

CAN STRATEGIC DATA COLLECTION IMPROVE THE PERFORMANCE OF POVERTY PREDICTION MODELS?

Satej Soman

School of Information
University of California, Berkeley
Berkeley, CA 94704, USA
satej@berkeley.edu

Emily Aiken

School of Information
University of California, Berkeley
Berkeley, CA 94704, USA
emilyaiken@berkeley.edu

Esther Rolf

Harvard Data Science Initiative &
Center for Research on Computation and Society
Harvard University
Cambridge, MA 02138, USA
erolf@seas.harvard.edu

Joshua Blumenstock

School of Information
University of California, Berkeley
Berkeley, CA 94704, USA
jblumenstock@berkeley.edu

ABSTRACT

Machine learning-based estimates of poverty and wealth are increasingly being used to guide the targeting of humanitarian aid and the allocation of social assistance. However, the ground truth labels used to train these models are typically borrowed from existing surveys that were designed to produce national statistics – not to train machine learning models. Here, we test whether adaptive sampling strategies for ground truth data collection can improve the performance of poverty prediction models. Through simulations, we compare the status quo sampling strategies (uniform at random and stratified random sampling) to alternatives that prioritize acquiring training data based on model uncertainty or model performance on sub-populations. Perhaps surprisingly, we find that none of these active learning methods improve over uniform-at-random sampling. We discuss how these results can help shape future efforts to refine machine learning-based estimates of poverty.

1 INTRODUCTION

Over the past decade, several influential studies have demonstrated how machine learning (ML) – when applied to digital data sources – can provide accurate, timely, and granular estimates of wealth and poverty (Blumenstock et al., 2015; Jean et al., 2016; Steele et al., 2017; Blumenstock, 2018; Burke et al., 2020; Yeh et al., 2020; Chi et al., 2022). These ML-based estimates, which can be produced at a fraction of the cost of traditional survey-based methods of data collection, can fill in data gaps in poor, fragile, and conflict-affected regions where recent surveys do not exist (Jerven, 2013). And increasingly, these estimates are being used to inform critical policy decisions, including the targeting of humanitarian aid (Aiken et al., 2022) and the allocation of social assistance (Smythe & Blumenstock, 2022; Gentilini et al., 2021).

Despite the important role ML-based poverty estimates now play in development policy, very little consideration has been given to the sampling strategy used to collect the ground truth labels used to train the ML models. Instead, the vast majority of poverty prediction models are trained on publicly-available household surveys that were collected for a different purpose (typically, to provide nationally representative demographic and health statistics). These standardized surveys typically use (stratified) random sampling to reduce measurement error (Howes & Lanjouw, 1998). However, sampling training data according to these strategies may not optimize predictive power for a machine learning algorithm subject to a budget constraint for data collection. In an *adaptive sampling* strategy, each incremental round of data acquisition is informed by an aspect of the performance of the model trained on data seen in previous rounds of data collection. In this paper, we

investigate whether adaptive sampling strategies for label data collection improve the performance of machine learning models trained to predict poverty from digital data sources, focusing on model accuracy and fairness.

This work builds on an emergency cash-transfer program run by the government of Togo that used poverty estimates from digital data sources to prioritize the distribution of emergency mobile money payments in response to the COVID-19 pandemic (described in more detail in Appendix A and Aiken et al. (2022)). In designing the program, the government sought to distribute cash to 60,000 of the country’s poorest individuals. At the time, however, the government did not have any traditional poverty data with which to determine eligibility. Instead, the Togolese government turned to satellite imagery analysis to identify the country’s poorest districts, and then used machine learning models using features derived from mobile phone service subscribers’ call data records to identify the poorest subscribers in eligible districts for humanitarian aid. The ground truth labels for the wealth prediction models came from surveys conducted in Togo prior to the program’s launch. An evaluation of the program’s targeting found that relative to alternative poverty targeting methods (such as geographic or occupation-based targeting), the phone-based approach was most effective at targeting cash transfers to the poorest people in the country (Aiken et al., 2022).

