}$ represents which district to send census workers to (for the sake of simplicity, only one visit per time step is allowed, but results apply without loss of generality to visiting multiple districts at once, since this can be approximated by increasing time step granularity.) If $a = 0$, then no districts are visited.

3.4 Transition Model

There are two ways that response rates can increase in our model: they can increase naturally – simulating the natural increase in response rate that comes from self-response – or they can increase as a result of workers being sent to a district.

On every time step t, we first apply the natural self-response increase to the response rate of every district i regardless of the action taken. This is done by sampling an increase from each district's individual self-response increase distribution, which is learned from the data.

Furthermore, we apply an additional increase to the response rate of the district i corresponding to the action a, as the result of census workers visiting the district. This increase is sampled from a distribution that is shared between all districts, that is, we assume that a census worker visit has the same impact on each district for the sake of simplicity. This could be extended in future work to include the effectiveness of visits in different districts.

Finally, we advance the time state by one step.

3.5 Reward Function and Fair Representation

The reward function is designed to reflect the three criteria that are critical to the census worker allocation: the cost of logistics of sending workers to a district, the utility of the total response rate, and the utility of the representation of population demographics.

At each time step t, to represent cost of sending workers to a district, we give a negative reward r_{cost} if the chosen action is to allocate worker to a district, i.e. $a \neq 0$. If the chosen action is to stay put, $a = 0$, then no negative reward is generated. In our simulation, r_{cost} is constant regardless of chosen district. In future work, it might be worth exploring r_{cost} as a function of the chosen district to better represent the variable cost of visiting districts with different demographics.

Furthermore, at the last time step, we give a positive reward r_{resp} based on the total response rate from all districts. We chose to do so instead of calculating positive reward during the running steps so to encourage exploration and avoid greedy policies that are short-sighted. After all, the total response rate and demographic outcomes only matter once the census reaches a deadline.

$$R(\mathbf{s} = (t, \mathbf{p}), a) = r_{cost} \times \mathbb{1}(a \neq 0) + \begin{cases} 0 & t \neq t_{end} \\ r_{resp}(\mathbf{p}) & t = t_{end} \end{cases}$$

We propose two variants of the reward function based on how we calculate the positive reward. In the population-based variant, the positive reward is the total number of visited households. We use the population count of each district sourced from the 2019 estimate by the census bureau, to calculate the total overall response rate.

$$r_{resp}(\mathbf{p}) = \sum_{i=1}^{n} p_i \cdot Population(i)$$

In the second variant, we attempt to reward fair representation as well as total response rate. The goal of this variant is to penalize situations where the sample's demographics deviates from the population's demographics. To do this, we want to compute the population and sample means of certain demographic statistics.

We represent the demographic information of all districts in matrix $D \in \mathbb{R}^{n \times m} = [\mathbf{d_1}, ..., \mathbf{d_n}]$, where $\mathbf{d_i}$ is the vector of demographic information for district i, containing m elements each corresponding to one numerical demographic value such as median income, percent ethnic distribution, and language spoken. This data is obtained from the previous census.

To compute the population mean $\mathbf{d_{pop}}$ of the demographics, we compute the average of the districts' demographics weighted by their populations:

$$\mathbf{d_{pop}} = \frac{\sum_{i=1}^{n} Population(i) \cdot \mathbf{d_i}}{\sum_{i=1}^{n} Population(i)}$$

To compute the sample mean $\mathbf{d_{sample}}$ of the demographics, we compute the average of the districts' demographics weighted by the population we have counted at that district:

$$\mathbf{d_{sample}} = \frac{\sum_{i=1}^{n} p_i \cdot Population(i) \cdot \mathbf{d_i}}{\sum_{i=1}^{n} p_i \cdot Population(i)}$$

Finally, we compute the positive reward by applying a hand-curated fixed weight θ_j to the ratio of the sample mean to the population mean of each demographic j, e.g. how our current sample's demographics compare to those of the population:

$$r_{resp}(\mathbf{p}) = \sum_{j=1}^{m} \theta_j \cdot \left| 1 - \frac{d_{sample_j}}{d_{pop_j}} \right|$$

3.6 Learning Self-Response Distributions

To simulate the response rate increases from self-response, we pulled self response rate data from the 2020 census starting from March 23rd and continuing up to October 17 (see **Data Source**). While the data we pulled was for each census tract, we decided – for the sake of computational efficiency – to let each district in our MDP correspond to a single U.S. State. Additionally, we chose to let each time step represent one week. We averaged cumulative response rates for each U.S. State at every week starting March 23rd. We used the cumulative self-response rates to calculate the weekly percent increase in response rates.

