

WORKING WOMEN AND CASTE IN INDIA: A STUDY OF SOCIAL DISADVANTAGE USING FEATURE ATTRIBUTION

Kuhu Joshi

International Food Policy Research Institute
South Asia Region, New Delhi, India
kuhu.joshi@cgiar.org

Chaitanya K. Joshi

School of Computer Science and Engineering
Nanyang Technological University, Singapore
chaitjo@gmail.com

ABSTRACT

Women belonging to the socially disadvantaged caste-groups in India have historically been engaged in labour-intensive, blue-collar work. We study whether there has been any change in the ability to predict a woman’s work-status and work-type based on her caste by interpreting machine learning models using feature attribution. We find that caste is now a less important determinant of work for the younger generation of women compared to the older generation. Moreover, younger women from disadvantaged castes are now more likely to be working in white-collar jobs.¹

1 INTRODUCTION

Working outside the house has historically been considered a social-stigma or a low-status activity for women in India (Eswaran et al., 2013). As a result, only the poorest women are working out of necessity in blue-collar jobs and once their family income increases, they withdraw from the workforce (Rao et al., 2010). Women tend to re-enter the work force only at high education levels that allow access to un-stigmatized white-collar jobs (Klasen & Pieters, 2012). In this cultural context, caste is an important determinant of a woman’s work-status. Since scheduled castes and scheduled tribes (Sc/St) are the socially and economically disadvantaged caste groups in India, Sc/St women have always had higher workforce participation, mainly in blue-collar jobs. Upper-caste (or general caste) women have historically been discouraged from working in order to maintain a higher social-status (Eswaran et al., 2013).²

We study whether the ability to infer a woman’s work-status based on her caste is changing over generations. We also study the change in the effect of caste on work-status over generations. Adopting strategy similar to Bertrand & Kamenica (2018), we use a nationally representative dataset from the National Family Health Survey (IIPS, 2017) for training ensemble Gradient Boosting Decision Tree models to predict women’s work-status. Our models uncover non-linear temporal patterns between caste and women’s work-status using the SHAP (SHapley Additive exPlanation) feature attribution framework (Lundberg & Lee, 2017).

SHAP values of a feature (such as belonging to the Sc/St caste) measure how important that feature is in predicting the outcome of a model (woman’s work-status). Using individual-level SHAP values, we find an upward trend across generations of working women, where caste is a less important determinant of work-status for younger women. Further, we unpack the impact of caste on blue-collar and white-collar jobs, finding that younger Sc/St women are moving out of blue-collar and into white-collar jobs in India.

Our paper is related to recent studies on machine learning for development (De-Arteaga & Dubrawski, 2017; De-Arteaga et al., 2018; Mullainathan & Spiess, 2017). We demonstrate a novel usage of the SHAP framework to study social disadvantage and how it affects women’s work in India. We also contribute to the aforementioned literature on caste and work by studying for the first time whether the younger generation of women are still facing the same level of disadvantage as older women.

¹Code and results available at <https://github.com/chaitjo/working-women>

²Bayly (2001) provides a comprehensive discussion of the caste system in India.

2 DATA

We use a nationally representative dataset from the National Family Health Survey conducted in 2015-16 (NFHS-4) consisting of 699,686 women aged 15-49 years across all 29 states and 7 union territories of India (IIPS, 2017). The survey has information on caste, work-status, and occupation types for a representative sub-sample of 111,398 women. We consider women over 21 years of age, leaving us with 81,816 women for whom we have detailed socio-economic information for our analysis.

We classify working women into those who have blue-collar type jobs (agriculture, skilled and unskilled manual labour, and domestic services) and those who have white-collar type jobs (professional, technical, managerial, clerical, and sales). We use sixteen socio-economic features for training machine learning models to predict work-status. Appendix A presents summary statistics of the dataset and the features used.

3 EXPERIMENTS

We design three binary classification experiments to predict a woman’s work: (1) having a job or not (*work-status*), (2) having a blue collar job or not (*blue-collar*), and (3) having a white collar job or not (*white-collar*). Using the sampling weights provided in NHFS-4, we create a test set with 5% of the data and use the remaining 95% for training our models in each experiment. We use stratified sampling to ensure that positive-to-negative class balance remains the same across training and test sets.

For each experiment, we train an ensemble Gradient Boosting Decision Tree (GBDT) model using LightGBM (Ke et al., 2017). We chose LightGBM due to generally strong empirical performance, fast training time, and easy compatibility with the SHAP framework. Optimal model hyperparameters are found using 5-fold cross-validated grid search over the training set. We use F1-score as the metric for choosing the best hyperparameters since it incorporates class imbalance. These hyperparameters are used to re-train each of the models on the entire training set.

