

Estimating and Forecasting Income Poverty and Inequality in Haiti

using Satellite Imagery and Mobile Phone data



**Cataloging-in-Publication data provided by the
Inter-American Development Bank
Felipe Herrera Library**

Estimating and forecasting income poverty and inequality in Haiti: using satellite imagery and mobile phone data / Neeti Pokhriyal, Omar Zambrano, Jennifer Linares, Hugo Hernández.

p. cm. — (IDB Monograph ; 824)

Includes bibliographic references.

1. Poverty-Haiti-Data processing. 2. Income distribution-Haiti-Data processing. 3. Geographic information systems-Haiti. 4. Machine learning-Haiti. 5. Haiti-Social conditions-Data processing. 6. Haiti-Economic conditions-Data processing. I. Pokhriyal, Neeti. II. Zambrano, Omar. III. Linares, Jennifer. IV. Hernández, Hugo. V. Inter-American Development Bank. Country Office in Haiti. VI. Series.
IDB-MG-824

JEL Classification: I3, I32, I38, O3, O31, O35, O54

Keywords: Haiti, Caribbean, big data, satellite imagery, poverty maps, measurement and analysis of poverty, economic development, income inequality, social innovation, machine learning, geographic information systems

Copyright © 2020 Inter-American Development Bank. This work is licensed under a Creative Commons IGO 3.0 Attribution-NonCommercial-NoDerivatives (CC-IGO BY-NC-ND 3.0 IGO) license (<http://creativecommons.org/licenses/by-nc-nd/3.0/igo/legalcode>) and may be reproduced with attribution to the IDB and for any non-commercial purpose. No derivative work is allowed.

Any dispute related to the use of the works of the IDB that cannot be settled amicably shall be submitted to arbitration pursuant to the UNCITRAL rules. The use of the IDB's name for any purpose other than for attribution, and the use of IDB's logo shall be subject to a separate written license agreement between the IDB and the user and is not authorized as part of this CC-IGO license.

Note that link provided above includes additional terms and conditions of the license. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the Inter-American Development Bank, its Board of Directors, or the countries they represent.



Estimating and Forecasting Income Poverty and Inequality in Haiti

using Satellite Imagery and Mobile Phone data

Authors:

Neeti Pokhriyal,
Omar Zambrano,
Jennifer Linares,
and Hugo Hernández

Content

Executive summary	4
1. Introduction	5
2. Literature Review	7
2.1 Estimating social indicators from census and household surveys	7
2.2 Estimating social indicators using auxiliary data	8
3. Snapshot of Haiti's social indicators	10
4. Methodology and Results	12
5. Policies to tackle social deprivations in the COVID-19 era and the recovery period	33
6. Conclusion	35

Acknowledgements

The authors would like to thank UN ECLAC, especially Fabiana Del Popolo and Alejandra Silva from CELADE, for access to Haiti's valuable census microdata through the use of their Redatam software. The authors would also like to thank Lisseth Escalante, Ricardo Benzecry and Salvador Traettino for their excellent research assistance, and Jean Marie Cayemitte and Gihanne Ambroise from the Banque de la Republique d'Haïti, Evans Jadotte from the World Bank, Raulin Cadet from Quisqueya University, Marta Ruiz-Arranz, Juan José Barrios, Agustin Filippo, José Antonio Mejía-Guerra, Nicola Magri, and Boaz Anglade from the Inter-American Development Bank, Jonathan Hersh from Chapman University, and Michael Mann from George Washington University for their valuable feedback. Finally, the authors would like to thank Duare Pinto for the graphic design of the report.



Prologue

Tracking poverty and shared prosperity is an essential task to inform a country's policymaking process. This is especially the case at present times, given that the economic crisis associated with the COVID-19 pandemic threatens to disproportionately affect vulnerable households. Nonetheless, doing so requires the frequent collection of household survey data, which can be a costly and labor-intensive effort. In addition, the granularity of the collected data is not sufficient for the proper identification of pockets of poverty and of income inequality. In Haiti's particular case, there has not been a household survey since 2012, and the most recently available data is only representative at the department level, potentially diluting important social deprivations present at the more local levels.

In view of these challenges, the IDB commissioned a machine learning framework that can estimate the distribution of income poverty and inequality in 140 *communes* and 570 *sections communales* in Haiti using features extracted from anonymized mobile phone data and satellite imagery. These features have proven to correlate with income poverty and other social indicators in several developing countries, including Brazil, Mexico, Belize, Cambodia, and Bangladesh. While these techniques are not intended to be a substitute of the valuable data obtained through household surveys, they represent an innovative and timely complement to the efforts of tracking social indicators.

Several trends were observed: First, that better-performing *communes* in Haiti were usually located in the Ouest department, where the capital of the country is located. Second, that some *communes* located in the Sud departments became increasingly more deprived than the rest over the last five years, likely due to Hurricane Matthew in late 2016. Third, that the Nord-Ouest *communes* increasingly became poorer than the rest over the past years, a trend consistent with the deterioration of the food emergency situation in this department reported by other international organizations. Finally, our machine learning framework identified that one out of four of the Haitian poor lived in 10 *communes* in the Artibonite, Ouest, Nord-Ouest and Nord departments, and that only three out of the 140 *communes* studied present low levels of income inequality.

Our resulting estimates, which were mapped for their convenient interpretation, provide evidence that there is a need of incorporating a territorial perspective to all growth strategies. We hope that this innovative approach to measuring social indicators is used for COVID-19 response efforts and for the design of a more inclusive growth agenda in Haiti. We also hope that this framework is adopted and further refined as new improvements to machine learning techniques are developed in order to complement future estimations of social indicators.

Yvon Mellinger

IDB Country Representative in Haiti



Introduction

Households surveys and censuses are the building blocks of informed policymaking as they contain relevant demographic and socioeconomic information of a country's population. In order to monitor the potential progress of a country in the achievement of development goals, both household surveys and census data must be collected on a regular basis. In fact, the Enhanced General Data Dissemination System (e-GDDS) of the International Monetary Fund encourages countries to collect population data (via population censuses) every ten years, and poverty data (via household surveys) every three to five years. Haiti carried out its last Population and Housing Census (*Recensement General de la Population et de l'Habitat*) in 2003 and its post-earthquake household living conditions survey (*L'Enquête sur les Conditions de Vie des Ménages Après Séisme* or ECVMAS) in 2012. Therefore, at the time of publishing of this report, Haiti had not carried out a census in seventeen years and a household survey in eight years¹.

Haiti is not the only country that has struggled to carry out household surveys and censuses on a regular basis. In fact, Serajuddin et al. (2015) estimate that over a third of the world's developing or middle-income countries had one or less poverty estimates between 2002 and 2011. This makes tracking poverty and shared prosperity a challenge in the developing world. One of the main reasons behind this challenge is the high cost of collecting ground data. According to Kilic et al. (2017), the average cost of a survey in a fragile state is about \$186 per household². If we were to take the sample of households surveyed in ECVMAS (about 5,000), the cost would be about \$930,000 in real 2014 US dollars (excluding any capacity building costs). Further, household surveys may not be representative at the desired geographical level—the 2012 ECVMAS, for instance, is representative at the national and departmental levels (10 geographic units for the latter), potentially masking variations at more granular levels of disaggregation—the *commune* level (144 geographic units) and the *section communale* level (571 geographic units). Moreover, safety concerns may limit the capacity of the government in carrying out these surveys in countries in conflict (Engstrom et al. 2017). These considerations make tracking development goals such as the Sustainable Development Goals (SDGs)—which, according to Kilic et al. (2017), rely more on household survey data than the Millennium Development Goals (MDGs)—a challenge for a country like Haiti.

In view of this, researchers have recently turned to auxiliary sources of data such as mobile phone records, satellite imagery, and remote sensing

¹ Prior to the 2012 ECVMAS, the household living conditions survey had been conducted in 2001.

² The average was estimated using the cost per household for three fragile countries: Afghanistan (\$109.25 per household), Iraq (\$149.23) and Yemen (\$298.42).



data to estimate social indicators such as income poverty (Blumenstock et al. 2015; Jean et al. 2016; Engstrom et al. 2017) Many of these auxiliary data sources have the advantage of being frequently updated at a lower cost and of allowing for a further disaggregation of a country's administrative areas. Further, combining several auxiliary datasets (Pokhriyal and Jacques, 2017 and Steele et al., 2017) like mobile phone and remote sensing data have proven to result, in some cases, in more accurate estimations than using each dataset separately. In spite of the novelty and cost-efficiency of these techniques, it is important to note that auxiliary sources are meant to be a complement to ground-truth data collection efforts— especially for the prolonged intervals between each household survey— and should not be considered a substitution, as they provide a snapshot of the distribution of poverty and other social deprivations within Haiti based on the 2012 national poverty line (as we will explain in the methodology section). Updated national poverty figures would require an update of the national poverty line and the determination of the updated income and consumption levels of the population, which can only be obtained from ground-truth data.

