



# Interpreting wealth distribution via poverty map inference using multimodal data

Lisette Espín-Noboa  
EspinL@ceu.edu  
Central European University  
Complexity Science Hub Vienna

János Kertész  
KerteszJ@ceu.edu  
Central European University  
Complexity Science Hub Vienna

Márton Karsai  
KarsaiM@ceu.edu  
Central European University  
Rényi Institute of Mathematics

## ABSTRACT

Poverty maps are essential tools for governments and NGOs to track socioeconomic changes and adequately allocate infrastructure and services in places in need. Sensor and online crowd-sourced data combined with machine learning methods have provided a recent breakthrough in poverty map inference. However, these methods do not capture local wealth fluctuations, and are not optimized to produce accountable results that guarantee accurate predictions to all sub-populations. Here, we propose a pipeline of machine learning models to infer *the mean and standard deviation of wealth* across multiple geographically clustered populated places, and illustrate their performance in Sierra Leone and Uganda. These models leverage seven independent and freely available feature sources based on satellite *images*, and *metadata* collected via online crowd-sourcing and social media. Our models show that combined metadata features are the best predictors of wealth in rural areas, outperforming image-based models, which are the best for predicting the highest wealth quintiles. Our results recover the local mean and variation of wealth, and correctly capture the positive yet non-monotonous correlation between them. We further demonstrate the capabilities and limitations of model transfer across countries and the effects of data recency and other biases. Our methodology provides open tools to build towards more transparent and interpretable models to help governments and NGOs to make informed decisions based on data availability, urbanization level, and poverty thresholds.

## CCS CONCEPTS

• **Computing methodologies** → *Neural networks*; • **Applied computing** → *Forecasting*; • **Information systems** → *Geographic information systems*; Location based services.

## KEYWORDS

high-resolution spatial inference, poverty maps, machine learning, deep learning, satellite images, online crowd-sourced data

## ACM Reference Format:

Lisette Espín-Noboa, János Kertész, and Márton Karsai. 2023. Interpreting wealth distribution via poverty map inference using multimodal data. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, April 30–May 04, 2023, Austin, TX, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3543507.3583862>



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

WWW '23, April 30–May 04, 2023, Austin, TX, USA  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-9416-1/23/04.  
<https://doi.org/10.1145/3543507.3583862>

## 1 INTRODUCTION

The first Sustainable Development Goal (SDG) set by the United Nations is to eradicate poverty by 2030 [47]. Although fewer people were living in extreme poverty around the world by 2018, the decline in poverty rates has slowed down ever since. This stagnation was partly due to the COVID-19 pandemic, but the ongoing impact of political turmoils, wars, and climate catastrophes set further barriers for progress in this direction [9]. Traditional data collection techniques fail to follow the effects of such rapidly changing circumstances, therefore, new data collection and analysis techniques are required. Further, the identification of places in need requires rapid, flexible and precise inference to inform the adequate allocation of resources, which are often misplaced due to coarse-grained and outdated statistics provided by census and survey data.

The fast penetration of mobile phones [4, 5] has enabled the collection and use of big data for social good. For instance, mobile phone call detailed records (CDR) and mobile airtime payment transactions have been used to infer several socioeconomic indicators [15, 18, 25], which in turn have been applied to map socioeconomic effects on the structure and dynamics of the underlying social networks [19, 40, 41]. One main limitation of CDRs, however, is that mobile phone data is proprietary and access is often granted via purchase or partnerships. On the contrary, the emergence of Web 2.0 technologies has opened new avenues for collecting and sharing online annotations and digital traces [35, 59]. Further, open data initiatives have strengthened collaborations between industry, academia, and the public sector [44], which have made it easier to study social, economic, and environmental issues through the estimation of socioeconomic indicators from online data [1, 3, 6, 42].

In recent years, satellite imagery has attracted great attention as a mean of inferring high-resolution poverty maps. As a first attempt, nightlight intensity of places has been shown to be a good estimator of economic activity [26], especially for non-extreme poor areas [54]. Moreover, they can even capture household consumption of the extreme poor when combined with daylight satellite images [36]. Other approaches leverage the visual qualities of street-level imagery and daylight satellite images to detect objects (e.g., vehicles, infrastructure, and terrain) and use them as proxies of wealth of neighborhoods [2, 7, 8]. In principle, any geospatial dataset can be used to build poverty maps from their spatial correlations with socioeconomic indicators. For instance, advanced deep learning methods trained on multiple features (e.g., satellite images together with population density maps) provide scalable and improved predictions of wealth at high-resolution for low- and middle-income countries around the globe [14, 39].

While these methods represent a great advance towards scalable, time-variant and fine-grained poverty maps, they commonly optimize performance over representativeness, and depend strongly on the availability of all data sources. As a result, they might perform

poorly in countries with scarce data and fall short in accuracy guarantees necessary for policy makers. In addition, most of these methods use machine learning algorithms with good performance conditional to specific data engineering and parametrization choices. This limits their generalization and transferability potential, especially in countries in emergency undergoing rapid demographic and environmental changes. Although the parametrization of these models is well documented, their robustness against data scarcity, skewness, and time is usually not evaluated or not transparent.

Here, we bridge this gap by performing a systematic study to better understand the capabilities of existing methods and feature sources, which to date have mostly been studied in isolation. Further, we propose regression models that combine features from seven freely available sources to show their overall performance and their strength at the intersection between socioeconomic classes and types of settlements. We train our models to infer not only the mean wealth of places, but also their standard deviation to give a more precise view of how wealth is distributed across households within each populated place. In line with these objectives, we frame our analysis around the following research questions.

- RQ1:** What type of model and features are best at predicting both the mean and standard deviation of wealth?
- RQ2:** Is the goodness-of-fit of models consistent across types of settlements and socioeconomic classes?
- RQ3:** How broadly is wealth distributed in each place? Is the best model able to capture the correlation between the mean and standard deviation of wealth?
- RQ4:** What is the trade-off between data availability and geographical model transfer?

We showcase performance in Sierra Leone (SL) and Uganda (UG), two Sub-Saharan African countries characterized by extreme poverty [9, 47]. Using the last two household surveys conducted by the Demographics and Health Program (DHS) [56], we compute the international wealth index (IWI) [64] of localized population clusters in these countries. In addition, we extract 173 metadata and 784 image-based features for each location, at no cost, from four data providers: Google, Meta (Facebook), OpenStreetMap, and OpenCellID. The image features refer to embeddings extracted from daylight satellite images while the metadata features include demographics, mobility, population density, nightlight intensity, mobile communication antennas, and ground infrastructure. We found that inference performance varies widely across models, feature sources, and geography suggesting that wealth and poverty have different characteristics across different countries. For instance, population density is the best predictor in SL while nightlight intensity is the best in UG. In both cases, however, mobility features are the second best predictors. Moreover, the combination of all metadata-based features outperforms the image-based and individual-source features in both countries. Interestingly, geographical model transferability pays-off when features are partially missing. Finally, while no model predicts wealth equally well at the intersection between socioeconomic classes (e.g., poor, rich) and types of settlements (i.e., urban, rural), we found that, there is a potential in combining multiple models to increase performance.

Our analysis sheds light on how to use sensor and online crowd-sourced data to reliably infer high-resolution poverty maps. To ensure transparency and reproducibility, we make our code and results openly available [20], and share an online visualization tool [69] to

interactively demonstrate that our inferred high-resolution poverty maps can be used to identify places in need.

## 2 RELATED WORK

Closest to our work, [14, 39] combine multiple feature sources to predict high-resolution poverty maps using wealth as a proxy of poverty. In terms of model architecture, we build upon [36, 39] by adding more features into the pipeline and combining two regressor models, one based on images and the other based on metadata (tabular data). Note that the advantage of having multiple features is that the limitations of some sources can be overcome by the strengths of others. For example, models can learn correlations between features characterizing the same location but coming from different sources [60]. In terms of analysis, only [14] goes beyond the overall performance of the model and assess feature importance, and measures the goodness-of-fit of geographical model transferability and different sampling methods. Compared to this work, we additionally show under which circumstances geographical model transferability pays-off given the available features. Further, we compare the performance of multiple models and show their weaknesses and strengths at the intersection between socioeconomic classes and types of settlements. Importantly, unlike in any other work, we infer not only the mean but the standard deviation of the distribution of wealth in populated places; in this way we provide information about the level of socioeconomic diversity at each location. For a profound review on the potential of combining new kinds of data with artificial intelligence (or machine learning algorithms) to achieve one or multiple SDGs, see [17, 25, 30–33, 45, 53].

