Food Fairness: An Artificial Intelligence Perspective for SNAP Allocation

Boli Fang[1] , **Miao Jiang**[1] , **Jerry Shen**[2] , **Pei-Yi Cheng**[1] and **Manju Chivukula**[1]
[1]Department of Computer Science, Indiana University, USA
[2]Price School of Public Policy, University of Southern California, USA
{bfang, miajiang, peicheng, mchivuku}@iu.edu, haoxuans@usc.edu

Abstract

To improve the food security of low-income households, Supplemental Nutrition Assistance Program (SNAP) provides help of food budgets to families in need. However, payment error rates of SNAP set a new decade high to 6.3 percent, revealing significant concerns with respect to mis-payment and unfair gains. In order to make sure that household living under poverty can fairly receive proper food and nutrition, we propose an algorithm that aims to include as many households experiencing food insecurity and poverty into the program as possible, while decreasing the number of households that do not need the food assistance. Using the legal notion of *disparate impact*, we construct fairgroups from the testing datasets to reflect the relative importance of different features, and apply logistic regression on these fairgroups. Our experiments show that our method effectively improves the outcome fairness of the distribution of scarce, common resources, while maintaining high accuracy in classification.

1 Introduction

"In a world of plenty, no one, not a single person, should go hungry. But almost 1 billion still do not have enough to eat. I want to see an end to hunger everywhere within my lifetime."
– Ban Ki-moon, Former United Nations Secretary-General

In the United States, 11.8 percent of (about 15 million) households are uncertain or unable to have enough food to meet their basic needs[Coleman-Jensen *et al.*, 2018]. While this problem is related to multiple causes, such as the existence of "food deserts"[Walker *et al.*, 2010], one of the factors for food insecurity is the lack of financial sources on household level.

To improve the food security of low-income households, Supplemental Nutrition Assistance Program (SNAP, formerly called the Food Stamp Program), the largest federal nutrition assistance program, provides help of food budgets to families in need. SNAP is currently one of the key components of the social safety net for low-income Americans. In recent years, fair allocation of SNAP resources affecting 40 million Americans is a crucial problem that needed to be solved. According to the requirements, households with incomes below the income-eligibility range and with elderly or disable members are the potential recipients for SNAP benefits. However, payment error rates, an indicator used to measure the integrity of the SNAP program, set a new decade high to 6.3 percent, revealing significant concerns with respect to fairness of the program.

Research community has witnessed machine learning and deep learning algorithms to help those who are at a disadvantage because of poverty, disability, etc. to obtain assistance while ensuring fairness [Morse, 2018]. Some of them focus on process fairness which concerns the expected allocation of resources, and others pay attention to outcome fairness which takes the final allocation of resources into account.

We propose an algorithm that aims at including as many households experiencing food insecurity and poverty into the SNAP program as possible, while decreasing the number of households that do not need the food assistance. In order to make sure that household living under poverty can fairly receive proper food and nutrition, outcome fairness is emphasized here. We optimize resource allocation by firstly deciding two groups of variables: protected features and unprotected features. Protected features are defined as prioritized features that play an important role in determining fairness, and unprotected features as other features that do not lead to fairness. After splitting variables into protected and unprotected features, we examine the representation of underprivileged groups with respect to protected features in the class of people receiving SNAP benefits. We construct fairgroup by computing feature relevance as revealed by correlation coefficients and classifying the data points with respect to each fair-group.

The contributions of this papercan be summarized as follows.

- An outcome-fairness algorithm is proposed to fairly allocate SNAP resources. We define fair-group and achieves fairness with respect to the protected features.
- Unprotected features are considered to make households with similar features that are not related to fairness would be classified in the same group.
- This fairness algorithm can be adapted to other fairness problems such as the earned income tax credit.