The objective of this study is to build a computational framework that can estimate the distribution of income poverty, income inequality, and standards of living deprivation³ in 2014⁴ and 2019 for 140 *communes* and 570 *sections communales* in Haiti using features extracted from anonymized mobile phone data and satellite imagery (referred to as auxiliary data hereafter). The estimation of these

social indicators was done in two stages: a) First, small-area estimates (SAE) of socioeconomic indicators were calculated using a combination of the 2003 census data and the 2012 ECVMAS, b) auxiliary data was used to estimate the social indicators in 2014 and 2019 using a machine learning model trained on the SAE from the first stage. The 2014 *commune* level estimates were validated using the SAEs from the first stage. Our trained model is also used to forecast the deprivations for 2019 both at *commune* and *sections communales*.

Three trends were observed from the resulting estimates: a) better-performing *communes* were usually located in the Ouest department in both 2014 and 2019; b) Some *communes* located in the Sud departments became increasingly more deprived than the rest between 2014 and 2019, likely due to Hurricane Matthew in late 2016; and c) the Nord-Ouest *communes* increasingly became poorer and more deprived than the rest in 2014 relative to 2019, a trend consistent with the deterioration of the food emergency situation in this department reported by the Food and Agriculture Organization (FAO) and the World Food Programme (WFP).

In the context of the COVID-19 pandemic, it is essential to have a better idea of where the most vulnerable population in the country resides in order to prioritize social interventions given the government's limited fiscal space⁵. Therefore, the present study lists and maps the names of the most vulnerable *communes* and *sections communales* in the 2014 and 2019.

3 Risk of malnutrition and ownership of durable goods, based on the definition of Santos and Villatoro (2018)

4 The 2014 maps are dynamic, as they were built using 2014 satellite imagery and mobile phone metadata from March to May in years 2016-2018.

5 The Haitian Ministry of Economy and Finance estimates that Haiti's non-financial public sector's deficit will be 6.2%.



Literature Review

Given the study's two-stage approach for the estimation of social indicators, the literature review section is divided into two subsections: the first subsection summarizes works related to estimating poverty from census and detailed household surveys to create a representative sample at more disaggregate levels, which act as regression targets for the poverty estimates generated from auxiliary data. The second subsection provides a summary of the literature related to using auxiliary data sources for poverty measurement and the estimation of other social dimensions.

2.1 Estimating social indicators from census and household surveys

Mapping socioeconomic conditions has historically been of interest for governments, development organizations, and academia. However, representative data at granular administrative levels is often unavailable in developing countries

as household surveys often use samples that are too small to produce representative estimates beyond national or regional levels. Small area estimations (SAE) is a frequently used indirect estimation technique for the estimation of social indicators, notably poverty, at more disaggregate levels. According to Das and Haslett (2019), basic SAE methods can estimate small area linear parameters, including means and totals, yet they are not well suited for the estimation of non-linear functions, such as distributions. Nonetheless, linear methods that incorporate unit records from a census or an administrative database can model non-linear functions with considerably increased accuracy.

One of such methods is the three-level Elbers, Lanjow, and Lanjow (2002) or ELL method, which was the first one developed for these purposes. The ELL method consists of the following steps: (i) First, similarities between census data and household surveys are identified in order to construct a common vector of independent variables, (ii) second, a parametric model of household income determinants is estimated based on the available information at the highest level of representation of the household survey, and (iii) Finally, the empirical distribution obtained in (ii) is simulated over the census data to obtain small area estimations of greater geographical representation. The ELL method (and variations of this method) has been successfully used for the estimation of social indicators⁶ in many developing countries, including Belize (Hersh, et al. 2020), Brazil (Elbers, et al. 2008), Mexico (Demombynes, et al. 2006), Thailand

⁶ Income poverty in most cases, yet some authors have also estimated indicators such as undernutrition.

⁷ According to Das and Haslett (2019), EBP is based on an area-specific two-level nested error regression model that includes small area level and household level effects, while ELL is based on cluster-specific two-level model using cluster and HH variability with contextual variables from the census.

(Healy et al., 2003), Cambodia (Fujii, 2004), South Africa (Alderman et al., 2002), Brazil (Elbers et al., 2004), Bangladesh (BBS & UNWFP, 2004; Haslett et al., 2014), the Philippines (Haslett & Jones, 2005) and Nepal (Haslett & Jones, 2006).

Another methodology for SAEs worth mentioning is the one proposed by Molina and Rao (2010), which estimates non-linear small area population parameters using the empirical Bayes or best prediction (EBP). According to Das and Haslett (2019), this methodology results in estimators with minimum mean squared errors (MSEs) that are “best predictors” through Monte Carlo approximation, assuming that the transformed welfare variable follows a nested error regression model⁷. This method has been used mostly in European Union countries.

According to Das and Haslett (2019), both of the aforementioned methods derive from standard linear random effects models with potentially strong distributional assumptions and formal specification of the random portion. Moreover, they may not be robust to outliers in the response variable, yet outliers used in the survey-based model do not necessarily imply outliers or non-robustness in the aggregate SAEs.

An alternative to SAE is the MQ approach proposed by Chambers and Tzavidis (2006), which is based on modelling quantile-like parameters of the conditional distribution of the target variable given the covariates. The MQ approach is distribution-free and avoids the strong assumptions associated with specification of random effects, allowing inter-area differences to be characterized by area-specific coefficients. Nonetheless, this method also has a shortcoming: it does not capture cluster-specific random effects by failing to differentiate between small area-specific random effects and cluster-specific effects. (Das and Haslett, 2019). MQ has been used to estimate poverty at the local administrative level in several countries in Europe, including Poland (Marchetti et al., 2018), Albania (Tzavidis et al., 2008), and Italy (Giusti et al., 2011).

A simulation ran by Das and Haslett (2019) uses the three aforementioned methods to reconstruct 2000

and 2001 poverty data in Bangladesh. The study found that the ELL performed better than the EBP and MQ in terms of relative bias and relative root mean squared errors when the majority of small domains have single clusters in the sample and when there is minor between-area variation in the population. Given its relatively simple methodology and widespread success in the estimation of social indicators in developing countries, a computational variant of the ELL method was our technique of choice for the first part of our exercise, combining data from population census with data from household surveys to generate a valid/robust estimations of per capita income at small levels of territorial disaggregation⁸.

2.2 Estimating social indicators using auxiliary data

Several studies have highlighted the relationship of poverty and information captured by satellites, including but not limited to, nightlights, weather, vegetation, and meteorological data. (Elvidge et al., 2009; Pokhriyal and Jacques, 2017; Jean et al., 2016) Moreover, the growth in the number of satellites and the enhancement of their sensing capabilities has enabled precise and timely reporting of various remote sensing measurements. In fact, many of the satellites have a revisit capability of 24 hours, enabling almost real time scanning of earth’s surface. Newer sensors aboard these satellites also capture geo-spatial measurements like cars, shadows (which act as proxy for building heights), density of buildings, transportation information, and roof types (which are useful in the mapping of slums). These measurements are captured at very fine spatial resolution of 25 cm or 30 cm, with varying spectral bands, which facilitate mapping and other downstream tasks from these images. The basic idea behind using satellite imagery to determine deprivations within a country is to extract high-level patterns or signals which correlate with socioeconomic deprivations of a geographical area. For example, satellite images of a given geographical area are scanned to find signs of urbanization, which are shown to correlate with economic development of a country.

⁸ In this approach, as in a traditional ELL method, at the first stage of the estimation a common support vector of variables merging household survey data and census data was determined. At the second stage, a predictive per-capita income model considering both cluster/locality effects, and idiosyncratic household effects, was predicted using a machine learning (ML) framework. The Machine Learning (ML) algorithm was trained using the smaller/representative data set (household survey, and then a probabilistic inference was projected over the out-of-sample larger dataset (census), in an analogue approach to a parametric SAE. The result was a complete set of predicted income or classification for all the households in the census sample, allowing the estimation of socioeconomic indicators at small territorial level. With no parametric structure to restrict the classification/imputation process, the machine learning approach allows for greater flexibility, focus on out-of-the-sample predictions, and systematic optimization of dimensionality problem to avoid overfitting.

Satellite imagery has been studied in the past to look for texture-based features. However, with advances in deep learning algorithms, researchers increasingly focus in extracting automated features from satellite imagery to learn signals of poverty deprivations. (Jean et al., 2016 and Pandey et al., 2018) Researchers have tested the spatial generalization capability of Jean et al., 2016 in developing countries (Head et al., 2017) such as Haiti and Nepal, and reported that the method is sensitive to hyper-parameter setting, and does not trivially generalize to other spatial locations, especially Haiti. They attribute it to the fact that Haiti has a high urban population density (55% of total, with 74% of these living in slums⁹) and thus poses unique challenges in learning signals of development from satellite imagery.

Additionally, the need to study temporal generalization and evolution of machine learning models that predict poverty and wealth based on satellite imagery is highlighted in Engstrom and Hersh (2017). Simultaneously, the need to have spatially high-resolution poverty metrics is highlighted by Watmough et al. (2019). Tang et al. (2018) employ Normalized Difference Vegetation Index (NDVI) data to predict the rate of change of poverty. However, given that it is based on a vegetation index, it is limited in its capability of predicting broader deprivations of poverty and is at coarser resolution of 250 m x 250 m.