## 3 METHODS

### 3.1 Ground-truth data

We use two types of surveys available through the Demographic and Health Surveys Program [56]: the Standard Demographic and Health Survey (DHS) and the Malaria Indicator Survey (MIS). Both are nationally-representative population-based surveys conducted at a household level. In particular, we focus on the housing characteristics questionnaire, which helps estimate the wealth of a household by considering the quality and quantity of available facilities or assets at home. In these surveys, household respondents are anonymized and assigned to a geo-located *cluster* from which we can compute mean wealth values and their respective standard deviations. For simplicity, we refer to both surveys as DHS surveys.

**Cluster locations.** Clusters are typically census enumeration areas selected with probabilities proportional to the size within each stratum [57], and often contain about 25-30 households. Among other attributes, clusters possess a type of settlement (urban/rural) and are geo-located with added noise to protect the exact location of respondents [55]. The displacement is random and varies according to the type of settlement: Up to 2 Km for urban and up to 5 Km for rural clusters. A further 1% of places in the latter group are displaced by a maximum of 10 Km. Note that most DHS clusters are located in rural areas, 66% in SL and 76% in UG, see Table A1.

**International wealth index (IWI).** The IWI score is a comparable asset-based index of household's material well-being, or economic status, that can be used for all low and middle income countries [64]. The IWI score for each household is computed using the answers of 10 questions taken from the housing characteristics or living standards questionnaire. Then, asset weights are derived using principal

component analysis (PCA) [34] and re-scaled to achieve IWI scores between 0 and 100. Households with IWI=0 have no assets and possess the lowest quality housing. In contrast, households with IWI=100 represent the richest end of the spectrum. The advantage of IWI scores over other wealth indices provided by the DHS, is that IWI scores are comparable across countries and years.

### 3.2 Populated Places

In order to obtain a high-resolution poverty map of a given country, we collect all its populated places and use them as target locations to infer their IWI scores with a selected model. Populated places are all cities, towns, neighborhoods, villages, hamlets and isolated dwellings that exist in OpenStreetMap (OSM) [51, 52], an online crowd-sourced platform that provides annotations or metadata of the entire world. As in the ground-truth data, most populated places are located in rural areas, 96% in SL and 99% in UG, see Table A1.

### 3.3 Features

For each location, given by the centroid of each cluster and populated place, we extract 957 features from 7 data sources described below (see Table A2 and [20] for technical details). The default bounding-box width or radius used to query all features within a location is 1 mile ( $\approx 1.6$  Km). In some cases, we also query features within 2.0, 5.0, and 10.0 Km to capture the original location of clusters, see Section 3.1. All features were extracted in June 2022.

**Type of settlement.** Rural areas often host the poorest population of a country, and despite being characterized by their relative abundance of natural resources, they lack good quality of services and infrastructure [67]. Thereby, we distinguish each cluster according to the urban and rural divide by using a flag provided by the DHS. In the case of populated places, we follow OSM’s standards [51] and define urban places as cities, towns, and neighborhoods, and rural places as villages, hamlets, and isolated dwellings.

**Daylight satellite images.** We download the daylight satellite image of each place using the Google Maps Static API [29]. These images are recorded in RGB bands with the resolution of 640x640 pixels corresponding to  $\approx 2.5$  m per pixel. We normalize all pixels, and remove the logo and copyright label to prevent erroneous generalizations.<sup>1</sup> The result image covers an area of  $\approx 1.6 \times 1.6$  Km<sup>2</sup>.

**Population.** High-resolution population density maps are obtained through the “Data for Good” platform by Meta [21]. In a nutshell, using state-of-the-art computer vision techniques, these maps estimate the number of people living within 30 m grid tiles in nearly every country around the world [66]. Using these maps, we build 9 *population-based features*, see Table A2. The “gravitational” features are motivated by the literature on population dynamics [37], and the selected  $\beta$  exponents are typical values found in real-world datasets [10]. We use the “closest tile” information and “total population” within different radii for each location to overcome the different resolutions between our locations and the population tiles.

**Mobility.** Movements between tiles are also obtained through Meta’s “Data for Good” platform [43]. These maps aggregate movement counts per day in three 8-hour windows to provide quick and localized responses after any natural disaster, and health crisis. The data provides the number of people moving between tiles at a

given day and time-window as well as a *baseline* reflecting the average movement between the same tiles, weekday, and time-window before the crisis. The resolution of these tiles is roughly 600x600 m<sup>2</sup> at the equator and they do not necessarily map the population tiles. We use the baselines from the COVID-19 maps, and generate mobility networks where every tile is a node, and every movement is an edge in the graph. In SL (UG) we found 544 (2.5K) nodes, 2.9K (10.8K) edges, and 238K (1.2M) movements. First, we assign to each location the closest tile, and then derive 27 *mobility-based features* (see Table A2). Note that while all these features align with previous findings on the relationship between the wealth in Sub-Saharan African countries and mobility patterns (frequency, distances, modalities, and purpose) [12], the “movements between tiles” dataset has not been used earlier to predict poverty maps.

**Demographics.** Using the Marketing API by Meta [24], we extract the number of monthly active users (MAU) on Facebook. Precisely, for each location we make 37 independent queries to obtain the number of MAU that match certain demographics or characteristics and whose home location lies within a radius of 1.6 Km. Motivated by previous work [22, 23, 27], the idea is to use these 37 *demographic-based features* as proxies of wealth, see Table A2. For instance, one hypothesis may claim that *the more smartphone and tablet owners in the area, the wealthier the users living in there*.

**Infrastructure.** Similar to the daylight satellite images, the infrastructure features capture the development of a place. Their main difference is that while infrastructure features are collected via online crowd-sourcing, satellite images are remotely sensed. Further, infrastructure features are computationally easier to handle, however, they often provide outdated or no information about unpopular places. Here, we collect 54 *infrastructure-based features* from OSM [46, 50, 52], e.g., number of buildings or ATMs within a 1.6x1.6 Km<sup>2</sup> bounding box around each location, and the distance to the closest road or a particular point of interest, see Table A2. The advantage of the “closest distance” features (without bounding boxes) is to help mitigating the missing data and unpopularity issues (e.g., the poorer the area, the longer the distance to the closest infrastructure). Note that although OSM features can be queried with a specific date, this date does not necessarily reflect the time since the feature has been physically available; instead, it shows its *last modified date* in the platform. Thus, we omit *time* in the query.

**Connectivity.** Another important proxy of wealth is provided by the number of antennas within an area, and the distance to the closest tower. For each location, we construct 9 *connectivity-based features* extracted from OpenCellID [38], see Table A2. As before, the *distance* features help us mitigate missing values.

**Nightlight intensity.** The intensity of light at night has been shown to be a good predictor of wealth, especially in wealthy areas [7, 36, 39, 54]. We leverage these correlations and construct 36 *nightlight-based features* for each location using the monthly average radiance composite images extracted from Google Earth Engine [13, 28]. In particular, we measure 9 statistics within 4 different buffers, see Table A2. We apply a mask to all images and retrieve all pixels whose average radiance is greater than or equal to a threshold,  $\text{rad} \geq 10$ , and compute the respective fraction of pixels, area, and radiance that fulfill that criteria. Though this threshold is arbitrary, it serves as a proxy to exclude empty areas. Finally, we rank all pixels based on their radiance and compute the mean among the top-30% and lowest-30%. Note that this is the only data

<sup>1</sup>We crop the image by 620x620 pixels starting from the top-left corner.



to 2019 and 2018 for SL and UG, respectively. Additionally, the most recent surveys only represent a few clusters, which might not be sufficient to train generalizable machine learning models. To mitigate the trade-off between sample size and data recency, we first evaluate our models on different temporal configurations of training and testing data using the latest two surveys as follows. (i) O-O: trains and tests on 2016, the *oldest* data. (ii) N-N: as before but on the *newest* data, 2018 in UG and 2019 in SL. (iii) O-N: trains on the *oldest* data, and tests on the *newest*. (iv) ON: combines both years in the train and test sets. We apply stratified sampling in (i), (ii), and (iv), see Section 3.5. Figure B1 (a) and (e) show the results of recency for SL and UG, respectively. We found that the best performance for predicting the mean IWI in both countries is achieved by ON, while N-N and O-O are best for predicting the standard deviation in SL and UG, respectively, followed by ON. Thus, we choose ON as a baseline for the upcoming experiments since overall provides the best performance in both countries with a larger sample. For demonstration, we developed an interactive tool [69] to compare the inferred IWI scores at two different points in time.