In this paper, we study whether using an adaptive strategy for sampling the poverty labels obtained from the phone survey in Togo could have improved the accuracy of the trained machine learning model. Our results speak most directly to recent efforts to estimate wealth and poverty from mobile phone data. More broadly, these results can inform the larger community of researchers and policy-makers interested in estimating social and demographic characteristics from digital data in contexts where traditional data are unavailable or out-of-date.

2 METHODS

We conduct experiments to simulate various data acquisition strategies using an in-person household survey conducted in Togo in 2018 ($N = 4,595$) with a stratified random sampling strategy¹. The survey collected information on daily per-capita consumption for each household interviewed (measured in USD/day), which we use as our ground truth measure of poverty. We match the household survey to anonymized features (such as number of text messages sent, proportion of international calls made, etc.) derived from call data records provided by the nation’s two cellular network operators. After partitioning the survey-labelled features into a label pool ($N_{LP} = 0.75N = 3,446$) and a holdout validation set ($N_V = 0.25N = 1,149$), we train random forest regressor models (Pedregosa et al., 2011) on subsets of the label pool of incrementally increasing size to simulate how the resulting model performs throughout the data acquisition process. The model’s performance is evaluated against the holdout set across all dataset sizes.

The baseline of acquiring data by uniform random sampling is compared against two adaptive data acquisition approaches, in which each data point still remaining in the label pool is assigned a heuristic sampling weight. In the *query-by-committee* strategy, this weight is proportional to the variance of the outputs of the individual decision trees within the random forest regressor; unlabeled points for which the ensemble’s constituent models disagree the most are upweighted in the next round of data acquisition. An alternative is *margin-based uncertainty sampling*, in which an independent logistic regression classifier is trained on the same features to predict whether a given individual falls below a relevant poverty threshold; points whose probability of classification is closer to 50% are upweighted in order to gain more information on points whose classification is close to random. Query-by-committee and uncertainty sampling are two standard methods for active learning for model accuracy (Settles, 2009). They have both been shown to substantially improve model performance subject to a budget constraint for data collection in small and relatively low-dimensional datasets unrelated to the poverty prediction task we consider in this paper (Settles, 2009; Settles & Craven, 2008; McCallumzy & Nigamy, 1998; Körner & Wrobel, 2006).

Inspired by literature on using active learning for improving fairness outcomes, (Abernethy et al., 2020; Cai et al., 2022), we also consider adaptively constructing the training dataset by increasing the sampling weights for the points from the demographic categories with the lowest accuracy

¹While the survey was stratified by region and urban/rural designation, our experiment treats this survey as if it were sampled uniform-at-random.

(*accuracy-weighted*), highest mean-squared-error (*MSE-weighted*), or highest exclusion/inclusion error rates (*disparity-weighted*). In these group-based strategies, each member of the demographic category is assigned the same weight, but each group’s weight differs by group-specific model performance in each round of simulated data collection.

Focusing on these demographic category weights also allows us to assess aspects of fairness of the model during the data collection process, as measured by minimum group accuracy, maximum group MSE, or total inclusion-exclusion errors. These goals are independent of increasing the model’s predictive power: even if deviation from uniform-at-random sampling provides no benefits in increasing accuracy or reducing MSE, there may be benefits in model fairness. In this investigation, the demographic category is geographic residence in each of the major top-level administrative regions of Togo (depicted in more detail in Figure 3a), based on previous analysis indicating there may exist slight geographic bias to the phone-based models (Aiken et al., 2022).

We run 50 simulations for each strategy, with incremental dataset sizes S_t logarithmically spaced between 0 and the full label pool size to better illustrate the effects of different strategies at smaller sample sizes. Within each simulation, the newly sampled points are appended to the previously-seen dataset (i.e. $S_t - S_{t-1}$ new data points are sampled at time t rather than S_t points sampled memorylessly).