One of the most recent initiatives feature extraction in this area is the Functional Map of the World (FMOW) challenge (Christie et al., 2018), which consists of creating automatic solutions to classify

a specific given location as one of 62 target classes (e.g. an airport, a flooded road, a place of worship, a construction site, and so on) or as none of them (false detections). FMOW images vary in quality and are distributed over more than 100,000 globe locations, which leads to high intraclass variations and considerable interclass confusion. This model was used as a feature extractor, as detailed in the methodology section of this report. This method is also compared to a direct night-light-based poverty estimation approach (Jean et al., 2016).

In addition to the use of satellite imagery as auxiliary data for poverty prediction, recent works like Steele et al., 2017, Pokhriyal and Jacques, 2017 have shown the efficacy of using a combination of anonymized mobile phone metadata (which allows for extraction of features that correlate with socioeconomic measures, such as mobility and activity features), and geographical information data for poverty prediction at high resolution. Given that both data sources are generated at different spatial resolutions, they complement each other well: the granularity of mobile data in urban areas compensates for the coarseness of geographical information data in these areas, and allows for estimations at much disaggregate geographical areas, such as neighborhoods (Steele et al., 2017). Using the combination of both data sources provided improved predictive power and lower errors than using these separately, as shown in Pokhriyal and Jacques, 2017. Given that Haiti has a high proportion of its population living in urban areas, we decided to use this method for our 2014 estimates to obtain more precise results.

⁹ According to the World Bank.



Snapshot of Haiti's social indicators

According to the 2012 ECVMAS, 59 percent of the Haitian population lived under the national poverty line (2.41 USD/day), while 24 percent lived in extreme poverty. About 45 percent of the Haitian population lives in the rural areas, where, according to (Ghayad et al., 2019), almost two-thirds of the population is considered chronically poor. Extreme poverty had declined from 31 percent in 2000 to 24 percent in 2012 at the national level, with urban poverty halving, while poverty levels remained about the same in the rural areas over the same period¹⁰. According to the World Food Programme's (WFP) 2020 Global Report on Food Crises, in Haitian rural areas, most vulnerable households lack agricultural work opportunities due to high labor costs and limited resources of farmers. Thus, vulnerable households recur to alternative sources of income such as migration, petty trade, or selling of charcoal.

There are significant differences in the provision of basic services among the urban and rural households. About 11 percent of the rural population has access to electricity, compared to 63 percent in urban areas. Likewise, only 16 percent of the population has access to improved sanitation facilities, while 48 percent did in the cities. Therefore, while poverty is present in both urban and rural areas in Haiti, it tends to be a rural phenomenon. Nonetheless, it is important to note that informal urban settlements have increased over the last decade.

In terms of income inequality, according to the World Bank, the gini coefficient in Haiti was 0.61 in 2012, a level which had remained practically unchanged since 2001¹¹. The gini coefficient estimated by the Haitian Institute of Statistics (IHSI) in 2012 was even higher than the one estimated by the World Bank, and stood at 0.68¹². According to IHSI, there is heterogeneity within each of the ten Haitian departments: In fact, 89 percent of household income inequality in 2012 was intraregional.

¹⁰ According to Ghayad et al. (2019), rural poverty is common among developing countries, as opportunities for education, employment, and access to technology tend to be concentrated in the urban regions.

¹¹ As reported in World Bank Group. 2014. "Poverty and Inclusion in Haiti: Social Gains at Timid Pace". A similar trend is observed when using The Standardized World Income Inequality Database (Solt, 2019).

¹² Note that these values significantly differ from the gini coefficients reported by the World Bank and the Standardized World Income Inequality Database in cross-country databases. In those databases, gini indices are standardized by welfare definition and adult-equivalence scale in order to facilitate comparability with other countries.

While the gini coefficient is a widely-used measure of income inequality, it is important to highlight that it does not feature subgroup consistency: in other words, if inequality declines in one subgroup (for instance, a region) and remains unchanged in the rest of the subgroups, the gini may not properly reflect this change. Therefore, the United Nations Development Program (UNDP) uses the Atkinson measure of inequality, which, in addition to displaying subgroup consistency, is also sensitive to inequality in the lower end of the distribution by putting more weight on the lower end. When using this measure, the UNDP shows that income inequality has been increasing over the last decade, with the most dramatic increase registered in 2018, when the sociopolitical crisis erupted (see figure 1).

In terms of human development, Haiti registered important gains over the last decade, with the human development index (HDI) increasing from 0.474 in 2010 to 0.503 in 2018. Nonetheless, the HDI is much lower when adjusted for income, education, and health inequalities (the inequality adjusted HDI, or IHDI). The IHDI still registered an improvement

between 2010 and 2018 (see figure 2), yet it has been at a slower pace. According to the 2019 UNDP Human Development Report, this was the case due to an increase in inequality in terms of years of schooling and life expectancy. Haiti is the fifth country in the world (out of 148) in terms of losses in HDI when controlling for these inequalities^{13,14}. Haiti had registered significant gains in both the HDI and the IHDI over the past decade. However, in 2018, the IHDI registered a decrease.

Finally, food insecurity is widespread in Haiti. According to the WFP, one out of every three Haitian were acutely food insecure before COVID-19 became a pandemic. This is equivalent to 3.7 million Haitians¹⁵. This was especially the case in the lower Northwestern region of the country and in the urban *commune* of Cité Soleil. Disruptions in the food supply chains due to the global pandemic, coupled with violent conflict, inflation and the currency depreciation registered over the past year could increase the severity of food insecurity and malnutrition in the coming months.

Figure 1: Income inequality (%) using the Atkinson Inequality Index. Source: UNDP-HDI 2019


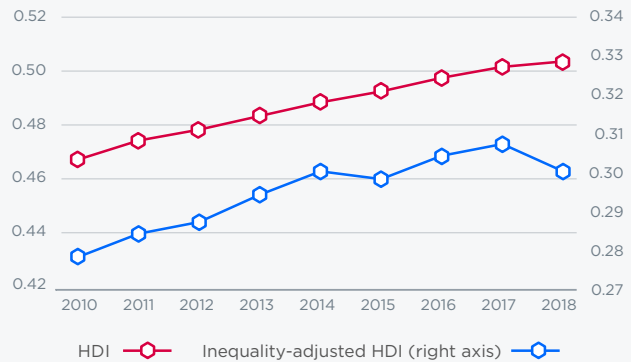



Figure 2: Evolution of the Human Development Index Source: UNDP-HDI 2019



13 The ranking was based on the average losses in HDI due to inequality between 2010 and 2018 for countries with 4 or less missing values. The top four countries with the biggest losses were: Comoros, Central African Republic, Guinea-Bissau, and Namibia.

14 Note that IHDI is based on the Atkinson index. IHDI is not association-sensitive, and therefore does not capture overlapping inequalities. In order to capture association sensitivities, all the data for each individual must be available from a single survey source, and this is currently not the case for Haiti.

15 Haiti is the tenth country with the highest number of people living in food insecurity crisis or worse after Yemen, DR of Congo, Afghanistan, Venezuela, Ethiopia, South Sudan, Syria, Sudan, and Northern Nigeria.



Methodology and Results

The objective of this study is to estimate income poverty, income inequality (Gini), Foster Greer Thorbecke index (FGT)¹⁶, and standards of living deprivation using two different sources of auxiliary data: anonymized mobile phone data and satellite imagery¹⁷. The estimation of these social indicators was done in two stages: a) First, small-area estimates (SAE) were calculated using a combination of the 2003 census data and the 2012 ECVMAS, b) auxiliary data was used to estimate the social indicators in 2014 and 2019 using a machine learning framework¹⁸. The 2014 estimates were obtained by training and validating using the SAEs from the first stage.

Stage 1: Estimating Computational Small-Area Estimates (SAEs) from Census and Household Survey Data

As discussed in the literature review, our method of choice for estimating SAEs is a computational variant of the ELL method, given its relatively simple methodology and widespread success in the estimation of social indicators in developing countries. In the vein of the ELL method, our computational approach consists of the following steps: (i) first, similarities between census data and household surveys are identified in order to construct a common vector of independent variables, (ii) second, a parametric model of household income determinants is estimated based on the available information at the highest level of representation of the household survey (in Haiti's case, at the department level), and (iii) finally, the empirical distribution obtained in (ii) is projected over the census data using Machine Learning (ML) trained predictive modeling to obtain small area estimations of greater geographical representation (in Haiti's case, to obtain estimations at the *commune* level)¹⁹.

¹⁶ The Foster-Greer-Thorbecke (FGT) indices are a family of poverty metrics. FGT_0 measures how widespread poverty is (headcount poverty), FGT_1 measures how poor the poor are (the expenditure discrepancy of poor people towards the poverty line), and FGT_2 gives an indication of how severe poverty is (as it puts higher weight on the poverty of the poorest individuals, making it a combined measure of poverty and income inequality).

¹⁷ Satellite imagery denotes both aerial imagery for year 2014 as well as the publicly available satellite imagery for 2019.