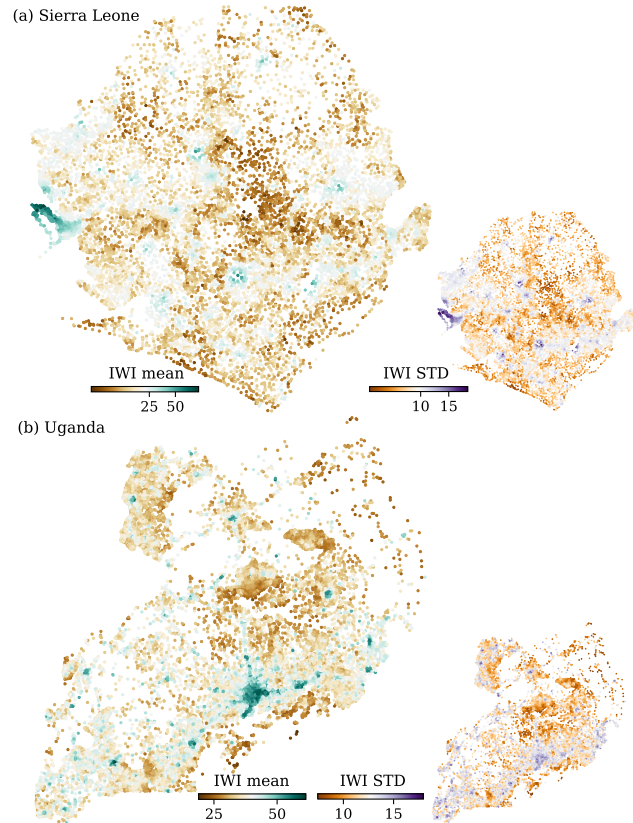
**Relocation.** Geo-located ground-truth data is anonymized by added noise to each cluster location to preserve the confidentiality of survey respondents [55]. While this is fundamental for ethical reasons, it induces uncertainty in the predictions. Previous work addresses this issue by either covering bigger areas around each cluster [14], or re-arranging locations in an iterative way while also adjusting their wealth [39]. While these options are plausible, it is unclear by how much they affect the prediction. Here, we move the noisy location of a cluster to the location of the closest populated place (obtained from OSM) without altering its IWI score. In case multiple clusters are assigned to the same populated place, we prioritize the cluster with fewer other potential matches. We repeat this iteratively until all clusters are re-arranged. We test two alternatives of this algorithm: (i) rc: where only rural clusters are relocated to their closest rural populated place, and (ii) ruc: where both rural and urban clusters are relocated accordingly. Figure B1 (b) and (f) show the results of relocation. In UG it is best to keep the noisy locations, while in SL performance improves when all or only rural clusters are relocated, though the gain compared to no relocation is minor. Thus, we keep the noisy locations for the upcoming experiments.

**Sample size.** The performance of deep neural networks often improves with larger amounts of training data, and the  $\approx 1000$  data points obtained by adding up the last two available surveys (see Table A1), might still not be enough. Thus, we apply offline augmentation during training to increase the training samples by 500%. In particular, for each cluster in the image-based model (see Section 3.4), we generate 5 new images as modifications of the original satellite image by applying different augmentation techniques.<sup>3</sup> Unfortunately, our experiments suggest that augmentation adds little to nothing to the performance, see Figure B1 (c) and (g). However, this can be improved by considering more sophisticated techniques [49, 62], which we leave for future work.

**Class imbalance.** A further complication comes from imbalanced ground-truth data. In terms of types of settlements, both clusters and populated places are mostly located in rural areas, see Table A1. Wealth is also unequally distributed across clusters (i.e.,  $Gini \approx 0.31$  for both countries<sup>4</sup>), making the poor areas the majority. Recall that

<sup>3</sup>brightness, noise, crop, erase, and rotation or flips.

<sup>4</sup>The Gini coefficient is a measure of income inequality in society [65].



**Figure 2: High-resolution poverty maps. Inferred mean (left) and standard deviation (right) of IWI scores using the best models,  $CB_w$  and  $CB$ , at the populated places of (a) SL and (b) UG, respectively. Maps on the left (right) are color-coded from brown-poor (orange-homogeneous) to green-rich (purple-heterogeneous) centered on their respective means. The capitals of both countries are among the wealthiest and more heterogeneous areas. Also, SL appears poorer than UG, which aligns with the estimated poverty rates by the World Bank [68]: 56.80% in 2018 for SL, and 21.40% in 2016 for UG.**

IWI scores are continuous real values that range from 0 (poorest) to 100 (richest). In order to account for socioeconomic class imbalances we first discretize these scores into 10 equal-width bins. Then, we apply a class-balanced loss based technique [16] to generate a weight for each data point in the training sample and pass these weights to the classifier.<sup>5</sup> We see in Figure B1 (d) and (h) that adding sample weights only benefits the metadata model in SL.

## 4.2 Performance

Our second contribution and main outcome of our methodology is the inferred high-resolution poverty maps of SL and UG, see Figure 2. These maps are explained in more detail, for all investigated models, in our online interactive tool [69].

<sup>5</sup>We control the imbalance issues on wealth and type of settlement separately and in combination using 4 techniques: heuristics, Inverse Number of Samples (INS), Effective Number of Samples (ENS), and `compute_sample_weight` from sklearn in python [63]. Due to space limitations, out of the 20 configurations, we report results using ENS with  $\beta = 0.9$ , the technique with best performance, see [20] for more details.

**Table 1: Performance by model and feature-source.** Each row shows the normalized root-mean-squared error (NRMSE=  $\epsilon$ ) of predicting the mean  $\mu$  and standard deviation  $\sigma$  of IWI scores. The CatBoost model (CB) is trained with metadata features (i.e., Demographics, Mobility, Population, Nightlight intensity, Antennas, and Infrastructure), and the CNN model (CNN) with daylight satellite images. The combined model (CNN+CB) uses all features. Model sub-indexes refer to “weighted samples” (CB<sub>w</sub>) and “offline augmentation” (CNN<sub>a</sub>), two techniques applied during training to correct for wealth imbalance and small samples, respectively. The best performance across models is shown in *blue*, and within each model type in **bold**.