We initially fit a single normal distribution to these increases for each state; however, we found that self-response rate increases were drastically higher in the first 7 weeks than in the last 21 weeks. To compensate for this bimodal behavior, we chose to fit two separate truncated normal distributions; one to the first 7 weeks of response rate increases, and one to the subsequent 21 weeks. We bound the random variable between 0 and whatever value increase would bring the response rate to 100%; we then sample from this truncated normal distribution. Visualizations for these distributions on a single state are included in the appendix (Figure 6).

3.7 Data Source

Our paper uses a combination of real data from the 2020 Census and synthetic data that we have created to simulate potential real world data. The need for synthetic data arises from a lack of publicly available data from the Census Bureau; however, we believe that this data is accessible internally within the bureau and therefore we demonstrate *how* real data could be used in our model by utilizing synthetic data.

We use real data to approximate the natural increase in response rates via self-response. The real data is pulled from this repo: `https://github.com/stuartlynn/census_2020_response_rates`. The repository used the Census Response Rate Data API to pull self-response rates by census tract for each day since March 22.

We assume arbitrary parameters for the normal distribution that represents the change in response rate produced by sending workers to a district since this data is not publicly available.

4 Results

We have implemented the above MDP using the *POMDPs.jl* Julia library, with source code available at `https://github.com/Census-MDP/Census-MDP`. We then produced policies using the Monte Carlo Tree Search solver with 10,000 iterations and sufficient rollout depth to reach the end time from the start. For all of the below runs, we divide each timestep in the original data (28 total) into two timesteps to increase the number of possible visits.

4.1 Self-Response Only

For our initial run, we introduce a negative-infinity reward for any action $a \neq 0$, that is, infinitely penalizing district visits, to produce a baseline of what the census outcome looks like without visits. Multiple runs with these parameters produced response rate averages of around 75%.

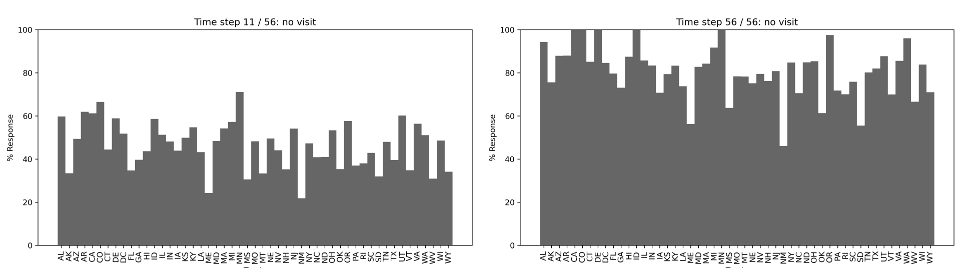

(a) Districts' Response Rates at $t = 11$ (b) Districts' Response Rates at $t = t_{end} = 56$

Figure 1: Results on MDP with visits discouraged

4.2 Rewarding Sample Population

For our next run, we introduce a small negative reward for any action $a \neq 0$, to penalize truly unnecessary visits (e.g. to districts already at 100%), and we use the population reward function defined earlier. Multiple runs with these parameters produced response rate averages of around 85%. More importantly, high response from more populated states such as CA, NY and TX is favored over less populated states, which is not necessarily desireable. The visited state is showed in black in the chart.

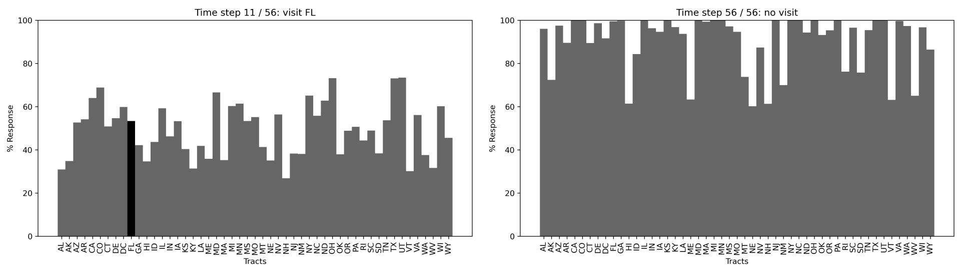

(a) Districts' Response Rates at $t = 11$ (b) Districts' Response Rates at $t = t_{end} = 56$