Next, we interpret the trained models using the SHAP feature attribution framework for tree ensembles. Unlike traditional feature importance methods, SHAP allows for fast computation of complex tree-based, non-linear models. We compute SHAP values for our entire dataset, following Lundberg et al. (2018) as a guide to using and interpreting the explanations obtained. These explanations allow us to understand how a single feature affects the model’s output, summarize relative feature importance over the entire dataset, and analyze higher order interactions among feature pairs.

Since we use a nationally representative dataset, results from our experiments represent actual country-level patterns. However, it is important to note that the accuracy of feature attribution methods is proportional to model performance. See Appendix B for a detailed discussion.

4 RESULTS

Global Feature Importance The SHAP summary plot for *work-status* experiment in Figure 1 shows the relative importance of features, the distribution of impacts of features on the model’s prediction, as well as how the feature’s value (Low to High) relates to its impact. Each dot represents a woman in the dataset and the x-axis position of the dot is the impact of that feature on the model’s prediction. Dots that do not fit on the row pile up to show density. Since the GBDT model has a logistic loss, the x-axis has units of log-odds. Summary plots for *blue-collar* and *white-collar* experiments can be found in Appendix C.

Belonging to the Sc/St caste is the sixth most important feature in predicting work-status. State, household wealth index, age, years of education, and number of children below 5 years of age, are the five most important predictors of work-status. The observed patterns are similar to those found in existing literature: women from wealthier households are less likely to work (Rao et al., 2010), more educated women are more likely to work (Bhalla & Kaur, 2011), and Sc/St women are more likely to work (Srivastava & Srivastava, 2010).




Figure 1: SHAP summary plot for *work-status* experiment. Categorical variable values are grey.




Figure 2: Mean of the magnitude of SHAP values of Sc/St Caste across 5-year age-bins.

Importance of caste over generations For each of the three experiments, we plot the mean of the magnitude of SHAP values of Sc/St Caste feature across 5-year age-bins along with the 99% confidence interval to analyze whether the importance of caste has changed over generations (Figure 2). We find that caste is more important in predicting work-status of older women than of younger women. This pattern is also observed for blue-collar jobs. For white collar jobs, there isn't a clear monotonic pattern, but overall, caste is more important for women younger than 35 years of age.

Dependence and main effects of caste The SHAP dependence plot in Figure 3 plots the SHAP values of Sc/St Caste feature from the *work-status* experiment and colours them by age to visualize underlying interaction effects. We can see the attributed importance of caste change as its value varies. Higher SHAP values represent higher probability of working: Sc/St women are more likely to be working than women of other castes. The dependence plot also captures vertical dispersion at a single value of Sc/St Caste feature due to interaction effects with other features in the model. Colouring each dot by age (as the interacting feature) we find that older Sc/St women are most likely to be working.

Next, we analyze the main effects plot for the Sc/St Caste feature by removing all the interaction effects with other features in the *work-status* experiment (Figure 4). We find that Sc/St women are highly likely to be working while women of other castes are not likely to be working.




Figure 3: SHAP dependence plot of Sc/St Caste coloured by age for *work-status* experiment.




Figure 4: SHAP main effects of Sc/St Caste for *work-status* experiment.

Interactions between caste and age SHAP interaction plots capture all of the vertical dispersion that was present in the SHAP dependence plot (Figure 3) but was missing from the main effects plot (Figure 4). We only focus on interaction plots between the Sc/St Caste feature and age (Figure 5). For younger Sc/St women, SHAP interaction values are larger negatives, implying a high prediction for not-working. For older Sc/St women, SHAP interaction values are larger positives, implying a high prediction for working. Plotting the SHAP interaction values using age on the x-axis instead of caste, we get a sharper visualization of the same pattern (Figure 6).

Note that SHAP interaction values are expensive to compute. Hence, main effect and interaction plots use a random sub-sample of 10,000 women sampled using the weights provided in NFHS-4.




Figure 5: SHAP interaction plot of Sc/St Caste and age for *work-status* experiment.




Figure 6: SHAP interaction plot of age and Sc/St Caste for *work-status* experiment.

Interactions for job types For fine-grained analysis, we plot the age and Sc/St Caste interaction plots for both *blue-collar* and *white-collar* experiments. We find that the pattern observed in the *work-status* experiment is only relevant for blue-collar jobs (Figure 7). We observe an opposite pattern for white-collar jobs (Figure 8). Younger Sc/St women are more likely to be working in white-collar jobs while older Sc/St women are less likely to be working in white-collar jobs. This pattern is prominent mainly for women younger than 26 and older than 45 years, respectively. We also observe an upward spike for women aged 35 and 40.