		Features							Sierra Leone		Uganda	
		De	Mo	Po	Nl	An	In	Im	$\epsilon_\mu$	$\epsilon_\sigma$	$\epsilon_\mu$	$\epsilon_\sigma$
Metadata-all	CB	✓	✓	✓	✓	✓	✓	-	0.46	0.81	<b>0.46</b>	<b>0.83</b>
	CB <sub>w</sub>	✓	✓	✓	✓	✓	✓	-	<b>0.44</b>	<b>0.78</b>	0.47	0.85
Metadata-single	CB <sub>w</sub>	✓	-	-	-	-	-	-	0.73	0.94	0.77	0.97
		-	✓	-	-	-	-	-	0.47	0.81	0.55	0.92
		-	-	✓	-	-	-	-	<b>0.45</b>	<b>0.79</b>	0.61	0.90
		-	-	-	✓	-	-	-	0.52	0.83	<b>0.50</b>	<b>0.86</b>
		-	-	-	-	✓	-	-	0.52	0.87	0.56	0.89
Image-only	CNN	-	-	-	-	-	-	✓	<b>0.57</b>	<b>0.90</b>	0.60	0.94
		-	-	-	-	-	-	✓	0.59	0.94	<b>0.59</b>	<b>0.91</b>
		-	-	-	-	-	-	✓	0.51	0.85	0.61	0.89
Combined	CNN+CB	✓	✓	✓	✓	✓	✓	✓	0.45	0.82	0.51	0.86
	CNN+CB <sub>w</sub>	✓	✓	✓	✓	✓	✓	✓	<b>0.45</b>	0.80	0.52	0.86
	CNN <sub>a</sub> +CB	✓	✓	✓	✓	✓	✓	✓	0.46	0.82	<b>0.47</b>	<b>0.83</b>
	CNN <sub>a</sub> +CB <sub>w</sub>	✓	✓	✓	✓	✓	✓	✓	0.45	<b>0.79</b>	0.47	0.85

Our third and last set of contributions addresses our research questions. Here, we disentangle the performance of our models by reporting the normalized root-mean-squared errors ( $NRMSE = \frac{RMSE}{\sigma_{y, true}}$ ) for the inference of both predicted mean and predicted standard deviation (STD). Results are based on the combination of the last two available surveys using noisy locations. For comparison, we report model performance with and without augmentation (CNN<sub>a</sub> vs. CNN), and with and without sample weights (CB<sub>w</sub> vs. CB).

**RQ1: Model selection and feature-source importance.** As shown in Table 1, we find that the *metadata-all* models, which use all non-image features, outperform the other models in terms of predicted mean ( $\epsilon_\mu$ ) and STD ( $\epsilon_\sigma$ ) of IWI scores. In SL, the CB<sub>w</sub> model, that balances the skewness of wealth in the data, provides the best performance, while the unbalanced CB model appears the best in UG (shown in *blue*). Interestingly, including all features in the combined model improves the performance only for UG as compared to the performance of individual feature-sources.

From the *metadata-single* results shown in Table 1, we find that while most of the features can provide relatively good predictive performance, the population (Po) and nightlight intensity (Nl) features alone are good predictors of wealth in SL and UG, respectively. This aligns with previous studies that show that population density [11] and luminosity [36] are strongly correlated with economic growth. Surprisingly, the *image-only* models achieve the worst performance, however, they can be improved by *combining* them with the metadata features.

**RQ2: Intersectionality.** The overall performance shown above is not informative about the intersectionality of the models, e.g., the prediction of wealth in poor-rural areas might not be as accurate as in rich-urban areas. Therefore, we assess the performance of our models at the intersection between types of settlements and socioeconomic classes; both derived from the DHS data. The former

is given, and the latter we infer by transforming the true mean IWI scores into quintiles (i.e., 5 equally-populated bins). Results based on the predicted mean wealth ( $\epsilon_\mu$ ) are shown in Table 2. We see that the CB models, which appeared to be the best in the overall performance (in Table 1), were found to be good only in certain regions. In SL, the lowest quintiles (Q1, Q2) are best captured by the CB model in rural areas, while the middle (Q2, Q3) and rich quintiles (Q4, Q5), regardless of their type of settlement, are best captured by the CNNs and the combined models, respectively. In UG, on the other hand, the richest quintile (Q5) is best captured by the CB model, while the middle quintiles (Q2-Q4) are best explained by the CNNs or combined models everywhere. Moreover, the poorest quintile (Q1) is best captured by the combined and CB models in rural and urban areas, respectively. These results demonstrate that there is no model that fits very well all types of settlements and wealth quintiles at the same time; thus, it is important to systematically explore them across different sub-populations.

**RQ3: Wealth variability.** When predicting aggregated economic indicators, mean values are often not enough to understand the distribution of wealth across households. Thus, by predicting the STD, and understanding its relation with the mean, we gain additional insights about the socioeconomic diversity of individuals living in the same populated places. From the results shown in Figure 3, we find that the best inference model in each country predicts well—with slight overestimation—the minimum mean wealth as compared to the true values (orange curves in Q1); crucial for policy making when targeting places in need. However, the inference somewhat underestimates the mean wealth at the richest end of the spectrum (Q5) as the models miss to capture the wealth of outlier rich locations; especially in urban areas (blue curves).

Meanwhile, we also find a positive and significant correlation between the mean and STD of wealth for both true and predicted

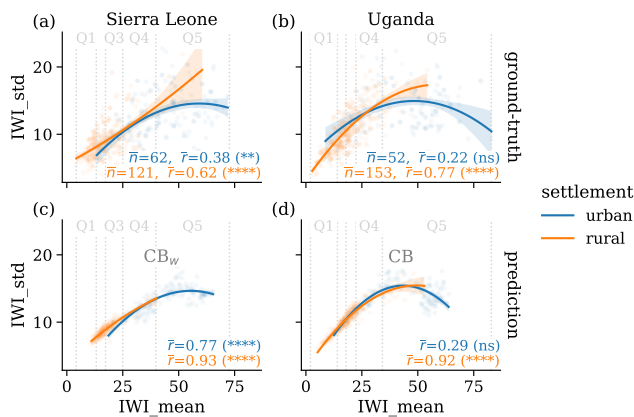
**Table 2: Performance by intersectionality.** For each country we divide the true mean IWI scores into quintiles and measure the root-mean-squared error (RMSE) of predicting the mean wealth ( $\epsilon_\mu$ ) at the intersection between countries, types of settlements, quintiles, and models. Bold values refer to the best performance in each quintile and type of settlement.

		Sierra Leone					Uganda				
		Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Q4	Q5
Rural	CB	<b>5.48</b>	<b>3.18</b>	5.77	10.46	23.20	7.01	<b>4.06</b>	3.63	8.04	<b>12.54</b>
	CB <sub>w</sub>	7.64	4.83	5.18	10.18	22.47	8.83	5.08	5.06	8.47	15.48
	CNN	7.72	4.48	<b>4.67</b>	10.59	22.56	9.84	5.51	4.87	7.55	16.47
	CNN <sub>a</sub>	8.96	5.07	4.80	10.30	23.41	9.66	5.11	4.61	<b>7.44</b>	16.72
	CNN+CB	6.24	3.49	4.78	9.81	<b>20.18</b>	7.20	4.49	4.26	7.67	15.07
	CNN+CB <sub>w</sub>	6.23	3.66	4.75	<b>9.74</b>	21.36	7.67	4.51	4.22	7.59	15.07
	CNN <sub>a</sub> +CB	6.74	3.38	5.03	10.55	22.64	<b>6.67</b>	4.26	<b>3.35</b>	7.54	13.78
	CNN <sub>a</sub> +CB <sub>w</sub>	6.55	3.45	5.02	10.54	22.11	6.87	4.15	3.59	7.59	13.84
Urban	CB	-	18.07	8.13	6.18	9.69	<b>6.42</b>	6.96	10.08	6.77	<b>9.33</b>
	CB <sub>w</sub>	-	10.30	8.16	8.67	11.34	11.70	7.38	12.19	9.68	12.83
	CNN	-	<b>4.54</b>	2.83	11.03	12.84	9.77	5.41	11.95	7.77	13.38
	CNN <sub>a</sub>	-	6.77	<b>2.21</b>	10.96	13.23	10.11	<b>5.10</b>	13.94	8.34	13.18
	CNN+CB	-	15.48	4.73	6.70	9.62	9.02	5.99	10.63	7.71	10.97
	CNN+CB <sub>w</sub>	-	14.15	6.53	6.77	9.33	8.66	7.13	10.96	8.13	11.08
	CNN <sub>a</sub> +CB	-	16.17	6.84	5.82	9.48	10.45	5.77	10.03	<b>6.67</b>	9.67
	CNN <sub>a</sub> +CB <sub>w</sub>	-	15.99	8.06	<b>5.75</b>	<b>9.31</b>	8.86	6.54	<b>9.86</b>	6.72	9.70

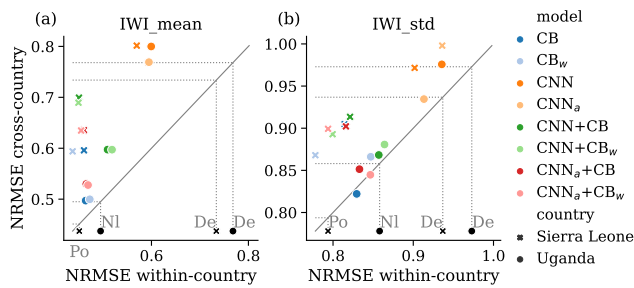
values; especially in rural areas (orange curves). This trend indicates that larger socioeconomic variation characterizes wealthier places. However, in urban areas, this correlation tends to follow a non-monotonic pattern. While the STD appears to be the smallest at the poorest places, after reaching a maximum for middle class locations, it gets smaller for the richest areas. This means that while middle class people tend to live in places with higher socioeconomic variations, the poorest and the richest live in more homogeneous

communities, signaling their segregation from the rest of the population. In Figure B2, we further confirm the positive correlation between population size and inferred wealth in populated places; aligned with our feature importance analysis. Moreover we find a strong distinction between rural and urban areas, the latter appearing more populated and wealthier, yet depicting the already mentioned segregated mixing patterns, especially in SL.