¹⁸ More precisely, both aerial and mobile data were used for 2014, while only satellite data was used for 2019, due to the unavailability of mobile phone records.

¹⁹ The basic problem structure of a typical SAE, which is a smaller data set (household survey) that contains information that has to be imputed over a larger data set (census), replicates the typical problem structure of predictive modelling from the Machine Learning framework. Predicting modeling algorithms exploit the value of information of smaller/representative data sets, where the algorithm learns the optimal parameters to solve estimation/classification problems (training sets), to then project a probabilistic inference over an out-of-sample larger dataset (scoring set). Given this shared structure, it is plausible to use a machine learning approach to solve the problem of SAE of socioeconomic indicators to produce poverty maps.

Two data sources were used for the construction of Haiti's SAE estimates: the 2012 household survey (ECVMAS) and the 2003 population census (the most recent one done in Haiti). The 2012 household data was obtained from the IHSI website, while the census data was obtained from the United Nations Economic Commission for Latin America and the Caribbean (ECLAC) Population and Development Division (CELADE), through their Redatam software.

Just like other household surveys, ECVMAS collects data on primary and secondary sources of income and transfers for the households, but it is only representative at the department level. On the other hand, the 2003 Population Census contains detailed information at the individual level but does not report income data. Nonetheless, both data sources have 129 household and individual variables in common (see table A.1 in the appendix), making it possible to construct a common support vector of poverty-related independent variables. This is the first step of the ELL framework.

For the second step, a per-capita income model was estimated for each household in every department, using the following structure:

$$(1) \quad Y_{h,c} = \beta X_{c,h} + \mu_{c,h}$$

$$(2) \quad \mu_{c,h} = \gamma_c + \varepsilon_{c,h}$$

Where in (1) and (2), $Y_{h,c}$ used household per capita income as a dependent variable with a constant (k=2,510 Haitian gourdes) linear adjustment and $\mu_{c,h}$ is the uncorrelated random error term composed by

the effect of γ_c , which is the cluster/locational effect, and $\varepsilon_{c,h}$ which is the idiosyncratic household effect.

Our predictive modeling process for generation of poverty indicators in this study included continuous, binary and multinomial specifications. A variety of Machine Learning algorithms were evaluated including decision trees, including Random Forest (RF) and Gradient Boosting Machines (GBM), linear models (GLS with Lasso), deep learning models (Neural Networks), and a combination of all of them (Stacked Learning). ML models were selected over the performance criteria of the cross-validation phase of the training set²⁰. Our preferred ML model, the linear GBM, minimized the sum of the squared errors between the estimates and actual headcount poverty rates at the statistically representative territorial portions of the ECVMAS 2012, thus accurately reproducing national and departmental headcount poverty rates (FGT_0) for the base year.

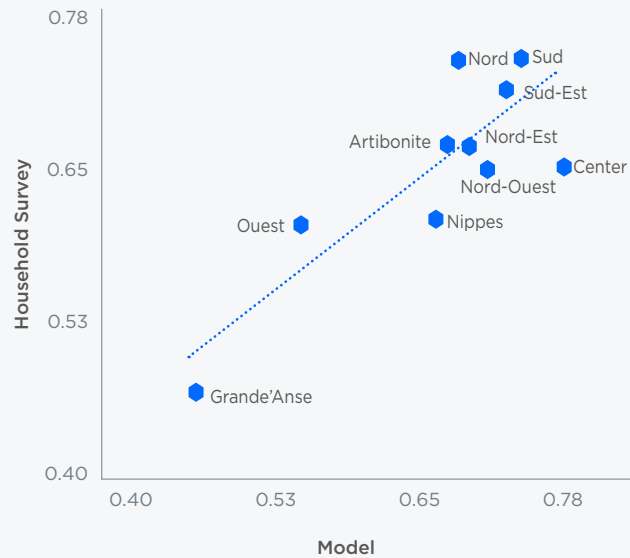
Table 1 reports the fittest models for each model specification, its estimated headcount poverty rate at the national level and a measure of each of their sum of squared errors with respect to the ground truth ECVMAS data. As previously mentioned, the selected model closely replicates the national poverty headcount rate and has the lowest sum of squared errors at the department level. In our case, this model is the GBM linear specification adjusted by a constant. As shown in figure 3, this GBM model specification closely replicates each department's headcount poverty rate, with a R^2 of 0.685.

Table 1: Estimated models for commune poverty rate in Haiti using ECVMAS data

Model	Resulting national poverty rate	Sum of Squared Errors (Department level)
Survey (ground truth)	0.597	0
Linear GBM specification adjusted by a constant	0.571	8.857
Linear specification adjusted by a constant for each income percentile	0.616	9.927
Dichotomous specification adjusted to the national poverty rate	0.631	20.475
Linear (natural logarithm) specification	0.759	48.716

²⁰ A Confusion Matrix for each evaluated model allowed to compare the models' out-of sample predictive capabilities in terms of accuracy, recall and precision metrics. Further, ROC-AUC Curves and Predictive Deciles Plots supported the metrics for the selection of learning models. Our preferred machine learning routine was *Gradient Boosting Machine*, a technique that combines predictions from multiple decision trees to generate a model that minimizes out-of-sample prediction error.

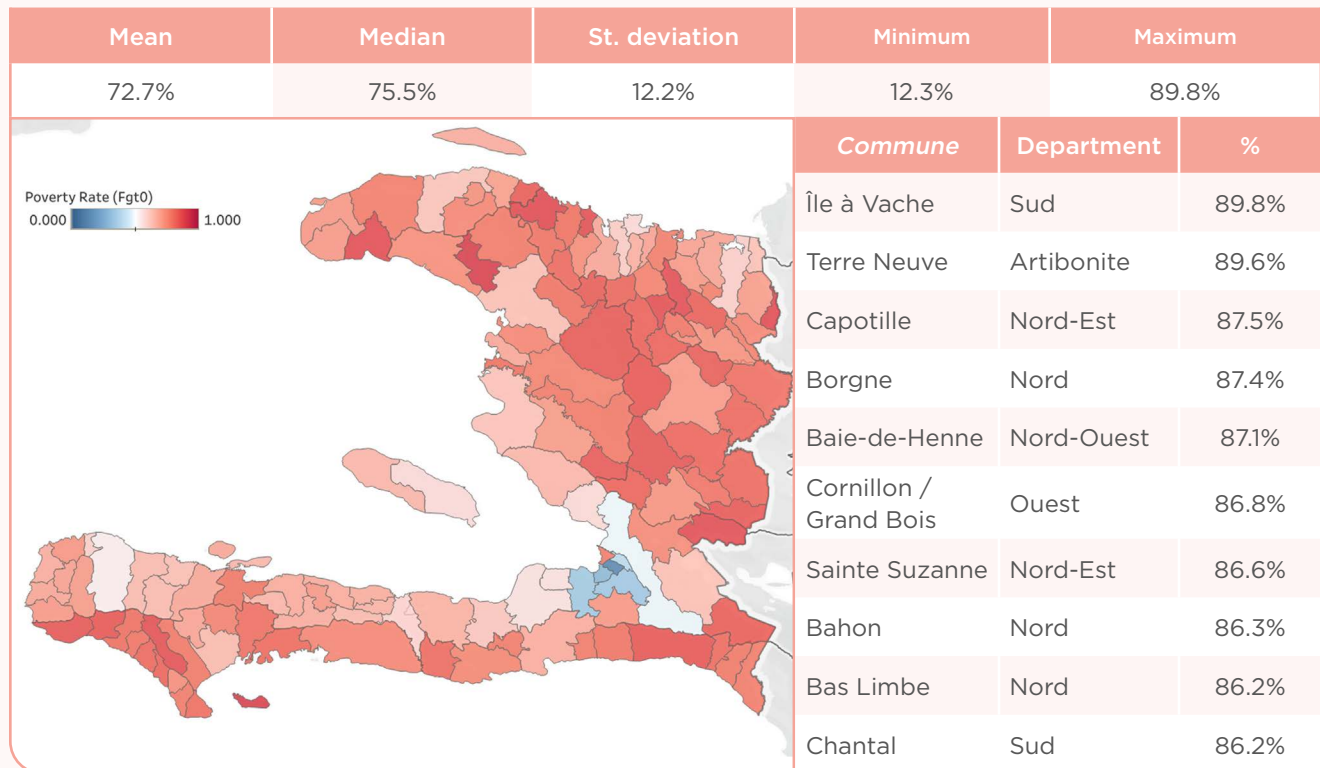
Figure 3: Poverty headcount ratio (using the 2012 national poverty line) R2 = 0.685



Results Income Poverty

Our model yields a national poverty rate of 57.1%, which is equivalent to approximately 5.8 million individuals. Figure 4 shows the resulting poverty rates at the *commune* level, which range from 12.3% (Delmas in the Ouest Department) to 89.8% (Île à Vache in the Sud Department). Note that there is a concentration of *communes* with very high levels of poverty in the eastern side of Artibonite, in the southernmost *communes* in the Sud and Sud-Est departments, and in the Centre department. However, notice also that the mean and median poverty levels among all *communes* are very high, at 72.7% and 75.5%, respectively. Among the 10 *communes* with the highest poverty rate in 2012, six were located in the northern region (which included the Nord-Est, Nord, and Nord-Ouest).