Figure 7: SHAP interaction plot of age and Sc/St Caste for *blue-collar* experiment.




Figure 8: SHAP interaction plot of age and Sc/St Caste for *white-collar* experiment.

5 CONCLUSIONS

We train machine learning models to predict women’s work-status in India and analyze the relationship between work-status and caste using the SHAP feature attribution framework. We find that, over generations, caste has become a less important determinant of younger women’s work-status, specially their participation in blue-collar jobs. Next, we unpack the impact of caste on work-status using SHAP interaction effects, finding that younger women of scheduled castes and scheduled tribes are less likely to be working in blue collar-jobs, and more likely to be working in white-collar jobs.

For nearly 70 years, the government of India has implemented caste-based quotas in education and government jobs to ensure equal economic opportunity. Our work can be used as a cost-effective tool to monitor the impact of the existing caste-quota policy. In future work, we will look deeper into geographical variations to identify regions where younger women are lagging, and where job creation can be targeted. Our methodology can similarly be used to study the persistence of, as well as the nuanced patterns underlying other types of social disadvantage and bias in both developing and developed country contexts.

REFERENCES

- Susan Bayly. *Caste, society and politics in India from the eighteenth century to the modern age*, volume 3. Cambridge University Press, 2001.
- Marianne Bertrand and Emir Kamenica. Coming apart? cultural distances in the united states over time. Technical report, National Bureau of Economic Research, 2018.
- Surjit Bhalla and Ravinder Kaur. Labour force participation of women in india: some facts, some queries. 2011.
- Maria De-Arteaga and Artur Dubrawski. Discovery of complex anomalous patterns of sexual violence in el salvador. *arXiv preprint arXiv:1711.06538*, 2017.
- Maria De-Arteaga, William Herlands, Daniel B Neill, and Artur Dubrawski. Machine learning for the developing world. *ACM Transactions on Management Information Systems (TMIS)*, 9(2):9, 2018.
- Mukesh Eswaran, Bharat Ramaswami, and Wilima Wadhwa. Status, caste, and the time allocation of women in rural india. *Economic Development and Cultural Change*, 61(2):311–333, 2013.
- IIPS. National family health survey (nfhs-4), 2015-16: India. Technical report, Mumbai: International Institute for Population Sciences and ICF, 2017.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, pp. 3146–3154, 2017.
- Stephan Klasen and Janneke Pieters. Push or pull? drivers of female labor force participation during india’s economic boom. 2012.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pp. 4765–4774, 2017.
- Scott M Lundberg, Gabriel G Erion, and Su-In Lee. Consistent individualized feature attribution for tree ensembles. *arXiv preprint arXiv:1802.03888*, 2018.
- Kristin Mammen and Christina Paxson. Women’s work and economic development. *Journal of economic perspectives*, 14(4):141–164, 2000.
- Sendhil Mullainathan and Jann Spiess. Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106, 2017.
- Nitya Rao, Arjan Verschoor, Ashwini Deshpande, and Amaresh Dubey. Gender, caste and growth assessment–india (report to uk department for international development). 2010.
- Nisha Srivastava and Ravi Srivastava. Women, work, and employment outcomes in rural india. *Economic and political weekly*, pp. 49–63, 2010.

A DATASET STATISTICS

Table 1 shows the distribution of various castes across work-status types in the representative sample dataset of women above 21 years of age. Overall, 34.1% of women are working, 28.3% of women are doing blue-collar jobs and 5.8% are employed in white-collar jobs. 42.7% of scheduled caste or scheduled tribe women are working, while only 23.9% of general caste women are working.

Table 1: Distribution of number of women from each caste across work-status types.

Work-status	General Caste	Sc/St Caste	OBC Caste	Unknown Caste	Total
Unemployed	13,978	17,480	22,254	230	53,942
Blue-collar	3,068	11,235	8,730	108	23,141
White-collar	1,341	1,787	1,592	13	4,733

Table 2 presents the summary statistics of the sixteen socio-economic features provided as inputs to machine learning models: each woman’s personal biographical data including age, years of education, state of residence and caste (split into four binary features); information about her household such as type of residence (rural/urban), household religion, household wealth index and household members; and her health/reproductive history represented by number of children ever born, number of children currently under five years of age, whether she is anemic, whether she is obese, and her frequency of watching TV.