**RQ4: Cross-country model transferability.** We test the transferability of our models by training them in one country to predict the wealth in the other without further training. Figure 4 shows the performance of the models in terms of normalized root-mean-squared errors (NRMSE) for both within country (x-axis) and cross-country (y-axis) validations. We found three main patterns: (i) Model transferability across countries achieves higher errors compared to the within-country counterpart, i.e., most data points are above the diagonal. Although this was expected, CB models provide surprisingly good performance in transferability, landing close to the diagonal in many cases. This can be due to the several metadata sources these models rely on since they may compensate data quality differences across layers. An interesting exception is for SL, where the STD is somewhat better predicted from a CB model trained on UG than the corresponding model trained on SL, see the dark blue dot below the diagonal in Figure 4 (b). (ii) Model transferability is asymmetric, i.e., training on SL to predict UG leads to smaller errors than vice versa. This may be due to the wider range of mean wealth in UG compared to SL, i.e., 99% of clusters in UG fall within the distribution of mean wealth in SL, see Table A1. Thus, models trained in UG may infer values that fall off the wealth range characterizing SL, resulting in larger inference errors. (iii) There is a trade-off between model transferability and feature availability when only one feature source is present. Naturally, model transferability pays-off if no feature is available. However, if only the best metadata-single features are available, i.e., nightlight intensity (NI) in UG and population



**Figure 3: Wealth variability.** Relationship between the mean (x-axis) and standard deviation (y-axis) of wealth in SL and UG (columns) for urban (blue) and rural (orange) clusters on the test sets. True values are shown at the top, (a) and (b), and predicted values using the best model in each country are at the bottom, (c) and (d). Curves are 2<sup>nd</sup> order polynomial fits, and labels show the average # of places  $\bar{n}$ , and the mean Pearson correlation  $\bar{r}$  between the two predicted variables.



**Figure 4: Cross-country model performance.** We test model transferability by using the models trained in one country to predict the wealth in the other without further training. For each response variable (columns a and b), model (color), and country (markers), we show model performance in terms of normalized root-mean-squared errors (NRMSE) for both within country (x-axis) and cross-country (y-axis) validations for SL (x) and UG (-). Dotted lines with labels indicate the within-country metadata-single performance for population (Po), nightlight intensity (NI), and demographics (De).

(Po) in SL (leftmost dotted lines in Figure 4), then model transfer should be avoided since the transferred model’s performance is worse than using the single available local feature source; except for  $CNN_a+CB_w$  and CB when inferring STD in UG. On the other hand, if only the worst metadata-single features are available, i.e., demographics (De) (rightmost dotted lines in Figure 4), then model transferability yields an advantage for almost all models.

## 5 LIMITATIONS AND FUTURE WORK

The main limitations of our methods are rooted in multiple aspects. First, the recency, quality, coverage, and biases of the different data sources that we use may hinder the performance of the inference task. We discussed in detail the effects of these issues and showed how to address them. Meanwhile, the multiple layers of collected features seem to account for the effects of low data quality in certain layers, making our models successful in predicting wealth even in data-sparse regions. Second, even though our models achieve high performance, comparable with previous work [14, 39], we cannot assume the poverty maps to be the possibly most precise representation of actual wealth. As we argued, data recency and sample size are important aspects to consider during training. Moreover, the timestamps of most features do not match the timestamps of the ground-truth data; the former cannot be queried retrospectively in time. Thus, more recent ground-truth data is required to infer up-to-date poverty maps. Third, the current implementation of our methodology most reliably works for low- and middle-income countries since it relies on DHS survey data. Nevertheless, our code and results are openly available [20], thus, new types of ground-truth data and features can be easily plugged into our framework, e.g., housing prices, mobile airtime payments, or street view imagery. Similarly, additional methods can be incorporated into our pipeline to extend the explanations of our results. For example, activation maps [2] can be used to understand the under-performance of the image-only models in poor areas, and geographical stratification [14] could guarantee a better spatial coverage during training in addition to the stratification of wealth (cf. Section 3.5).

Note that, while this and similar methodologies have been implemented for a social good, they may also be misused, e.g., for commercial purposes which may increase the inequality gap between the rich and the poor. One way to mitigate this issue is by implementing better synergies between data providers and analysts to keep track of whom is using the data and for what purposes.

In the future, we aim to tackle these limitations and explore the spatial and temporal transferability of single- and multi-source models with missing data layers at the target population. Further, we will consider combining the best of each model to produce more stable and consistent results across different sub-populations.

## 6 CONCLUSION

In this work, we showed how to produce high-resolution poverty maps to interpret wealth distributions using multimodal data. We proposed machine learning models that learn to predict the mean and standard deviation of wealth in places from low- and middle income countries, and showcased their capabilities in Sierra Leone and Uganda. These models rely on a multimodal architecture composed by sensor and online crowd-sourced data. Moreover, we systematically studied four major issues namely data recency, noisy locations, small training samples, and class imbalance, and conducted an ablation study to evaluate their performance.

Our findings show that: (1) The combined data features outperform satellite images in solving the inference task. Further, population and nightlight intensity features are the strongest predictors of wealth in SL and UG, respectively, followed by the mobility features in both countries. (2) There is no single model or feature source that can predict wealth very accurately at the intersection between all socioeconomic classes and types of settlements. However, in general, the poorest areas are best inferred by the CB models using metadata-only features, while wealthier areas are best inferred by the CNN or combined models using daylight satellite images. (3) Our models capture well the mean and standard deviation of wealth, their correlation, and non-monotonous dependencies. This way they provide information about the socioeconomic heterogeneity at given places and reflect segregation between different classes. (4) There is an asymmetric trade-off between geographical transferability and data availability where countries with sparse information may benefit from reusing models trained in other countries.

To conclude, our results shed light on the importance of producing accountable models to guarantee accurate predictions across all types of settlements and socioeconomic groups. Finally, beyond our scientific results, we made our code and results openly available [20], and built an online interactive visualization tool [69] to help identifying places in need for the development of fair policy and intervention designs.

## ACKNOWLEDGMENTS

We thank Liuhuaying Yang for developing the visualization tool, and Gilles Vandewiele and Philipp Singer for their feedback. We are also grateful for the support of the Data for Good Team at Meta, the SoBigData++ project (H2020-871042), Dataredux ANR project (ANR-19-CE46-0008), and the CHIST-ERA project SAI: FWF I 5205-N. This material is based upon work supported by the Google Cloud Research Credits program with the award GCP19980904. The presented computational results have been achieved in part using the Vienna Scientific Cluster (VSC).