Figure 4: 10 Poorest *communes* in 2012 (income poverty in USD)



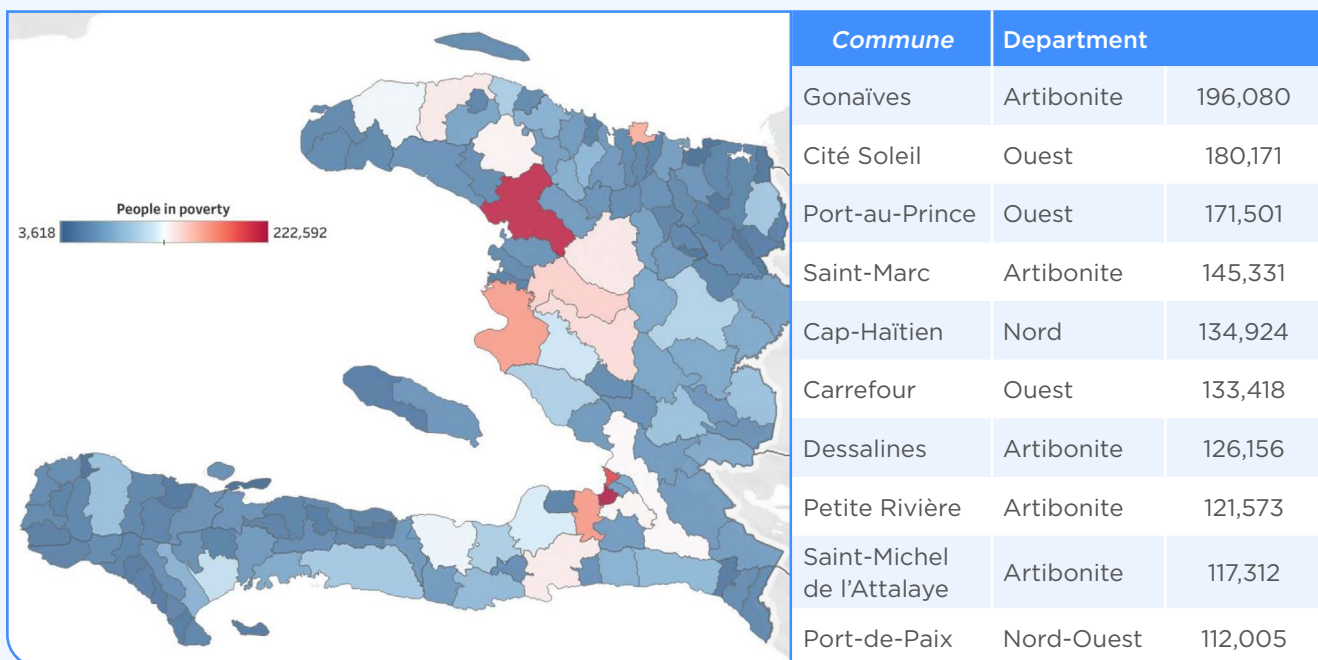
An alternative way to report the results is the number of people living in poverty. Among the 10 *communes* with the highest number of people living in poverty in 2012 (Figure 5), most of them were concentrated in two departments, Artibonite (in Gonaïves, Saint-Marc, Dessalines, Petite-Rivière and Saint-Michel de l'Attalaye) and Ouest (Cité Soleil, Port-au-Prince, and Carrefour). In addition, one out of four of the Haitian poor lived in the 10 *communes* listed in figure 5.

In terms of income inequality, the *communes* with the highest levels of inequality were located in the Sud Department and Artibonite. Nonetheless, notice that the average commune in Haiti has a gini index of about 0.70²¹. These high levels of income inequality are reflected in figure 6, which shows that *communes* in most departments have similar levels of income inequality. In fact, only three out of the 140 *communes* examined in this exercise had a gini coefficient of less than 0.50. These were Cité Soleil, Delmas, and Tabarre in Ouest. It is important

to note that the fact that a *commune* registered a lower gini coefficient (like in the case of Cité Soleil, which has a poverty level of 77.8%) is not indicative of the level of wealth of its households, rather, a comparison of income between households within this particular *commune*. Thus, a low gini coefficient in Cité Soleil indicates that most households within this *commune* have similar levels of (low) income.

The FGT₂ index, which shows how severe poverty is²², shows that the *communes* with the most intense poverty levels in 2012 were located in the Sud and Artibonite departments. The *communes* with less poverty intensity were those located in the Ouest department (Port-au-Prince, Pétion-Ville, Carrefour, Delmas) and Cap-Haïtien (Nord). Notice in figure 7 that there is a concentration of *communes* with severe poverty in Centre, the eastern region of the Artibonite department, and the southernmost *communes* of the Sud department.

Figure 5: 10 Communes with the largest number of people living in income poverty in 2012



21 The gini index or coefficient is a measure of income inequality, where 0 represents perfect equality and 1 represents perfect inequality. *Communes* with gini coefficients approaching 1 feature the highest levels of income inequality.

22 This measure gives greater weight to those that fall far below the poverty line than those that are closer to it.

Figure 6: 10 communes with the highest levels of income inequality in 2012 (Gini Index)

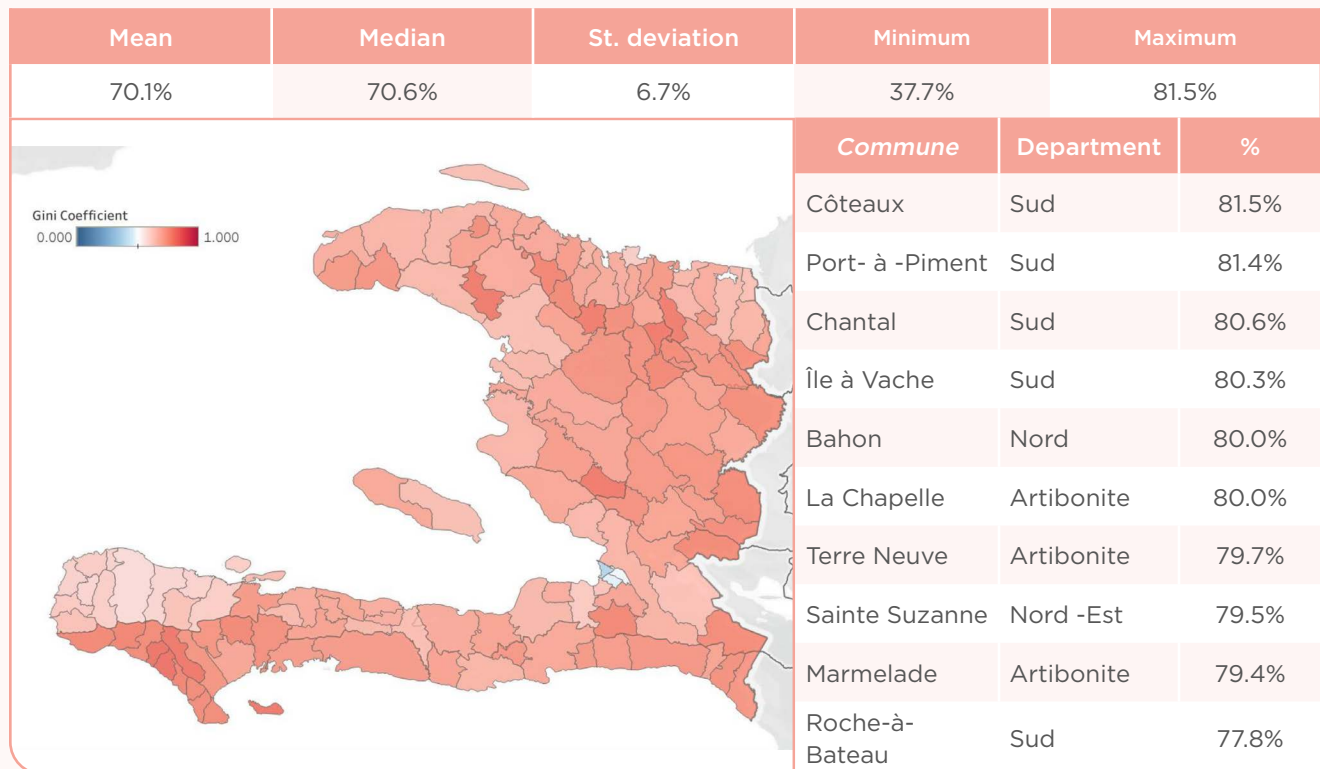
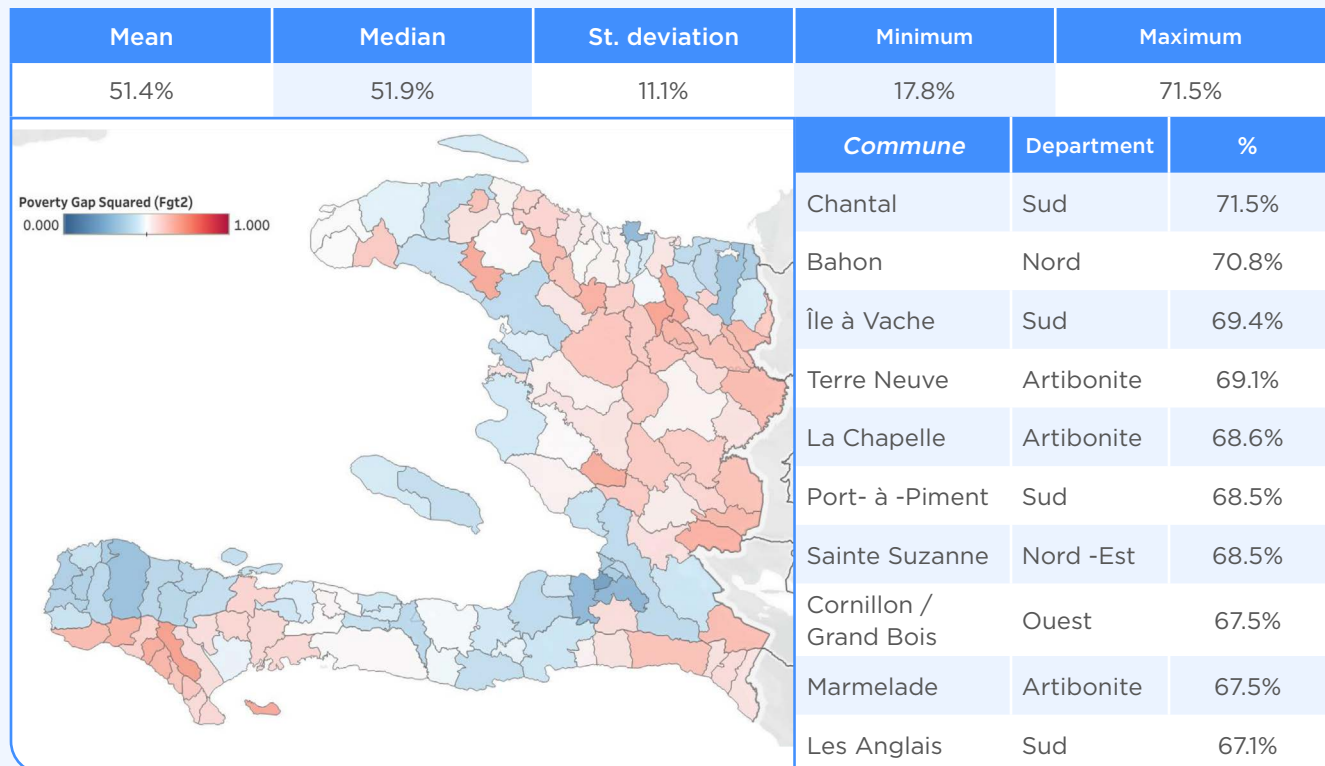