3 RESULTS

We present the results of each of the sampling strategies described, comparing individual-level active learning methods (query-by-committee and uncertainty sampling) and group-level active learning methods (accuracy-, MSE-, and disparity-weighting) to a baseline of uniform-at-random sampling. Figures 1 and 2 show the average performance of each group of sampling methods across 50 simulations, with shaded regions representing bootstrapped 95% confidence intervals. In Figure 1, we show the performance of the point-based sampling strategies, evaluated across a number of standard machine learning accuracy metrics (Spearman ρ , mean-square-error, AUROC, accuracy, precision, and recall).² It is clear that none of the point-based strategies provide substantial gains to overall model performance above uniform at random sampling. With some strategies, such as the query by committee approach, we do see some gains to model precision, but these are on the order of 1% and are offset by reduction in recall (which is a more important metric for aid prioritization, since a model with lower recall will mistakenly exclude more recipients who should receive aid).

Our group-based sampling strategies rely on upweighting groups (regions of Togo) for which our model performs worse. We evaluate these strategies based on the same metrics as above, and three additional fairness-related metrics calculated for the major geographic regions in Togo. In Figure 2, we plot minimum group accuracy, maximum group MSE, and a summary statistic for targeting inclusion/exclusion errors (absolute demographic deviation; ADD, explained in more detail in Appendix B). Again, we observe negligible differences in fairness between our active learning approaches and the uniform-at-random baseline sampling strategy.

Additional experiments in Appendix C detail similar results when using a feature set of reduced dimensionality and a hyperparameter tuning procedure across dataset sizes.

4 DISCUSSION

Our results indicate that a well-implemented, uniformly-sampled survey may be close to optimal as a source of training data for machine learning models that aim to predict poverty levels from digital trace data. In particular, neither the model-based data acquisition strategies nor the group characteristic weighting approaches we considered showed appreciable gains on either accuracy or fairness metrics over simple uniform random sampling.

This result has important implications for the growing research literature – and increasing number of policy applications – that rely on estimates of poverty predicted from digital trace data. Nearly all such estimates have been generated from “found” data, i.e., survey data that were collected for

²Poverty predictions are continuous and produced with regression models; we binarize predictions using the international poverty line of US\$1.90/day to obtain binary classification metrics.

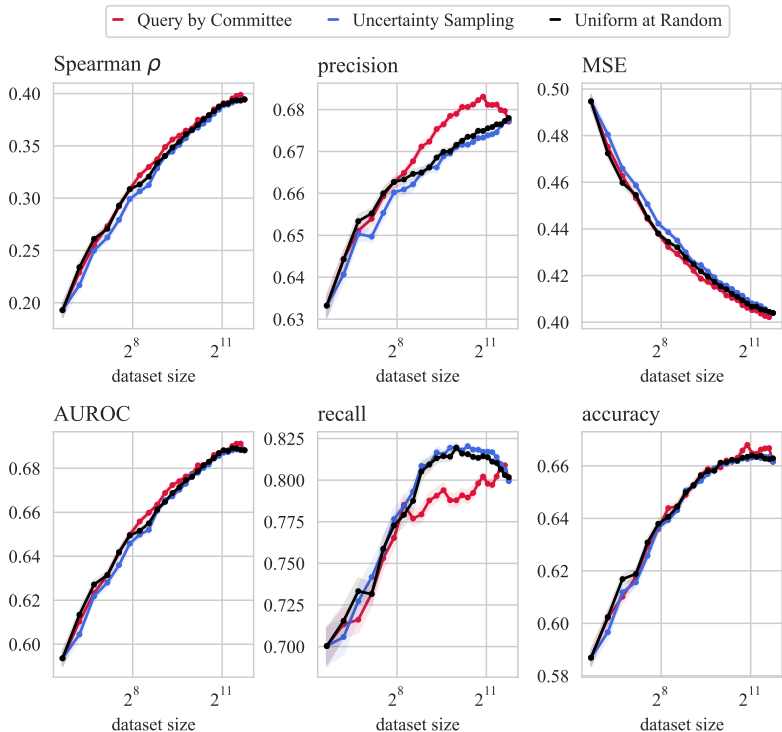


Figure 1: Comparison of model-based data acquisition strategies on key machine learning performance metrics as dataset size incrementally grows.