Figure 2: Results on MDP with population-based rewards

4.3 Rewarding Sample Demographics' Similarity to Population Demographics

For our final run, we introduce a small negative reward for any action $a \neq 0$, to penalize truly unnecessary visits (e.g. to districts already at 100%), and we use the demographic-based reward function defined earlier. Multiple runs with these parameters also produced response rate averages of around 85%. However, unlike the previous model, high response from of more populated states such as CA, NY and TX is no longer favored over less populated states, since doing this causes a skew towards those states' demographic means. The new reward function reduces the variance of the states' response rates and produces a sample that better approximates their demographics, as shown in the next section. The visited state is showed in black in the chart.

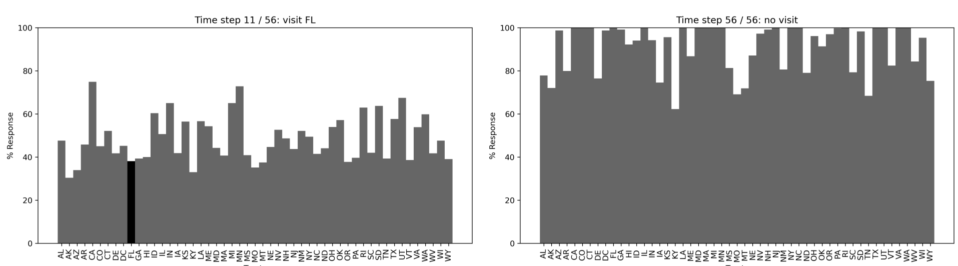

(a) Districts' Response Rates at $t = 11$ (b) Districts' Response Rates at $t = t_{end} = 56$

Figure 3: Results on MDP with demographic-based rewards

4.4 Comparing Population-Based Reward to Demographics-Based Reward

It was already shown that the two reward methods produce similar average response rates across the states. However, the difference between the two approaches can be seen in the below figure comparing each demographic statistic's population and sample means for both methods used:

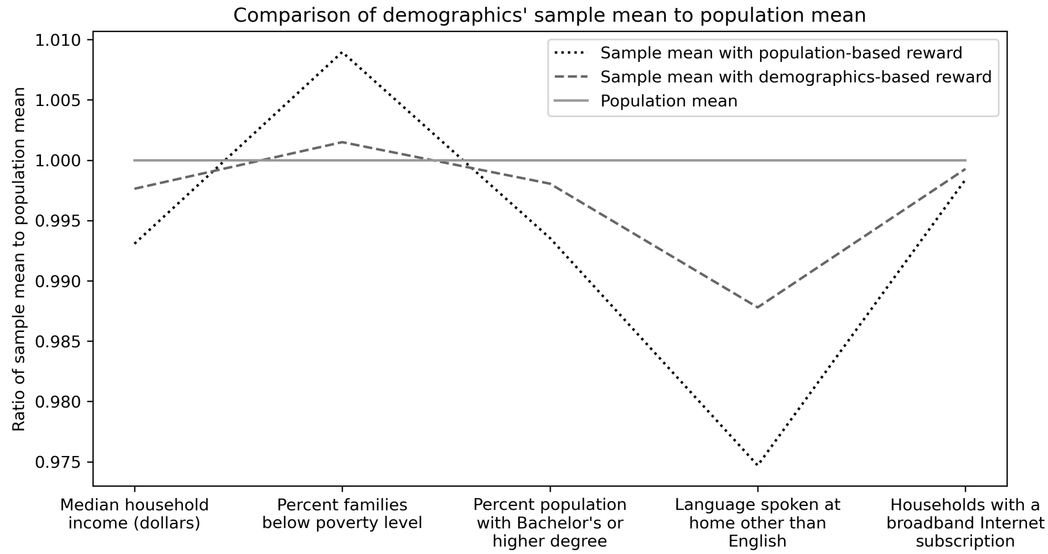

Figure 4: Comparison of effect of reward functions on sample demographics

It is clear that the demographics-based reward, without necessarily impacting the total number of households visited, is able to produce a sample that better approximates (e.g. deviates less from) the demographics of the population than the population-based reward is. With further tuning of the weights, this effect can be further improved.

5 Conclusion and Future Work

We have defined census worker allocation as a Markov Decision Process, with a visit to at most one possible district per time step. Then, we have defined our generative transition function based on distributions of response rate increase that we learned from existing census response data. Finally, we have explored policies learned from both a population-based reward function and a demographic-based reward function, and showed how the demographic-based reward function can be used to produce a sample that better approximates the population's demographics, which we hope can be used to produce more representative census results.