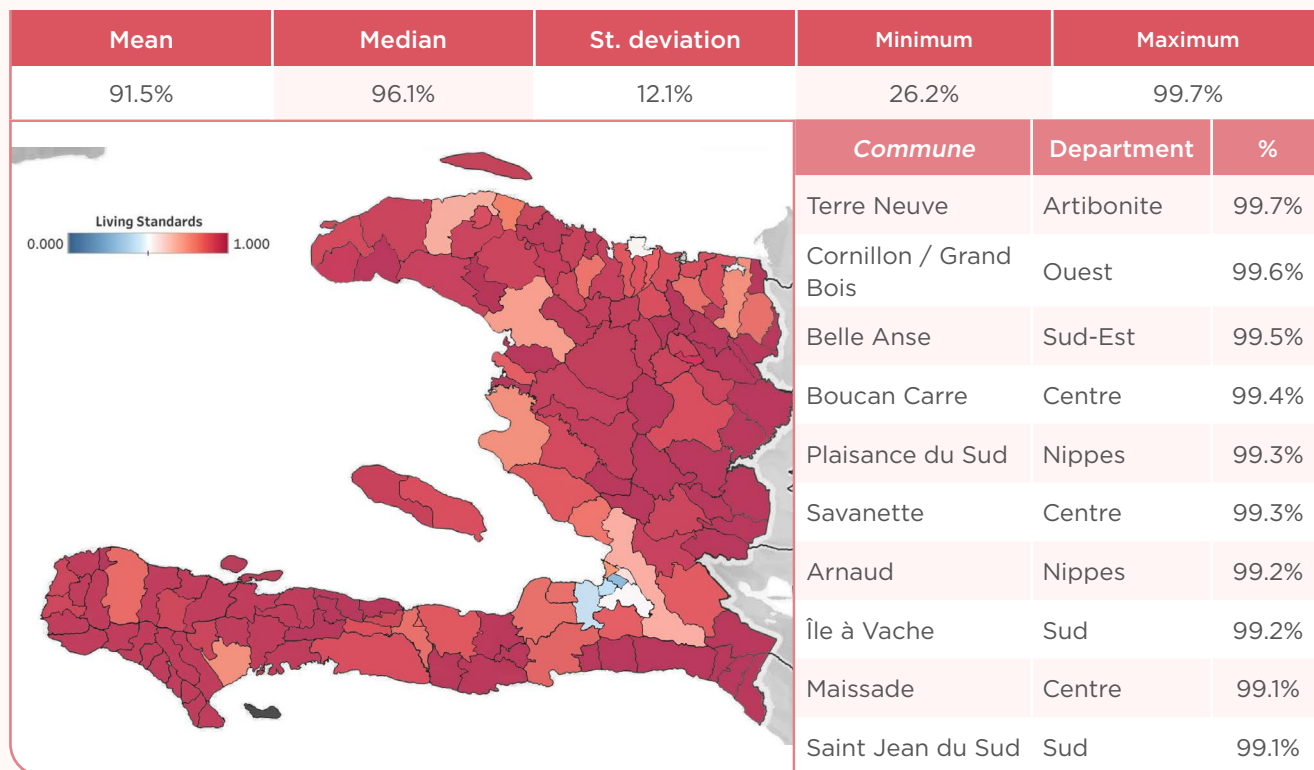
Figure 7: 10 Communes with highest FGT₂ poverty (in USD)



In terms of standards of living deprivation²³, most *communes* registered high levels of deprivation (over 90%, on average). The *communes* with the least deprivations were Delmas, Port-au-Prince, Carrefour, Pétion-Ville, Cap-Haïtien, and Tabarre.

Most of these are located in the Ouest department (except for Cap-Haïtien, which is located in the Nord department). These *communes* are highly urban.

Figure 8: 10 Communes most deprived of standards of living



²³ The standards of living deprivation is a subcomponent of the multidimensional poverty index according to the Santos and Villatoro (2018) definition. The results discussed in this report, however, only account for the incidence (H) of the standards of living deprivation, and not the intensity (A).

Stage 2: Estimating Social Deprivations using Mobile Phone Data and Satellite Imagery

Mobile phone data description, feature extractions, and data aggregating procedures

The results obtained from the previous stage are used to validate more contemporary estimates of poverty and other social deprivations using mobile phone metadata and satellite imagery. This section provides the data description for each source, along with feature extraction and building targets for deprivations. Please note, that since our study was performed both at the spatial granularity of *commune* and *sections communales*. We will refer to these collectively as “microregions” hereafter, unless otherwise stated.

The anonymized mobile phone data was obtained from FlowMinder²⁴ and it includes data from March 1 to May 31 for 2016, 2017, and 2018. In addition to allowing for population estimates, the dataset allowed for the estimation of road usage by producing a set of trips from mobile data, and then mapping these trips. From this data, we get the estimated road usage per micro region, by finding the road segments that lie within microregional boundaries and aggregating the road usage for those segments.

Moreover, the mobile phone dataset allowed for the estimation of the number of subscribers in a given location across different time points, thus producing an estimate of population-wide movement for the years 2016, 2017 and 2018 available at microregional levels (for more information, please refer to the appendix box A.1).

We conceptualize the migration data as a directed weighted graph G , with each microregion as a node, V , and an edge exists between two nodes V_a and V_b if there is a population flow from V_a to V_b . The amount of population flow determines the weight of the edge. The graph G encodes the migration structure across the micro regions. To ascertain a metric with each node, we run centrality measures, namely the Pagerank algorithm that calculates the influence or importance of a node, recursively, using the importance of the nodes connected to it.

Pagerank is an important centrality measure and has been used in context of quantifying information accessibility from mobile phone data to study socioeconomic deprivations. (Pokhriyal and Dong, 2015; Soto et al., 2011).

Since we have access to very limited mobile features, aggregating them to *communes* and *sections communales* depends on the specific feature. We calculated three features, one each from population estimates, migration and road usage data, respectively. The first resulting feature is the log of population density for each microregion. The second feature is the network-centric pagerank feature, as discussed earlier. The third feature is the road usage statistics per microregion and is calculated from road usage data.