another purpose, and which were subsequently re-used as labels for training a poverty prediction algorithm (e.g., Blumenstock et al., 2015; Jean et al., 2016; Steele et al., 2017; Blumenstock, 2018; Burke et al., 2020; Yeh et al., 2020; Chi et al., 2022). Our hypothesis in initiating this analysis was that such re-use was sub-optimal; if the poverty prediction use case was known *ex ante*, we expected there to be performance gains from a more strategic approach to obtaining training data. However, our results indicate that — at least for the adaptive approaches that we tested — no such gains exist.

Whether this is a general finding, or one that is specific to our particular data context (i.e., mobile phone data from Togo), is an area of future work that we are continuing to explore. We also see several new directions in which to extend this analysis. For instance, all of the simulations described here assume an identical labeling cost for all data points. This uniform-cost model is appropriate in settings where the cost of interviewing an individual is the same regardless of their physical location in the country — as is the case with phone-based surveys. By contrast, many settings involve highly heterogeneous sampling costs — such as household surveys where some households are easily accessed and others are remote. In such settings, factoring in the costs of traveling to survey locations and conducting surveys may provide more realistic strategies for collecting training labels in household surveys. In the current investigation, in which groups are defined geographically, group-specific sampling is especially cost-infeasible because where the group to be sampled may change in every round of incremental data acquisition. In continued work, we plan to investigate cost-sensitive data acquisition strategies, and generate theoretical explanations for the lack of gains in any of the non-uniform sampling strategies.

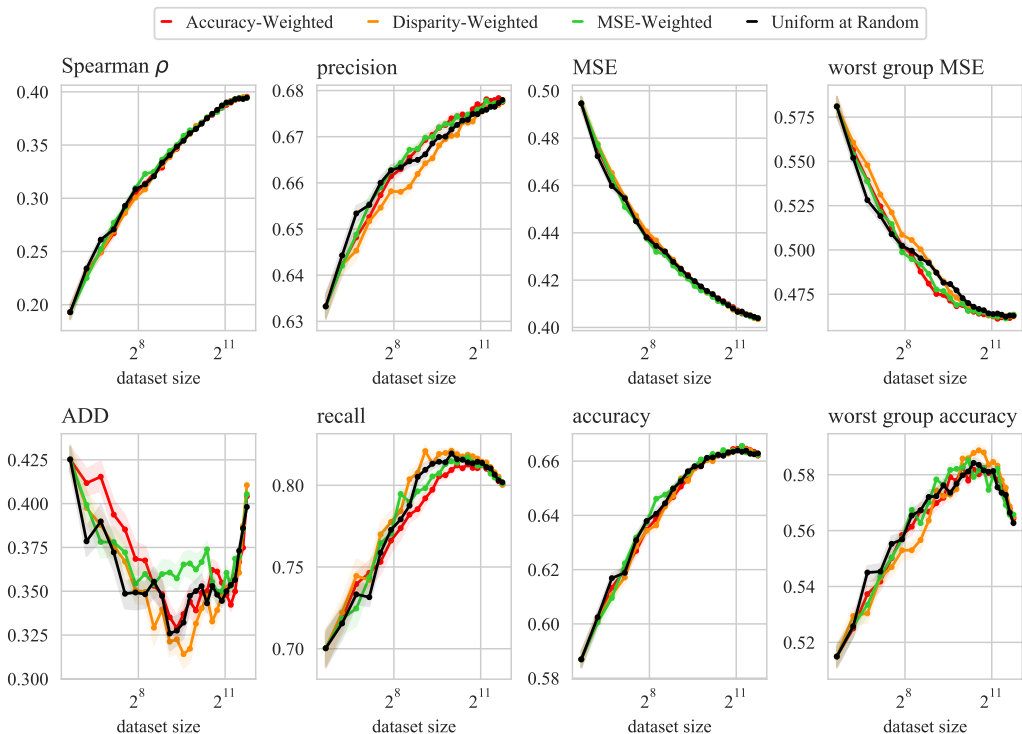


Figure 2: Comparison of group weighting strategies on standard machine learning metrics, as well as fairness-oriented metrics such as minimum group-level accuracy, maximum group-level mean-square-error, and total inclusion/exclusion errors (ADD).