## REFERENCES

- [1] JL Abitbol and AJ Morales. 2021. Socioeconomic Patterns of Twitter User Activity. *Entropy (Basel, Switzerland)* 23, 6 (2021).
- [2] Jacob Levy Abitbol and Márton Karsai. 2020. Interpretable socioeconomic status inference from aerial imagery through urban patterns. *Nature Machine Intelligence* 2, 11 (2020), 684–692.
- [3] Jacob Levy Abitbol, Márton Karsai, and Eric Fleury. 2018. Location, occupation, and semantics based socioeconomic status inference on twitter. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 1192–1199.
- [4] Lishan Adam and Michael Minges. 2018. ICTs, LDCs and the SDGs: Achieving Universal and Affordable Internet in the Least Developed Countries. , 128 pages.
- [5] Areppim AG. 2017. Mobile Phone Market Forecast. [https://stats.areppim.com/stats/stats\\_mobilex2017.htm](https://stats.areppim.com/stats/stats_mobilex2017.htm). Accessed: 2022-02-02.
- [6] Nikolaos Aletras and Benjamin Paul Chamberlain. 2018. Predicting Twitter User Socioeconomic Attributes with Network and Language Information. In *Proceedings of the 29th on Hypertext and Social Media (Baltimore, MD, USA) (HT18)*. Association for Computing Machinery, New York, NY, USA, 20–24. <https://doi.org/10.1145/3209542.3209577>
- [7] Kumar Ayush, Burak Uz kent, Marshall Burke, David Lobell, and Stefano Ermon. 2021. Generating Interpretable Poverty Maps Using Object Detection in Satellite Images. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (Yokohama, Yokohama, Japan) (IJCAI'20)*. Article 608, 7 pages.
- [8] Kumar Ayush, Burak Uz kent, Kumar Tanmay, Marshall Burke, David Lobell, and Stefano Ermon. 2021. Efficient Poverty Mapping from High Resolution Remote Sensing Images. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 1 (May 2021), 12–20. <https://ojs.aaai.org/index.php/AAAI/article/view/16072>
- [9] World Bank. 2018. Decline of Global Extreme Poverty Continues but Has Slowed: World Bank. <https://www.worldbank.org/en/news/press-release/2018/09/19/decline-of-global-extreme-poverty-continues-but-has-slowed-world-bank>. Accessed: 2022-02-02.
- [10] Marc Barthélemy. 2011. Spatial networks. *Physics reports* 499, 1-3 (2011), 1–101.
- [11] Luis MA Bettencourt. 2013. The origins of scaling in cities. *science* 340, 6139 (2013), 1438–1441.
- [12] Deborah F Bryceson, Tatenda C Mbara, and David Maunder. 2003. Livelihoods, daily mobility and poverty in sub-Saharan Africa. *Transport reviews* 23, 2 (2003), 177–196.
- [13] Earth Engine Data Catalog. 2022. VIIRS Stray Light Corrected Nighttime Day/Night Band Composites Version 1. [https://developers.google.com/earth-engine/datasets/catalog/NOAA\\_VIIRS\\_DNB\\_MONTHLY\\_V1\\_VCMSLFCG](https://developers.google.com/earth-engine/datasets/catalog/NOAA_VIIRS_DNB_MONTHLY_V1_VCMSLFCG). Accessed: 2022-05-03.
- [14] Guanghua Chi, Han Fang, Sourav Chatterjee, and Joshua E. Blumenthal. 2022. Microestimates of wealth for all low- and middle-income countries. *Proceedings of the National Academy of Sciences* 119, 3 (2022). <https://doi.org/10.1073/pnas.2113658119>
- [15] Eduardo Cruz, Carmen Vaca, and Monica Villavicencio. 2021. Estimating urban socioeconomic inequalities through airborne top-up transactions data. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 4265–4272.
- [16] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9268–9277.
- [17] Assunta Di Vaio, Rosa Palladino, Rohail Hassan, and Octavio Escobar. 2020. Artificial intelligence and business models in the sustainable development goals perspective: A systematic literature review. *Journal of Business Research* 121 (2020), 283–314.
- [18] Yuxiao Dong, Yang Yang, Jie Tang, Yang Yang, and Nitesh V Chawla. 2014. Inferring user demographics and social strategies in mobile social networks. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 15–24.
- [19] Nathan Eagle, Michael Macy, and Rob Claxton. 2010. Network diversity and economic development. *Science* 328, 5981 (2010), 1029–1031.
- [20] Lisette Espin-Noboa. 2023. PovertyMaps. <https://github.com/lisette-espino/PovertyMaps>. Accessed: 2023-02-10.
- [21] Facebook Connectivity Lab and Center for International Earth Science Information Network, CIESIN, Columbia University. 2016. High Resolution Settlement Layer (HRSL). Source imagery for HRSL © 2016 DigitalGlobe. Accessed:08-07-2021 (SL), 10-07-2021 (UG).
- [22] Masoomali Fatehikia, Benjamin Coles, Ferda Ofli, and Ingmar Weber. 2020. The relative value of facebook advertising data for poverty mapping. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 934–938.
- [23] Masoomali Fatehikia, Isabelle Tingzon, Ardie Orden, Stephanie Sy, Vedran Sekara, Manuel Garcia-Herranz, and Ingmar Weber. 2020. Mapping socioeconomic indicators using social media advertising data. *EPJ Data Science* 9, 1 (2020), 22.
- [24] Meta for Developers. 2022. Reach Estimate. <https://developers.facebook.com/docs/marketing-api/reference/reach-estimate/>. Accessed: 2022-02-02.
- [25] Jian Gao, Yi-Cheng Zhang, and Tao Zhou. 2019. Computational socioeconomics. *Physics Reports* 817 (2019), 1–104.
- [26] Tilottama Ghosh, Rebecca L Powell, Christopher D Elvidge, Kimberly E Baugh, Paul C Sutton, and Sharolyn Anderson. 2010. Shedding light on the global distribution of economic activity. *The Open Geography Journal* 3, 1 (2010).
- [27] Serena Giurgola, Simone Piaggese, Márton Karsai, Yelena Mejova, André Panisson, and Michele Tizzoni. 2021. Mapping urban socioeconomic inequalities in developing countries through Facebook advertising data. *arXiv preprint arXiv:2105.13774* (2021).
- [28] Google. 2022. Earth Engine. <https://earthengine.google.com/>. Accessed: 2022-02-02.
- [29] Google. 2022. The Maps Static API. <https://developers.google.com/maps/documentation/maps-static/overview>. Accessed: 2022-02-02.
- [30] Arash Hajikhani and Arho Suominen. 2022. Mapping the sustainable development goals (SDGs) in science, technology and innovation: application of machine learning in SDG-oriented artefact detection. *Scientometrics* (2022), 1–33.
- [31] Ola Hall, Francis Dompae, Ibrahim Wahab, and Fred Mawunyo Dzanku. 2023. A review of machine learning and satellite imagery for poverty prediction: Implications for development research and applications. *Journal of International Development* (2023).
- [32] Antonio Hidalgo, Samuel Gabaly, Gustavo Morales-Alonso, and Alberto Uruena. 2020. The digital divide in light of sustainable development: An approach through advanced machine learning techniques. *Technological Forecasting and Social Change* 150 (2020), 119754.
- [33] Jacinta Holloway, Kerrie Mengersen, and Kate Helmstedt. 2018. Spatial and machine learning methods of satellite imagery analysis for Sustainable Development Goals. In *Proceedings of the 16th Conference of International Association for Official Statistics (IAOS)*. International Association for Official Statistics (IAOS), 1–14.
- [34] Harold Hotelling. 1933. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology* 24, 6 (1933), 417.
- [35] Andrew Hudson-Smith, Michael Batty, Andrew Crooks, and Richard Milton. 2009. Mapping for the masses: Accessing Web 2.0 through crowdsourcing. *Social science computer review* 27, 4 (2009), 524–538.
- [36] Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353, 6301 (2016), 790–794.
- [37] Mikhail Yurievich Khavinson and Matvey Pavlovich Kulakov. 2017. Gravitational model of population dynamics. *Bulletin of the South Ural State University. Series: Mathematical modeling and programming* 10, 3 (2017), 80–93.
- [38] Unwired Labs. 2021. OpenCellID - The world's largest Open Database of Cell Towers. <https://www.opencellid.org>. Accessed: 2022-02-02.
- [39] Kamwo Lee and Jeanine Braithwaite. 2022. High-resolution poverty maps in Sub-Saharan Africa. *World Development* 159 (2022), 106028. <https://doi.org/10.1016/j.worlddev.2022.106028>
- [40] Yannick Leo, Eric Fleury, J Ignacio Alvarez-Hamelin, Carlos Sarraute, and Márton Karsai. 2016. Socioeconomic correlations and stratification in social-communication networks. *Journal of The Royal Society Interface* 13, 125 (2016), 20160598.
- [41] Yannick Leo, Márton Karsai, Carlos Sarraute, and Eric Fleury. 2018. Correlations and dynamics of consumption patterns in social-economic networks. *Social Network Analysis and Mining* 8, 1 (2018), 1–16.
- [42] Alejandro Llorente, Manuel Garcia-Herranz, Manuel Cebrian, and Esteban Moro. 2015. Social media fingerprints of unemployment. *PLoS one* 10, 5 (2015), e0128692.
- [43] Paige Maas, Shankar Iyer, Andreas Gros, Wonhee Park, Laura McGorman, Chaya Nayak, and P Alex Dow. 2019. Facebook Disaster Maps: Aggregate Insights for Crisis Response & Recovery.. In *KDD*, Vol. 19. 3173.
- [44] Meta. 2022. Data for Good. <https://dataforgood.facebook.com/>. Accessed: 2022-02-02.
- [45] David Mhlanga. 2021. Artificial intelligence in the industry 4.0, and its impact on poverty, innovation, infrastructure development, and the sustainable development goals: Lessons from emerging economies? *Sustainability* 13, 11 (2021), 5788.
- [46] Franz-Benjamin Mocnik, Amin Mobasheri, and Alexander Zipf. 2018. Open Source Data Mining Infrastructure for Exploring and Analysing OpenStreetMap. <https://github.com/mocnik-science/osm-python-tools>. *Open Geospatial Data, Software and Standards* 7, 3 (2018). <https://doi.org/10.1186/s40965-018-0047-6>
- [47] United Nations. 2020. Sustainable Development Goals - Goal 1: End poverty in all its forms everywhere. <https://www.un.org/sustainabledevelopment/poverty/>. Accessed: 2022-02-02.
- [48] Faizaan Naveed. 2019. Satellite Imagery Classification Using Deep Learning. <https://medium.com/data-driven-investor.com/patch-based-cover-type-classification-using-satellite-imagery-a67edeae7e24>. Accessed: 2022-05-04.
- [49] Humza Naveed. 2021. Survey: Image mixing and deleting for data augmentation. *arXiv preprint arXiv:2106.07085* (2021).
- [50] OpenStreetMap. 2021. Infrastructure Elements. Nodes: <https://wiki.openstreetmap.org/wiki/Node>, Ways: <https://wiki.openstreetmap.org/wiki/Way>. Accessed: 2022-02-02.
- [51] OpenStreetMap. 2021. Populated Places. <https://wiki.openstreetmap.org/wiki/Key:place>. Accessed: 2022-02-02.
- [52] OpenStreetMap. 2022. Explorer. <https://www.openstreetmap.org/>. Accessed: 2022-02-02.
- [53] Iván Palomares, Eugenio Martínez-Cámara, Rosana Montes, Pablo García-Moral, Manuel Chiachio, Juan Chiachio, Sergio Alonso, Francisco J Melero, Daniel Molina, Bárbara Fernández, et al. 2021. A panoramic view and swot analysis of artificial intelligence for achieving the sustainable development goals by 2030: Progress and prospects. *Applied Intelligence* 51, 9 (2021), 6497–6527.