Among these features, we are particularly interested in looking at the relationship between work-status and caste, which may take four values: general caste, Sc/St caste, other backward class (OBC), or unknown. To facilitate detailed analysis using feature attribution, we split caste into four binary features for each woman: whether she is from the general caste, whether she is from the Sc/St caste, whether she is from the OBC caste, and whether her caste is unknown. Having a positive value for any one of these four features means that the values of the other features are zeros.

Table 2: Summary statistics for the features used as inputs to machine learning models.

Feature	Age	Yrs. of Education	State	Residence Type	Hh. Religion	Wealth Index	Hh. Members	Freq. of TV
Type	Numeric	Numeric	Categorical	Binary	Categorical	Numeric	Numeric	Numeric
Values/Range	[21, 49]	[0, 20]	36 Categories	{0, 1}	10 Categories	[0, 4]	[1, 39]	[0, 3]
Mean	33.765	6.301	-	0.309	-	2.088	5.546	2.047
Std. Dev.	8.153	5.401	-	0.462	-	1.389	2.553	1.256

Feature	Total Children	Children Below 5	Anemic	Obese	General Caste	Sc/St Caste	OBC Caste	Unknown Caste
Type	Numeric	Numeric	Binary	Binary	Binary	Binary	Binary	Binary
Values/Range	[0, 15]	[0, 9]	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}
Mean	2.426	0.632	0.520	0.234	0.224	0.372	0.398	0.004
Std. Dev.	1.749	0.907	1.389	0.499	0.417	0.483	0.489	0.065

B MODEL PERFORMANCE AND SHAP EXPLANATIONS

Table 3 presents the performance of the three GBDT models described in Section 3 on the respective training and test sets. We evaluate our models based on binary accuracy and F1-score. Models for *work-status* and *blue-collar* perform acceptably and do not overfit to the training set. However, F1-score for the *white-collar* model clearly indicates that it has overfit to the training set and generalizes poorly to unseen data. We believe this is due to extremely low class balance.

It is important to note that the SHAP plots explain how GBDT models work, but the models are not guaranteed to be causal. Outcomes of predicting work-status may have been impacted by several hidden factors not provided to the model or for which data collection was impossible. Hence, model performance and generalization to unseen data is a good indicator of how much we should trust the explanations produced by SHAP. In practice, we observe that using the sampling weights from NFHS-4 leads to explanations that are coherent with existing literature on gender and labour in India.

Table 3: Training and test set performance for GBDT models. Class balance is computed as number of samples from the positive class divided by total number of samples.

Experiment	Class Balance	Training Set		Test Set	
		Accuracy	F1-score	Accuracy	F1-score
<i>work-status</i>	0.341	0.692	0.604	0.673	0.578
<i>blue-collar</i>	0.283	0.716	0.596	0.681	0.551
<i>white-collar</i>	0.058	0.898	0.510	0.853	0.305

C ADDITIONAL VISUALIZATIONS

Summary plots for work-types Figures 9 and 10 show the SHAP summary plot for *blue-collar* and *white-collar* experiments, respectively. For blue-collar jobs, belonging to the Sc/St caste is the sixth most important feature after years of education, household wealth index, state, age, and number of children below 5 years of age. Highly educated women or those from wealthy households are unlikely to be employed in blue-collar jobs.

Years of education, state, age, total number of children, and household members are the five most important predictors of white-collar employment. None of the features related to caste are important predictors. Highly educated women are likely to be employed in white-collar jobs, whereas women with many children or large household sizes are not.

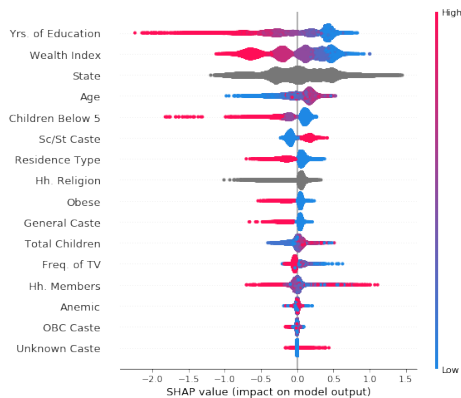


Figure 9: SHAP summary plot for *blue-collar* experiment.




Figure 10: SHAP summary plot for *white-collar* experiment.

Household wealth and work-status The dependence plot of household wealth index in Figure 11 shows that most of the Sc/St women in the dataset are from the bottom three wealth levels. For women belonging to wealthier households, our model gives a high prediction for not-working. This is consistent with socio-economic literature on women’s work in India (Mammen & Paxson, 2000).