REFERENCES

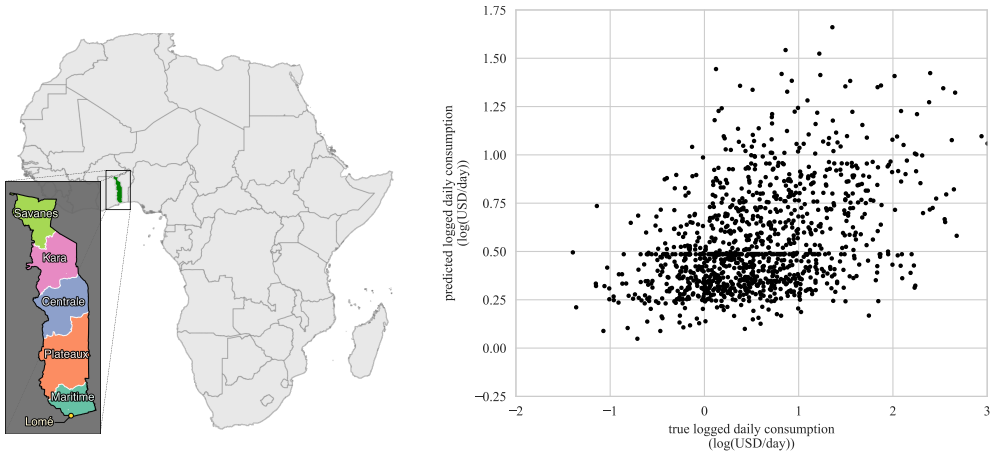
- Jacob Abernethy, Pranjal Awasthi, Matthäus Kleindessner, Jamie Morgenstern, Chris Russell, and Jie Zhang. Active sampling for min-max fairness. *arXiv preprint arXiv:2006.06879*, 2020.
- Emily Aiken, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock. Machine learning and phone data can improve targeting of humanitarian aid. *Nature*, 603(7903):864–870, 2022.
- Joshua E Blumenstock. Estimating economic characteristics with phone data. In *AEA papers and proceedings*, volume 108, pp. 72–76, 2018.
- Joshua E Blumenstock, Gabriel Cadamuro, and Robert On. Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264):1073–1076, 2015.
- Marshall Burke, Anne Driscoll, David Lobell, and Stefano Ermon. Using Satellite Imagery to Understand and Promote Sustainable Development. Technical Report w27879, National Bureau of Economic Research, October 2020. URL <https://www.nber.org/papers/w27879>.
- William Cai, Ro Encarnacion, Bobbie Chern, Sam Corbett-Davies, Miranda Bogen, Stevie Bergman, and Sharad Goel. Adaptive sampling strategies to construct equitable training datasets. *arXiv preprint arXiv:2202.01327*, 2022.
- Guanghua Chi, Han Fang, Sourav Chatterjee, and Joshua E. Blumenstock. Microestimates of wealth for all low- and middle-income countries. *Proceedings of the National Academy of Sciences*, 119(3), January 2022. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2113658119. URL <https://www.pnas.org/content/119/3/e2113658119>. ISBN: 9782113658118 Publisher: National Academy of Sciences Section: Social Sciences.