Imagery data description, feature extractions, and data aggregating procedures

In the case of the imagery, we used 3,246 aerial images from 2014. Each of the images was 12,000 x 12,000 pixels and each pixel represents 0.25 x 0.25 sq. m. on the ground. Thus, each image sums up to 3 x 3 sq. km on the ground. Given the high resolution of satellite images, several pre-processing steps are needed that involve coarsening to a required resolution and dividing the image into units that can be taken as input by the feature extractor model. For feature extraction, we use two state-of-the-art methods described as follows:

- 1. Functional features-based extractor:** It is based on the Functional Map of the World (FMOW), which is a dataset consisting of 1 million satellite images spanning over 200 countries to train machine learning methods that can predict the functional purpose of the buildings/land in the images. The idea is to classify each image as belonging to one of the 63 functional classes (like airport, construction site, shopping mall, or residential unit). We use a state-of-the-art deep neural network model for image analysis called DenseNet (Huang et al., 2016), which is trained on the FMOW dataset images (for details of the model, see Christie et al., 2018) To extract features from the satellite imagery, we divide the imagery into smaller units, each covering an area of 112 X 112 sq. m. The model takes in

²⁴ Inter-American Development Bank mimeo: Flowminder final consulting report: Caracol Industrial Park and National Road Network. An analysis of commuting and migration patterns.

satellite imagery and runs it through its network. We intercept the model at the 6th penultimate layer and use the output at that stage as our feature vector for that image.

2. Nightlight-based extractor: As a second feature extractor model, we use another state-of-the-art model, based in Convolutional Neural Network (CNN) that is trained to learn the relationships between daytime satellite imagery and its corresponding nighttime light imagery in the context of predicting poverty and socio-economic deprivations. It is trained for over 3,000 DHS (Demographic and Health Surveys) clusters in Africa (For details, see Jean et al., 2016). This model takes as inputs satellite imagery at resolution of roughly a 1-sq. km area for the satellite imagery of years 2014 and 2019 and extracts features corresponding to it.

During our experiments, we found that the night-light (NL) based feature extractor performed better than the FMOW. It is what we had expected, as the NL feature extractor is tuned for predicting poverty, while the other extractor is mostly tuned on high resolution images of the developed world, and is built to look for functional forms or uses of land (Christie et al., 2018). Hence, for the subsequent analysis, we used the features from the NL feature extractor.

Since both feature extractor models produce a long feature vector (order of 2,000 -4,000-dimensional length), we use principal component analysis (PCA)

to automatically reduce the dimensions that retain 90% of the variance in the data. The benefit of dimensionality reduction is that it considerably helps reduce computational cost. Also, since we have limited spatial locations (*communes* and sectional *communes*), reducing the dimensions helps prevent overfitting the model to the data.

It is important to note that the satellite imagery is at a higher resolution than the policy planning units in Haiti for which poverty and other social deprivations is ideally reported. Therefore, as a next step, the features from the satellite images were aggregated to the desired microregions using the following approach.

Satellite images are in the form of rectangular grids, whereas the microregions are available as shapefiles containing their irregular polygon boundaries. Numerous rectangular grids fall within a polygon (see Figure 9). All grids that fully lie within the boundaries of a microregion are assigned to that microregion. Grids that lie within the intersection of multiple microregions are assigned to the microregion where it has the most intersections. To get the image features from a microregion, we estimate an average of all the image features corresponding to satellite images that are assigned to that microregion. This gives an aggregated feature vector summarizing the satellite imagery for a microregion.

BOX 1: Important Data Considerations and Challenges

There are several challenges that result from working with auxiliary data sources, namely satellite imagery and mobile phone data. These datasets, usually, exist at different spatial and temporal granularity. Mobile phone data are available for each subscriber, while environmental data have mixed spatial resolution, from very accurate vector data to low-resolution satellite imagery. The specific challenge in dealing with data heterogeneity is to identify the optimal spatial resolution for the satellite imagery that can capture the signal for development and can be exploited to predict poverty. Since the satellite imagery exists in high-resolution grids, one must decide the resolution which will act as the unit of analysis. We believe it is dependent on the task and the country in question, as all spatial locations within a country will not be homogeneous in containing the signal for development. An example is the highly dense urban areas versus the mostly agriculture-based rural areas.

Another consideration is the computational expertise and resources needed to analyze the auxiliary data. If state of the art deep neural network based methods are to be explored, then it is important to keep in mind that training a deep model is a computationally expensive task, both in terms of time, money, and computational expertise and resources availability of computational resources and expertise.

While satellite imagery at a certain resolution is publicly available, mobile phone data is always proprietary and is thus not openly available to researchers. It also has privacy concerns, though all analyses in this work have used anonymized and aggregated metadata.

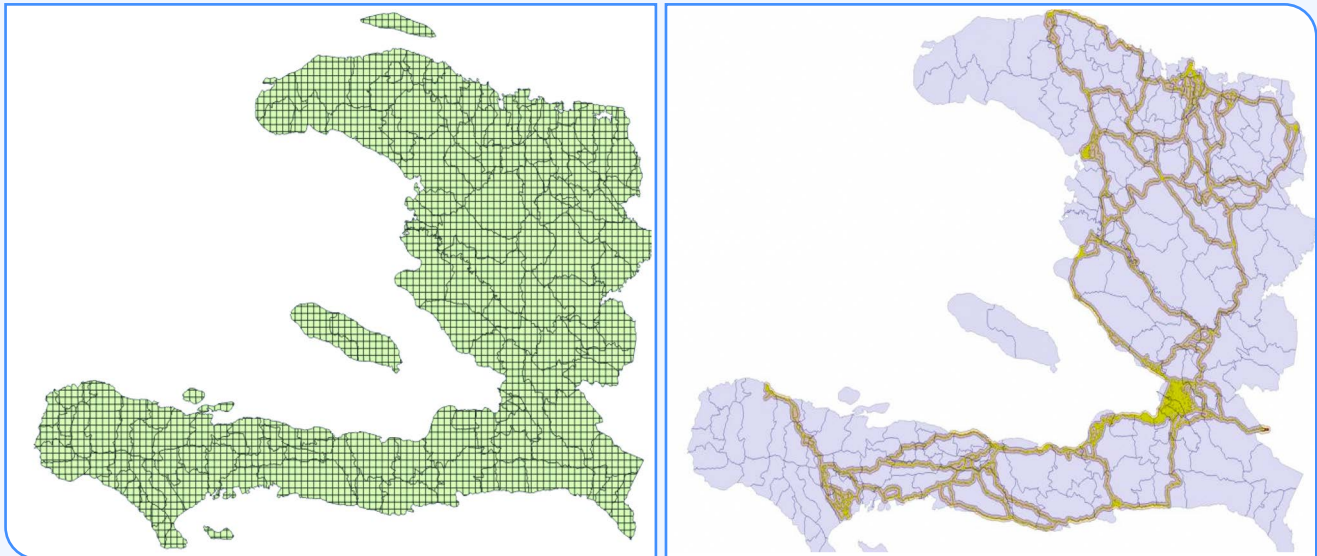
Another major challenge is how to validate the model estimates at finer spatial granularity of *sections communales*. This is owing to the unavailability of ground truth data that can be extracted from census/surveys at that spatial granularity and/or period.

A note about bias in the datasets: It is important to note that there are biases in the auxiliary datasets, which might impact the inferences that are drawn. Mobile phone data suffers from selection bias, as it represents only subscribers of a particular telecom provider. Thus, important demographics like children and the ultra-poor (who may not have access to a mobile phone) may not be properly represented in the data. Also, results might be biased towards urban areas, rather than rural ones, due to satellite imagery mostly picking up urbanization signals as well as the access to electricity for mobile phone service. Given that Haiti has a higher urban density than other low-income countries, this should not be as big of a challenge as in other countries.

We did see improved performance accuracy when different auxiliary data sources were combined, especially for targets such as average income and FGT indicators. Thus, combining satellite imagery and mobile phone data might assist in mitigating some of the biases existing in these datasets, but further longitudinal studies need to be undertaken to get a conclusive evidence.

Finally, our model estimates and uncertainties are also affected by the poor data quality owing to temporal and spatial resolution in satellite imagery. However, better quality of input data should benefit model estimates.

Figure 9: Left - Satellite Images for Haiti for 2014.
Right - Road Map of Haiti for 2016 calculated from the mobile phone metadata.



Methodology

Our research question is expressed as a regression, with the independent variables/covariates being the features extracted from the data sources and the dependent variable being the target that we want to estimate: Poverty (in USD), FGT indices, gini, and standards of living deprivations. The targets are calculated using the 2012 ECVMAS and the 2003 census.

Our idea is to build individual regression models on each source of data, and then combine those predictions in a Bayesian weighted manner. A naïve way to combine different data sources is to concatenate the different feature spaces together, but such concatenation with few data points leads to overfitting (Xu et al., 2013). We use Gaussian Process regression, which provides uncertainty with its predictions, which are then used to get the combined predictions. The advantage of this approach lies in its modularity, as more data sources can be easily added (since we run regression on individual data and their predictions combined). Additionally, each data source remains private to its ecosystem, and only the output predictions need to be shared. This is an important concern when working with datasets that lies within different private and public entities and has privacy concerns (like mobile phone metadata).