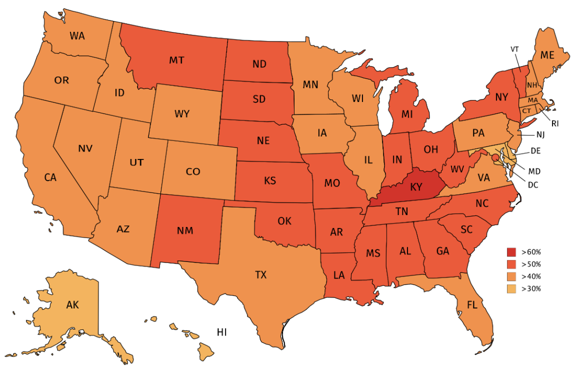

Figure 1: Percentage of SNAP Receiving Households below Poverty Line

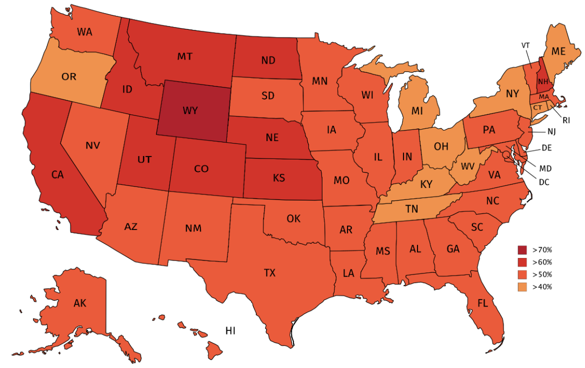

Figure 2: Percent of Impoverished Households NOT Receiving SNAP

Table 1: SNAP Recipient Poverty vs Not

	#	%
Below Poverty Line	7,420,946	49.4
Above Poverty Line	7,608,552	50.6
Total	15,029,498	100

Table 2: Impoverished HH Receiving vs Not Receiving SNAP

	#	%
Receiving SNAP	7,420,946	45.28
Not Receiving SNAP	8,969,163	54.72
Total	16,390,109	100

2 Existing Problems

We discovered, by observing the US Census American Community Survey Household Microdata, that there are two potential shortfalls in the current distribution of governmental subsidies used for food. These are:

1. A larger percentage of SNAP recipients are above poverty line than below poverty line.
2. A larger percentage of impoverished households are not currently receiving SNAP.

As shown in Table 1 & 2, on the National level, 50.5% of the SNAP recipients are living above poverty line, amounting to around 7.6 million households. On the other hand, currently, there are around 9 million impoverished households not receiving SNAP, which is 54.72% of all impoverished households.

On the State level, the intensity of the problem varies. States in the Midwest regions usually have a larger number of SNAP recipients living below poverty line, with Kentucky having over 60% of its SNAP recipients from impoverished households. By contrast, states in the Northeast region have a lower portion of their SNAP recipients from households below poverty line. At the same time, states such as Wyoming have more than 60% of its impoverished households not receiving food stamps.

On microdata level, there are also noticeable cases which doesn't make much sense. For example, while a family of 2 in California with no children making over $600,000 is receiving SNAP, another family of 2 in Wyoming with 1 children and an annual income of $9,500 is not. Such disparity signifies the need for a potential of change, and a possibility for artificial intelligence algorithms to intervene.

3 Related Work

Previous work on fairness in machine learning can be largely divided into two groups. The first group has centered on the mathematical definition and existence of fairness. Along this track, alternative measures such as statistical parity, disparate impact, and individual fairness [Chierichetti *et al.*, 2017] have been produced. Moreover, [Kleinberg *et al.*, 2016] suggested that it's not possible to achieve some desired properties of fairness at the same time.

The second group has centered on algorithms to achieve fairness. Along the route of disparate impact, [Feldman *et al.*, 2015] has described algorithms to spot the presence of disparate impact through Support Vector Machine, while [Chierichetti *et al.*, 2017] applied the notion of disparate impact to design an algorithm that achieves balance in unsupervised clustering algorithms. This paper also introduces the notion of *protected and unprotected features*.

4 Model

Since the variable corresponding to the actual SNAP allocation is binary, we can frame the allocation problem as a decision/classification problem involving various factors in the data set. Under such a setting, we present a novel strategy called *fair-grouping* to achieve fairness in classification. Our strategy adopts the notion of fairness as defined by *disparate impact* [Feldman *et al.*, 2015], where practices based on neutral rules and laws may still more adversely affect individuals with one protected feature than those without.

4.1 Preliminaries

We first define the terminology to be used in subsequent description. A *protected feature* is a feature that carries special importance and is of priority when making relevant decisions. An *unprotected feature*, on the other hand, is of relative minor importance in decision making. Since the problem in our paper primarily focuses on discrete label classification with discrete features, we assume, without loss of generality and for sake of simplicity, that the protected traits are binary and that the classification label class is also binary. Given a protected feature A along with the dataset, the *balance* B of the dataset with respect to A is defined as

$$Bal(A) = \min\{\frac{\#\{A=0\}}{\#\{A=1\}}, \frac{\#\{A=1\}}{\#\{A=0\}}\} \in [0,1],$$

where $Bal(A)=0$ refers to the case of all data points having the same feature value of A, and $Bal(A)=1$ refers to the case where $\#\{A=0\} = \#\{A=1\}$. A dataset is *α-fair* with respect to feature A if the balance of A does not go below a certain number $\alpha \in [0,1]$. In other words, a dataset is α-disparate with respect to A if the groups with 2 different values in A have a bounded and relative balanced numerical ratio between $\frac{1}{\alpha}$ and α. Following the doctrine of disparate impact as stated in [], we say that a classification is (α, i)-fair if the group corresponding to label i in the classification class $L=\{+,-\}$ is α-fair, meaning that the protected feature is fairly represented with balance at least α in group i.