- Ugo Gentilini, Saksham Khosla, and Mohamed Almenfi. Cash in the city. Technical report, World Bank, 2021.
- Stephen Howes and Jean Olson Lanjouw. Does sample design matter for poverty rate comparisons? *Review of Income and Wealth*, 44(1):99–109, 1998.
- Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.
- Morten Jerven. Poor numbers. In *Poor Numbers*. Cornell University Press, 2013.
- Christine Körner and Stefan Wrobel. Multi-class ensemble-based active learning. In *European conference on machine learning*, pp. 687–694. Springer, 2006.
- Andrew Kachites McCallumzy and Kamal Nigamy. Employing em and pool-based active learning for text classification. In *Proc. International Conference on Machine Learning (ICML)*, pp. 359–367. Citeseer, 1998.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison, 2009.
- Burr Settles and Mark Craven. An analysis of active learning strategies for sequence labeling tasks. In *proceedings of the 2008 conference on empirical methods in natural language processing*, pp. 1070–1079, 2008.
- Isabella S. Smythe and Joshua E. Blumenstock. Geographic microtargeting of social assistance with high-resolution poverty maps. *Proceedings of the National Academy of Sciences*, 119(32): e2120025119, August 2022. doi: 10.1073/pnas.2120025119. URL <https://www.pnas.org/doi/10.1073/pnas.2120025119>. Publisher: Proceedings of the National Academy of Sciences.
- Jessica E. Steele, Pål Roe Sundsøy, Carla Pezzulo, Victor A. Alegana, Tomas J. Bird, Joshua Evan Blumenstock, Johannes Bjelland, Kenth Engø-Monsen, Yves-Alexandre de Montjoye, Asif M. Iqbal, Khandakar N. Hadiuzzaman, Xin Lu, Erik Wetter, Andrew J. Tatem, and Linus Bengtsson. Mapping poverty using mobile phone and satellite data. *Journal of The Royal Society Interface*, 14(127):20160690, February 2017. ISSN 1742-5689, 1742-5662. doi: 10.1098/rsif.2016.0690. URL <http://rsif.royalsocietypublishing.org/content/14/127/20160690>.
- Alexandru Tifrea, Jacob Clarysse, and Fanny Yang. Uniform versus uncertainty sampling: When being active is less efficient than staying passive. In *Workshop on Adaptive Experimental Design and Active Learning in the Real World*, 2022.
- Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature communications*, 11(1):1–11, 2020.

A ADDITIONAL BACKGROUND

Togo is a West African country with a population of 8 million, over half of whom live below with international poverty line. In March 2020, the government imposed a lockdown order within the country in response to the first confirmed COVID-19 cases in the country. These lockdowns proved economically disruptive, with risks of food insecurity rising as work became scarce for Togolese citizens employed informally (Aiken et al., 2022).

After an initial pilot of an emergency cash transfer program, named “Novissi”, implemented through mobile money payments to residents in the greater metropolitan area surrounding the capital city of Lomé, the Togolese government sought to expand the program to 100 of Togo’s poorest 397 cantons. Registrants were required to use a SIM card to register for the expanded program. Historical call records and cell phone usage metadata tied to those SIM cards were used to predict consumption, as illustrated in Figure 3b.



(a) Location of Togo (green) in Africa and major regions of the country (inset).

(b) True versus predicted log daily consumption for the holdout validation set, with a random forest trained on phone usage features.

Figure 3: Geographic context and poverty prediction model output for Togo.

Given the geographic distribution of socioeconomic characteristics and potential differences in cell phone penetration and usage patterns across the country, it is not surprising that models trained using these data show differential performance across populations living in different regions (Fig. 4) (or other demographic breakdowns). Aiken et al. (2022) found that aid targeting errors using this approach, despite the variance in performance, were lower than other, non-machine-learning-based approaches. Nevertheless, the difference in performances across regions was the starting point for our investigations into improving fairness characteristics, described below.

B FAIRNESS METRICS EXPLAINED

B.1 ACCURACY

We follow the standard definition of accuracy, but calculate it for every group g in the demographic category under consideration, G :

$$A_g = \frac{TP_g + TN_g}{TP_g + TN_g + FP_g + FN_g} .$$

where TP_g , TN_g , FP_g , and FN_g are the true positive, true negative, false positive, and false negative rates, respectively, calculated for group g .