- [54] Maxim Pinkovskiy and Xavier Sala-i Martin. 2016. Lights, camera... income! Illuminating the national accounts-household surveys debate. *The Quarterly Journal of Economics* 131, 2 (2016), 579–631.
- [55] The DHS Program. 2022. Analyzing DHS Data. [https://dhsprogram.com/data/Guide-to-DHS-Statistics/Analyzing\\_DHS\\_Data.htm](https://dhsprogram.com/data/Guide-to-DHS-Statistics/Analyzing_DHS_Data.htm). Accessed: 2022-04-12.
- [56] The DHS Program. 2022. The Demographic and Health Surveys (DHS) Program. <https://dhsprogram.com/>. Accessed: 2022-02-02.
- [57] The DHS Program. 2022. GPS Data Collection. <https://dhsprogram.com/methodology/gps-data-collection.cfm>. Accessed: 2022-02-02.
- [58] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. 2018. CatBoost: unbiased boosting with categorical features. *Advances in neural information processing systems* 31 (2018).
- [59] Anat Rafaeli, Shelly Ashtar, and Daniel Altman. 2019. Digital traces: New data, resources, and tools for psychological-science research. *Current directions in psychological science* 28, 6 (2019), 560–566.
- [60] Anil Rahate, Rahee Walambe, Sheela Ramanna, and Ketan Kotecha. 2022. Multi-modal co-learning: challenges, applications with datasets, recent advances and future directions. *Information Fusion* 81 (2022), 203–239.
- [61] Parikshit Ram and Kaushik Sinha. 2019. Revisiting kd-tree for nearest neighbor search. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*. 1378–1388.
- [62] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of big data* 6, 1 (2019), 1–48.
- [63] Sklearn. 2022. Compute Sample Weight. [https://scikit-learn.org/stable/modules/generated/sklearn.utils.class\\_weight.compute\\_sample\\_weight.html](https://scikit-learn.org/stable/modules/generated/sklearn.utils.class_weight.compute_sample_weight.html). Accessed: 2022-10-30.
- [64] Jeroen Smits and Roel Steendijk. 2015. The international wealth index (IWI). *Social indicators research* 122, 1 (2015), 65–85. Accessed: 2022-02-02, Play with IWI: <https://globaldatalab.org/iwi/>.
- [65] StatsDirect. 2000. Gini Coefficient of Inequality. [https://www.statsdirect.com/help/default.htm#nonparametric\\_methods/gini.htm](https://www.statsdirect.com/help/default.htm#nonparametric_methods/gini.htm). Accessed: 2020-11-09.
- [66] Tobias G Tiecke, Xianming Liu, Amy Zhang, Andreas Gros, Nan Li, Gregory Yetman, Talip Kilic, Siobhan Murray, Brian Blankespoor, Espen B Prydz, et al. 2017. Mapping the world population one building at a time. *arXiv preprint arXiv:1712.05839* (2017).
- [67] Steve Wiggins and Sharon Proctor. 2001. How special are rural areas? The economic implications of location for rural development. *Development policy review* 19, 4 (2001), 427–436.
- [68] World Population Review. 2022. Poverty Rate by Country 2022. <https://worldpopulationreview.com/country-rankings/poverty-rate-by-country>. Accessed: 2022-05-12.
- [69] Lihua Yang and Lisette Espin-Noboa. 2023. Interactive Poverty Map Visualization. <https://vis.csh.ac.at/poverty-maps/>. Accessed: 2023-02-10.

## A APPENDIX TABLES

**Table A1: Locations.** Household and cluster data are gathered via the Demographics and Health Survey program (DHS), and populated places from OpenStreetMap (OSM). Rural and urban labels of clusters are directly assigned by the DHS, and for populated places we infer them from their type: *urban* if they are cities, towns or neighborhoods, and *rural* if they are villages, hamlets or isolated dwellings. Statistical significance (\*\*\*) , p-value  $\leq 0.0001$ .