Our MDP setup has certain limitations. We chose to limit the number of districts we can visit per time step to 1. If we allow multiple district visits per time step, our action space would increase from $|\mathcal{A}| = n$ to an upper bound of $|\mathcal{A}| = O(2^n)$. The high memory usage and processing time with the expanded action space means this approach was beyond the scope of this project. However, multiple district visits definitely more accurately reflect the reality of census follow-up, and is worth exploring in the future with a single-agent MDP and possibly multi-agent MDPs.

We set several constants in our reward functions. We assumed a static cost for each district visit. In the demographic-based reward function, we hand-curated the weights for each demographic type. In the future, we can instead learn these parameters by understanding the value of each demographic using methods such as Utility Elicitation.

6 Contributions

Cem Gokmen implemented the MDP and solver in Julia, executed multiple models, and contributed the approach and result sections in this report. Lynn Kong built the graph visualization and wrote the conclusion section and parts of the introduction and approach sections of this report. Connor Toups parsed the census data into distributions to be used by the MDP and wrote the abstract, background and parts of the introduction and approach section. All three team members met regularly, contributed to the MDP conceptualization, and edited the report.

References

[1] U.S. Census Bureau. 2020 census paid temporary workers, 2020.

[2] U.S. Census Bureau. Why we conduct the decennial census, Apr 2020.

[3] William Cai, Johann Gaebler, Nikhil Garg, and Sharad Goel. Fair allocation through selective information acquisition. *arXiv.org*, Sep 2020.

[4] Dartunorro Clark. The states with the best and worst census response rates, Sep 2020.

[5] Mark Mather, Linda A Jacobson, Alicia Vanorman, and Matthew Martinez. 2020 census self-response rates are lagging in neighborhoods at risk of undercounting young children, Mar 2020.

[6] Jeffrey Mervis. Money, politics, and abandoned homes: Why the 2020 census might be in jeopardy, Dec 2017.

[7] U.S. Government Accountability Office. Substantial efforts needed to achieve greater progress on high-risk areas, Mar 2019.

[8] Herbert F Stackhouse and Sarah Brady. Census 2000 mail response rates, Jan 2003.

7 Appendix

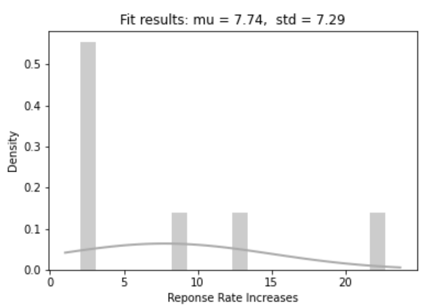

(a) First 7 Weeks

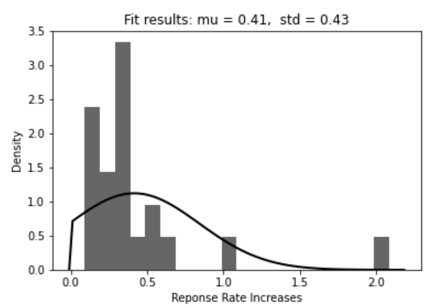

(b) Last 21 Weeks

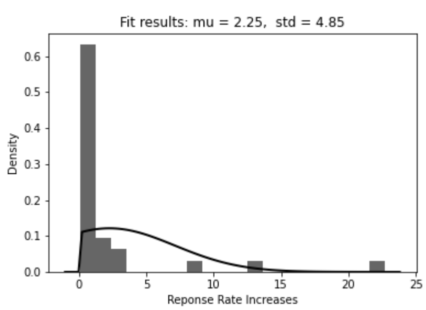

(a) All weeks with 1 truncated normal fit

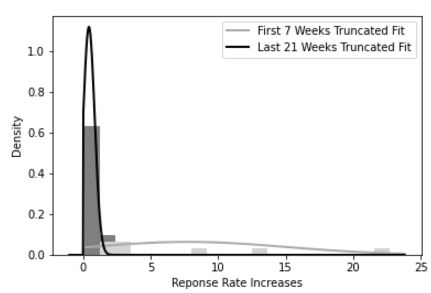

(b) All weeks with 2 truncated normal fits

Figure 6: Fitted transition distributions over real census data for a sample U.S. State