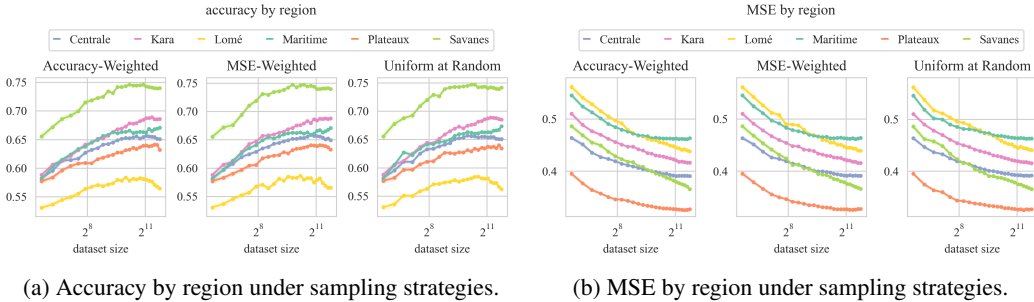


Figure 4: Per-group accuracy and MSE for each region of Togo. Note the differential performance of the model across regions, and the slight differences in the crossover points at which the model’s performance for one subgroup improves compared to other subgroups, such as the MSE in Maritime vs Centrale.

The summary statistic, worst-case group accuracy, is the minimum over all groups: $\min_{g \in G} A_g$. Because we want to sample from the groups with the lowest accuracy, the weighting for each unseen data point i in the label pool belonging to group g is proportional to the group accuracy subtracted from one: $w_i \propto 1 - A_g$.

B.2 MEAN SQUARED ERROR

As with accuracy, we calculate mean squared error for each demographic group:

$$\text{MSE}_g = \frac{1}{|g|} \sum_{i \in g} (y_i - \hat{y}_i)^2 .$$

The summary statistic across groups, worst-case group MSE, is the maximum, and the sampling weight for each point i is directly proportional to the group MSE.

B.3 DEMOGRAPHIC DISPARITY AND ABSOLUTE DEMOGRAPHIC DEVIANCE

The final fairness metric we present is based on a measure of targeting exclusion developed in Aiken et al. (2022), namely group demographic parity (DP)³: defined as the proportion of each group predicted to live below the US\$ 1.90 threshold, minus the true proportion of the group living below the same threshold:

$$\text{DP}_g = \frac{(\text{TP}_g + \text{FP}_g) - (\text{TP}_g + \text{FN}_g)}{N_g} .$$

A positive value for DP indicates a given group is over-represented in the model-based targeting, while a negative value means the group is insufficiently targeted for aid receipt. For the disparity weighting strategy, the sampling weight is, like accuracy, proportional to one minus the group demographic parity value: $w_i \propto 1 - \text{DP}_g$.

For a summary statistic, we define absolute demographic deviance (ADD) as the sum of the absolute values of the demographic parity for each group, with the intent of capturing both errors of inclusion as well as errors of exclusion in a single number:

$$\text{ADD} = \sum_{g \in G} |\text{DP}_g| .$$

³This definition of demographic parity differs from other definitions in the relevant literature.

C DIMENSIONALITY REDUCTION AND CROSS-VALIDATION

C.1 DIMENSIONALITY REDUCTION

One potential issue with active learning approaches is that the dimension of the feature space makes it difficult to isolate the important dimensions along which new point selection would be beneficial (Tifrea et al., 2022). To test this hypothesis in our context, we used a principal components analysis to reduce the dimension of our feature space from approximately 850 features to 50. As seen in Figure 5, no appreciable gains were seen to using strategies other than uniform-at-random sampling in the PCA dataset (lighter colors), either.