State-wise patterns of work-status The dependence plot in Figure 12 shows differences in the caste-composition of working women across the states of India. It also shows state-wise differences in the likelihood of working. Women in Andhra Pradesh, Telangana, Maharashtra, Manipur, and Mizoram are most likely to be working. Women in Bihar, Assam, and Jammu and Kashmir are least likely to be working.

Figure 13 plots the mean of the magnitude of SHAP values of Sc/St Caste feature across states of India. It shows the distribution of the importance of caste in predicting work-status across states.




Figure 11: SHAP dependence plot of household wealth index coloured by Sc/St Caste for *work-status* experiment.




Figure 12: SHAP dependence plot of state coloured by Sc/St Caste for *work-status* experiment.

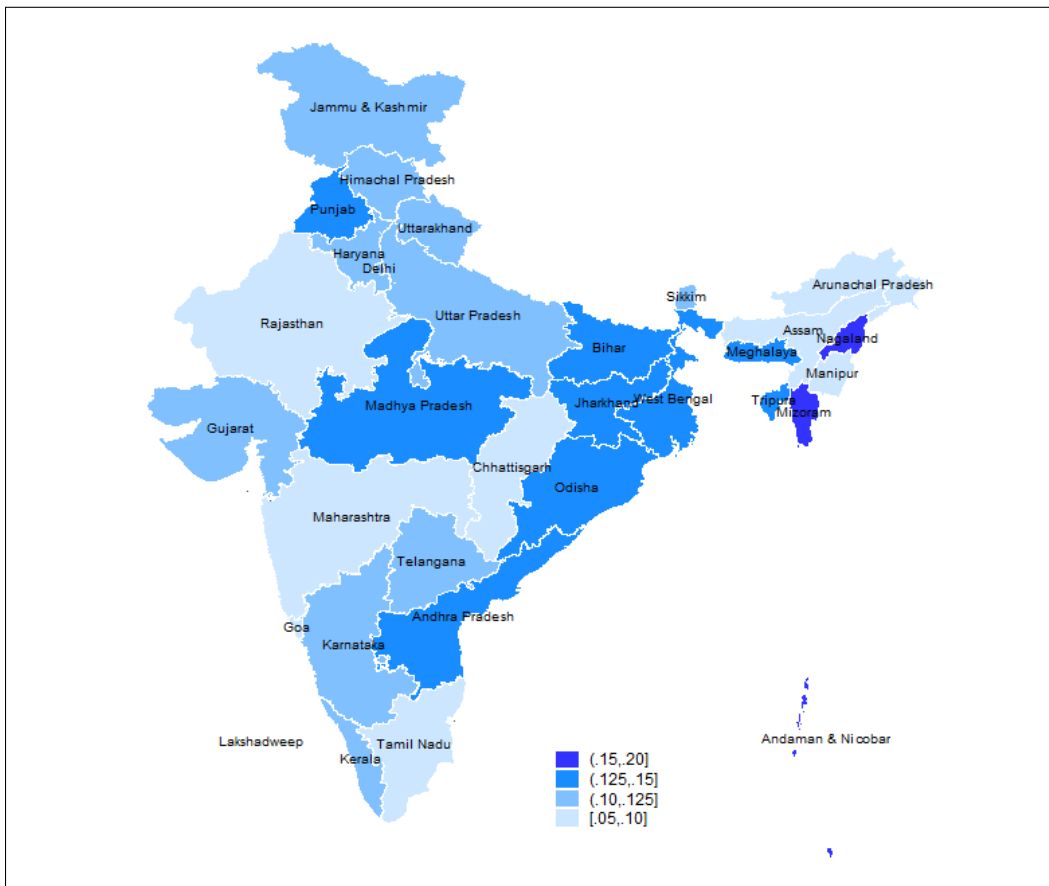


Figure 13: Mean of the magnitude of SHAP values of Sc/St Caste across states of India for *work-status* experiment

Interaction heatmap Figure 14 shows a heatmap of magnitude of SHAP interaction values for the *work-status* experiment. SHAP interaction values are a generalization of SHAP values to higher order interactions. The SHAP framework computes a matrix for every prediction, where the main effects are on the diagonal and the interaction effects are off-diagonal. The main effects are similar to the SHAP values for a linear model. The interaction effects capturing all the higher-order interactions are divided up among the pairwise interaction terms.

The heatmap shows that pairwise interactions between the more predictive features for the model, such as state, wealth index and years of education, are the most important at capturing interaction effects. For the Caste Sc/St feature, we see that state and age are the most important interaction effects.




Figure 14: Heatmap of magnitude of SHAP interaction values for *work-status* experiment. Shades from blue to yellow denote increasing strength of interaction.