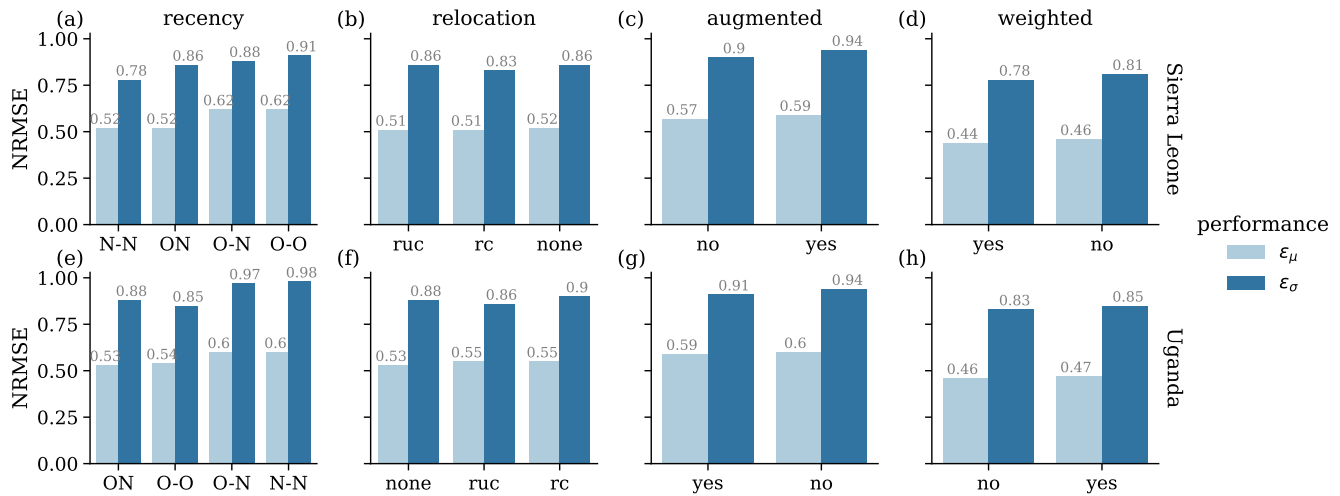
Country	Sierra Leone	Uganda
Years	2016, 2019	2016, 2018
DHS households	19975	27798
<b>DHS clusters</b>	<b>893</b>	<b>1001</b>
- oldest, newest	336, 557	685, 316
- urban	308 (34%)	242 (24%)
- rural	585 (66%)	759 (76%)
- IWI mean ( $\mu$ )		
min, max, mean, std.dev.	4.2, 72.7, 26.2, 15.6	2.0, 82.9, 24.8, 14.9
- IWI std.dev. ( $\sigma$ )		
min, max, mean, std.dev.	3.8, 22.8, 10.4, 3.5	3.9, 22.4, 11.4, 3.4
- Pearson corr. ( $\mu$ and $\sigma$ )	0.71 (***)	0.64 (***)
- Gini IWI mean	32	31
WorldBank Gini index	35.7 (2018)	42.7 (2019)
<b>OSM populated places</b>	<b>9881</b>	<b>27791</b>
- urban	366 (4%)	348 (1%)
- rural	9515 (96%)	27443 (99%)
<b>Relocated DHS cluster</b>	<b>705 (79%)</b>	<b>733 (73%)</b>
- urban	120 (39%)	71 (29%)
- rural	585 (100%)	662 (87%)

**Table A2: Features.** Population, mobility and demographics are datasets provided by Meta. Infrastructure features are extracted from OpenStreetMap. Connectivity features refer to the antennas provided by OpenCellID. Nightlight intensity scores are provided by Google as well as the daylight satellite images. For more technical details see [20].

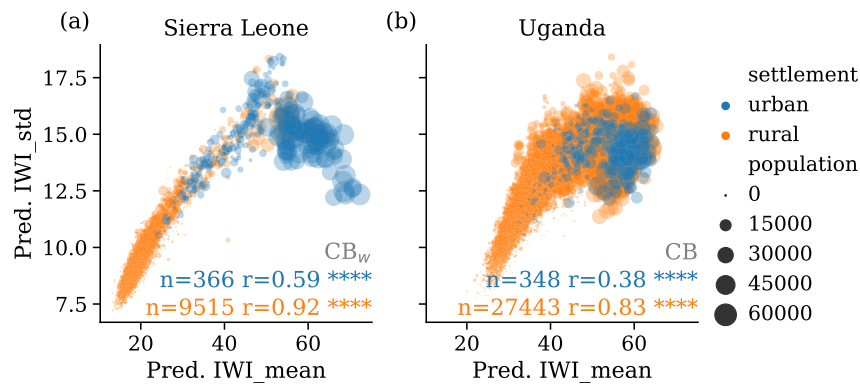
Name	Source	#	Features
Population <sup>†</sup>	Meta	9	distance_to_closest_tile, population_in_closest_tile, total_population_within(radius), and gravitational_closest_tile <sup><math>\beta</math></sup> = population/distance <sup><math>\beta</math></sup> , where $\beta \in \{1.0, 1.5, 2.0\}$ and radius $\in \{1.6, 2.0, 5.0, 10.0\}$ Km.
Mobility <sup>‡</sup>	Meta	27	distance_to_closest_tile, average_distance_in, average_distance_out, people_flow_in <sup><math>\beta</math></sup> , people_flow_out <sup><math>\beta</math></sup> , in_degree <sup><math>\beta</math></sup> , out_degree <sup><math>\beta</math></sup> , pagerank <sup><math>\beta</math></sup> , and weighted_pagerank <sup><math>\beta</math></sup> , where $\beta \in \{\text{None}, 1.0, 1.5, 2.0\}$ . When $\beta = \text{None}$ , we assume raw values (without the $\beta$ exponent), otherwise we adapt the gravitational formula [37] as before: metric/distance <sup><math>\beta</math></sup> , where the distance between the location and the closest tile is measured in meters using KD-tree nearest neighbor search [61].
Demographics <sup>§</sup>	Meta	37	Behavior (5 mobility, 1 business, 4 network, 6 technology), Life events (2), Industry (4 employment), User device (4 assets), User-OS (5 assets), Education (4), Interests (2)
Infrastructure <sup>†</sup>	OpenStreetMap	54	Roads (4), Buildings (2), POIs (24 counts, 24 distances)
Connectivity <sup>†</sup>	OpenCellID	9	distance_to_closest_cell, number_of_cells $\in$ radius, number_of_towers $\in$ radius, where radius $\in \{1.6, 2.0, 5.0, 10.0\}$ Km.
Nightlight intensity <sup>‡</sup>	Google	36	min, max, mean, median, frac_pixels, frac_area, frac_sum_rad, t30_mean, and l30_mean for radius $\in \{1.6, 2.0, 5.0, 10.0\}$ Km.
Daylight satellite images <sup>§</sup>	Google	-	A 640x640 image for each place at zoom level 16 (e.g., buildings and streets)

<sup>†</sup> Open access <sup>‡</sup> Upon request <sup>§</sup> API costs can be covered by free-tiers, sand-box accounts, or Google Cloud research credits for Google products

**B APPENDIX FIGURES**



**Figure B1: Ablation results.** Results on (a,e) recency, (b,f) relocation, (c,g) augmentation, and (d,h) weighting are shown for SL and UG, respectively. Each bar shows the performance of each model as the normalized root-mean-squared error (NRMSE) for the prediction of the mean  $\mu$  (light blue) and standard deviation  $\sigma$  (dark blue) of the IWI scores. This normalization allows comparing the two outcomes, and we see that  $\mu$  is easier to infer than  $\sigma$ . Recency and relocation values are averages across the CB and CNN models (i.e., neither weighted nor augmented samples). Additionally, experiments across different recency configurations use no relocation (relocation=none). The rest of experiments use the combined recency (recency=ON) since in both countries, combining the last two available surveys as ground-truth yields the smallest  $\epsilon_\mu$  across all recency configurations. In terms of relocation, the noisy locations are best for UG (relocation=none), and while relocation is better in SL (relocation $\neq$ none), we use the noisy locations in further experiments since the difference in performance is minimal. Augmentation only benefits UG, and weighted samples benefit only SL. Note that the combined models (CNN\*+CB\*) are excluded from this analysis.



**Figure B2: Mean wealth, variability and population size:** These are the predicted values of wealth on the populated places of (a) SL and (b) UG using their respective best models  $CB_w$  and CB. The x-axis shows the predicted mean IWI scores while the y-axis shows the predicted standard deviations. We see that urban places (blue) are often wealthier than the rural ones (similar as in the ground-truth data, see Figure 3), and wealthier places tend to be more populated than the poor. Here, population refers to the number of total people within the populated place in a diameter of  $d \approx 1.6$  Km from its centroid. Annotations refer to the number of populated places  $n$ , and the Pearson correlation  $r$  between the two predicted variables.