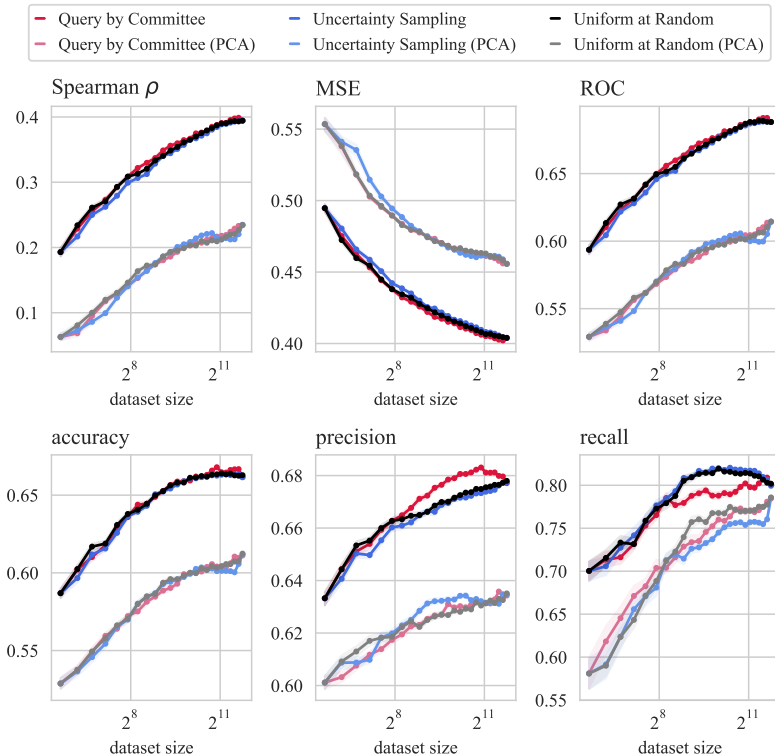


Figure 5: Comparison of point-based strategies on the original data (darker colors) versus a PCA-compressed representation of the data (lighter colors).

C.2 CROSS-VALIDATION

To maintain a simpler experimental setup, we skipped the usual cross-validation steps on models trained at each incremental dataset size S_t . To ensure that the lack of cross-validation did not artificially constrain the model performance, we present the results of cross-validating the model using three-fold cross-validation on the training set to find optimal tree depth at each incremental dataset size. As with dimensionality reduction, cross-validation did not show appreciable gains to model characteristics.

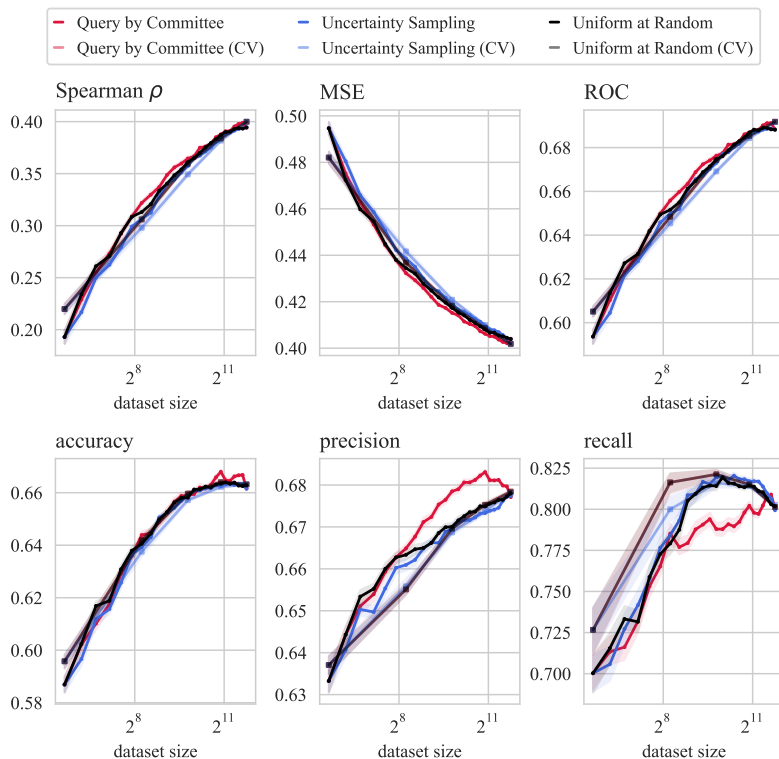


Figure 6: Comparison of point-based strategies without cross-validation (darker colors, circular markers) versus the same strategies, with three-fold cross-validation run at each incremental dataset size (lighter colors, square markers).