4.2 Fair-group construction

We provide in this section the details of the algorithms we will use to achieve fairness in classification. Assume that we already have a classifier C which yields predictions for data points and might not yield α-fair classification results. Overall, our algorithm constructs fair-groups from testing data, and conducts classification on the data points with C while taking the properties of the fairgroups into consideration.

The sections below provide more details of our method.

Correlation Computation

Most of the social decision problems involve different features of varying degrees of relevance and importance to the goal. Therefore, we need a measure to describe the similarity. To achieve this goal, we compute the correlation coefficient between feature X_i and the outcome Y to determine the contribution of each feature to the final classification outcome:

$$Corr(X_i, Y) = \frac{E[X_iY] - E[X_i]E[Y]}{\sqrt{Var(X_i)Var(Y)}}.$$

We then rank all the features by an increasing order of the absolute values of correlation coefficients, because higher correlation values indicate greater statistical significance in either positive or negative directions. Then, we assign to each feature X_i a weight w_i which is equal to the rank by increasing values of the correlation coefficients. The weight w_i reflects the significance of feature X_i in the classifier.

After constructing the relative weight w_i of each feature X_i from the correlation coefficients, we examine the actual values of X_i for each data point j, here denoted by x_{ij}. If a feature X_i is positively correlated with Y, then we rank all data by the decreasing order of the corresponding x_{ij}'s of the feature X_i, and define r_{ij} as the rank of x_{ij} in the set of all values of X_i's. Alternatively, if a feature has negative correlation, the the data is ranked in increasing order of x_{ij}, and r_{ij}'s are defined accordingly. Intuitively, the rank r_{ij}'s show how much influence each feature X_i in data point j has to the final classification prediction. These ranks are constructed in a way to make sure that the data points with higher values of X_i are given enough consideration, since higher feature values in socialogical datasets are often likely to correspond to special cases requiring extra attention.

Finally, for each attribute X_i in corresponding to data point j, we define $r'_{ij} = w_i r_{ij}$ as the *feature importance index*, and define $\mathbf{r}'_j$ as the *feature importance vector* corresponding to data point j. The feature importance vector reveals information about the relative importance of data point j, and such information will be used to construct fairgroups for subsequent fair classification.

Fairgroup construction

With each data point now represented in the form of feature importance vectors, we now examine how close these data points are in terms of the influence each data point might exert to the final classification outcome, and how data points with similar features can be grouped together for easier analysis. To achieve these goals, we define a suitable distance between two vectors and consider a clustering problem where similar data points are grouped together.

Notice that each of the entries in the feature importance vectors are integers corresponding to different rankings, and that closer ranks imply similarity in one feature. Thus, we make use of the Manhattan-L1 distance to describe the distance between feature importance vectors $\mathbf{r}'_p, \mathbf{r}'_q$:

$$d(\mathbf{r}'_p, \mathbf{r}'_q) = \sum_{i=1}^{N} |r'_{ip} - r'_{iq}| = \sum_{i=1}^{N} w_i |r_{ip} - r_{iq}|,$$

Here N refers to the number of unprotected features.

Afterwards, we consider a k-median cluster algorithm to divide the entire dataset into k groups, each containing points with similar feature values. Within each cluster, we look at the protected features. Without loss of generality, we assume that the protected feature is binary, and that our goal is to maintain the balance of the protected feature A does not go below a certain threshold t. Since this requirement implies that the ratio between $\#\{A=0\}$ and $\#\{A=1\}$ falls between t and $\frac{1}{t}$, we match as many $A=0$ and $A=1$ data points as possible on condition that the ratio between $\#\{A=0\}$ and $\#\{A=1\}$ in each match falls between t and $1/t$. A set consisting of data points in such matches is denoted as a *fairgroup*.

Classification with respect to each fairgroup

It is now clear that within each fairgroup, the data points are similar and the ratio of points in different classes of protected attributes is balanced. For each fair-group we have thus constructed, we randomly pick a point to be classified by C. If the point is labeled as $+$, we apply the same label to all other data points in the group. Alternatively, if the point is labeled

Table 3: Features Used in the Experiment

VARIABLE	FEATURE	SAMPLE NON-RECIPIENT	SAMPLE RECIPIENT
DIVISION	DIVISION CODE	8 - MOUNTAIN REGION	9- PACIFIC
REGION	REGION CODE	4 - WEST	4 - WEST
ST	STATE CODE	56- WYOMING	6-CALIFORNIA
TEN	TENURE	3-RENTED	3-RENTED
HHL	HOUSEHOLD LANGUAGE	1-ENGLISH	1-ENGLISH
HINCP	HOUSEHOLD INCOME	9500	613000
HUGCL	HOUSEHOLD WITH GRANDPARENT LIVING W GRANDCHILDREN	0 - NO	0 - NO
NOC	# OF CHILDREN	1	0
NPF	# OF PERSONS IN FAMILY	2	2
R18	WHETHER SOMEONE IS UNDER 18 YO	1 - YES	0-NO
R60	WHETHER SOMEONE IS ABOVE 60 YO	0 - NO	0 - NO
WIF	WORKERS IN FAMILY	1	2
FS	FOOD STAMP RECIPIENCY	2 - NO	1 - YES

as $-$, we need to take into consideration the properties of the protected feature to determine whether other data points in the same fair-group will be given the same label. In our case of SNAP allocation, protected features such as poverty should be treated as a protected feature only in the positive label class, because our primary goal is to ensure that people receiving food stamps are mainly composed of people living under the poverty threshold, and it is relatively irrelevant to consider fairness out of the people who are rejected from receiving SNAP benefits.

Moreover, to reduce the negative effect of potential misclassification as much as possible, we construct as many fair-groups as possible by first expressing t and $\frac{1}{t}$ as ratios $\frac{p}{q}$ and $\frac{q}{p}$, where p, q are co-prime integers. Starting from $\frac{\#\{A=0\}}{\#\{A=1\}}$, we iteratively match p data points where $A = 0$ with q data points where $A = 1$(or q data points where $A = 0$ with p data points where $A = 1$) depending on whether $\frac{p}{q}$ or $\frac{q}{p}$ is smaller than and closer to the ratio of unmatched $\frac{\#\{A=0\}}{\#\{A=1\}}$. These matched $p + q$ points will form a fairgroup, and corresponding numbers of $A = 0$, $A = 1$ points will be moved from the unmatched point set. We repeat the procedure until all the points are matched or unmatchable.This procedure ensures that we create maximal numbers of fairgroups, so that even when one fairgroup is misclassified due to the misclassification of the randomly drawn point, the effects on the overall fairness and consistency can be minimal.

5 Data and Variables Used

To conduct experiments using the model explained above, we use the United States Census American Community Survey data. Consisting 7487361 entries, the household level micro-data displays important features, including geographical location, living condition, and household socio-economic status. The list of features used is listed above in Table 3.

6 Results and Conclusion

As indicated in Table 4, when using pure logistic regression, the percentage of SNAP recipients that are of low income is relatively low, with only around 36.4 percent of household having low income. By contrast, our method increases the percentage of low income SNAP recipients significantly, while maintaining a healthy model accuracy as compared with that obtained through pure logistic regression.

Table 4: Experiment result

METHODS	% OF POVERTY GET SNAP	MODEL ACCURACY
LOGISTIC REGRESSION	36.4	88.2
OUR METHOD	**79.3**	85.1

7 Conclusions and future work

In this work we present a novel approach to solve the current shortfalls of SNAP allocation through logistic-regression classifiers that achieve fairness in outcome. To achieve our goal, we propose the strategy of *fair-group* construction. As a part of our future work, we hope to apply our method to address other current social problems related to inequality and inequity in both the developed and developing world involving decisions in scarcity.

References

[Bureau, 2017] US Census Bureau. American community survey 2017 5-year estimate. 2017.

[Chierichetti *et al.*, 2017] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems*, pages 5029–5037, 2017.

[Coleman-Jensen *et al.*, 2018] Alisha Coleman-Jensen, Christian Gregory, and Matthew Rabbitt. Key statistics graphics, Sep 2018.

[Feldman *et al.*, 2015] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International*

Conference on Knowledge Discovery and Data Mining, pages 259–268. ACM, 2015.

[Kleinberg *et al.*, 2016] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.

[Morse, 2018] Susan Morse. Artificial intelligence helps insurers identify medicare members who also qualify for medicaid, Nov 2018.

[Walker *et al.*, 2010] Renee E Walker, Christopher R Keane, and Jessica G Burke. Disparities and access to healthy food in the united states: A review of food deserts literature. *Health & place*, 16(5):876–884, 2010.