Each of our individual regression model is of the form:

$$(3) \quad y_i = \beta^T x_i + f(x_i) + \epsilon$$

where y_i is the target value and x_i is a vector of independent variables derived from the particular data source for the i^{th} microregion. The first term is a linear combination of the independent variables. The function $f(x_i)$ models the non-linear relationship between y_i and x_i . The residual term, ϵ , models the remaining unexplained noise, and is modeled as a zero-mean Gaussian random variable, i.e., $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$. Without the non-linear term, $f(x_i)$ in the equation, the model is equivalent to an ordinary linear regression. However, instead of assuming a fixed parametric form for $f()$, we adopt a non-parametric approach, by assuming a Gaussian Process (GP) prior on $f()$. For details on the generative process and the methodology, please refer to the appendix.

Combining source-specific models

To predict poverty metrics for a microregion, we used the model specified in the previous equation. If only one data source (mobile data or imagery) was available for this region, then only the model specific

to that data source can be employed. For D data sources, D independent models are employed. Now, for each data source, denoted as d (where $d \in D$), our model produces a posterior Gaussian distribution, denoted by $y_{id} \sim N(\bar{y}_{id}, \sigma_{id}^2)$. The combined poverty estimate, y_i , is assumed to be a mixture distribution consisting of d Gaussians, defined above. The mixing weights for a given data source, d (where $d \in D$) and microregion i , are defined as:

$$(4) W_{id} = \frac{\frac{1}{\sigma_{id}^2}}{\sum_{v=1}^D \frac{1}{\sigma_{iv}^2}}$$

The weights assign greater importance to the source that provides a smaller predictive variance, signifying higher confidence in the prediction for the particular microregion. The mean and the variance for the combined poverty estimate for a microregion are given as:

$$(5) E[y_i] = \sum_{v=1}^D w_{iv} \bar{y}_{iv} \text{var}[y_i] = \sum_{v=1}^D w_{iv} \sigma_{iv}^2 + \sum_{v=1}^D \sum_{v'=1}^D w_{iv} w_{iv'} (\bar{y}_{iv} - \bar{y}_{iv'})^2$$

Out-of-sample generalization using spatial validation

To measure the extrapolation capacity of the model on out-of-sample data, we used spatial cross-validation (CV) techniques, which are more robust (Deville et al., 2014; Bahn and McGill, 2013). In spatial CV, the training and evaluation/test sets are taken from geographically distinct regions. Such validation is more robust than standard cross-validation techniques (Blondel et al., 2012; Bahn and McGill, 2013), and thus proves that our model is generalizing well to out-of-sample instances. For Haiti, during each CV run, our model is trained on

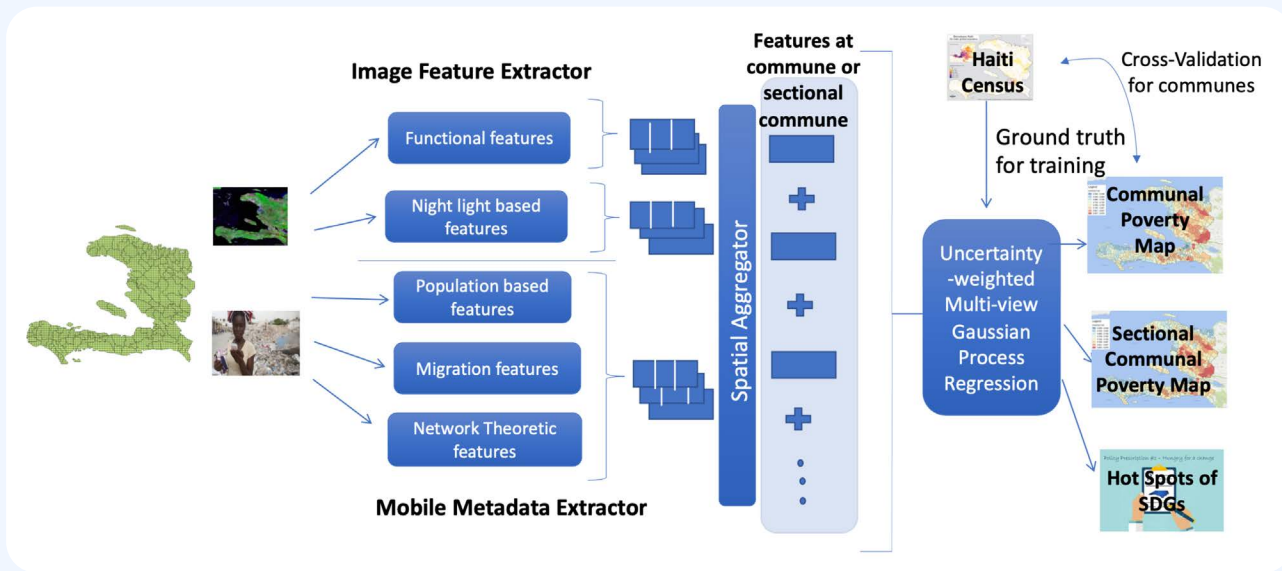
9 of the Haitian departments and tested on the remaining department. Our CV procedure is run 10 times, to ensure that all *communes* are tested.

It is important to note that, since we do not have access to SAEs at the *section communale* level, we cannot validate these predictions with the ground data. This lack of validation of the extrapolatory power of the poverty estimation models is a challenge as discussed by Rodriguez Castelan et al., 2019. However, these estimates can be validated when newer surveys are done.

We use two metrics of evaluation- Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). MAE is the average of the errors for a set of predictions. RMSE is square root of the mean squared error. Both MAE and RMSE are in the same units as the target variable, and lower values of them indicate a better fit. Both are good measures of how accurately the model predicts the target. However, since in RMSE the errors are squared, its values can be higher if some errors are higher. The methodology detailed above is summarized in figure 10 below.

Our model produces estimates of poverty and its associated deprivations for 2014 at *commune* and *sections communales*, as well as forecasts for 2019. The results and maps that are named as 2014, are built using 2014 satellite imagery and mobile phone metadata spanning few months for 2016/17/18, and are thus dynamic in nature. We understand that it would be better for the model to have concurrent mobile phone and satellite imagery, but owing to problems with data acquisition, we decided to make the best use of data at hand. Also, it should be noted that poverty and social deprivations tend to evolve slowly over time, and thus one needs to be mindful that the input data should not capture short-lived variations. The forecasted maps for 2019 are produced using only satellite imagery for 2019, due to the unavailability of recent mobile phone data.

Figure 10: Summary of the methodology for the estimation of social indicators using auxiliary data.



Tables 2 and 3 summarize the cross-validated results at *commune* level for each of the targets in 2014, related to poverty and its deprivations. To assess the efficacy of our model, we experimented with the following settings:

1. Multiview Gaussian Process (GP) Model: This is our GP model, with two auxiliary sources of data- mobile and satellite imagery.

2. Linear Model: Linear Regression model.

3. Mobile: It is our GP model, and the independent variables are the features extracted from mobile phone metadata.

4. Satellite Imagery: It is our GP model, and the independent variables are features extracted using the NL model from the satellite images.

Table 2: Mean Absolute Error (MAE) values from spatial cross-validation at *commune* level.

Model	Statistics	Average Income	Poverty	Extreme Poverty	Poverty (in USD)	Standards of Living Deprivation
	Min	1375.20	0.11	0.06	0.10	0.26
	Max	10012.20	0.88	0.79	0.87	1.00
Multiview GP	MAE	620.73 (215.72)	0.07 (0.02)	0.10 (0.03)	0.08 (0.03)	0.04 (0.02)
Linear	MAE	622.57 (160.59)	0.07 (0.02)	0.09 (0.03)	0.08 (0.02)	0.04 (0.01)
Image Only	MAE	778.47 (430.08)	0.09 (0.04)	0.11 (0.04)	0.09 (0.04)	0.08 (0.07)
Mobile Only	MAE	641.57 (206.80)	0.08 (0.03)	0.10 (0.03)	0.08 (0.03)	0.03 (0.01)
Model	Statistics	FGT ₁ HTG	FGT ₂ HTG	Gini	FGT ₁ USD	FGT ₂ USD
	Min	14.73	217.11	0.32	0.35	0.12
	Max	87.83	7713.61	0.82	2.09	4.38
Multiview GP	MAE	6.67 (2.42)	916.10 (323.98)	0.06 (0.03)	0.16 (0.06)	0.52 (0.18)
Linear	MAE	7.85 (3.14)	1043.55 (369.26)	0.07 (0.03)	0.19 (0.08)	0.59 (0.21)
Image Only	MAE	6.92 (2.47)	968.12 (318.92)	0.06 (0.02)	0.16 (0.06)	0.55 (0.18)
Mobile Only	MAE	7.30 (2.95)	1006.47 (399.38)	0.06 (0.03)	0.17 (0.07)	0.57 (0.22)

Note: Mean values and the corresponding standard deviation (in parenthesis) over 10 cross-validation runs are reported.

Table 3: Root Mean Squared Error (RMSE) values from spatial cross-validation at commune level.

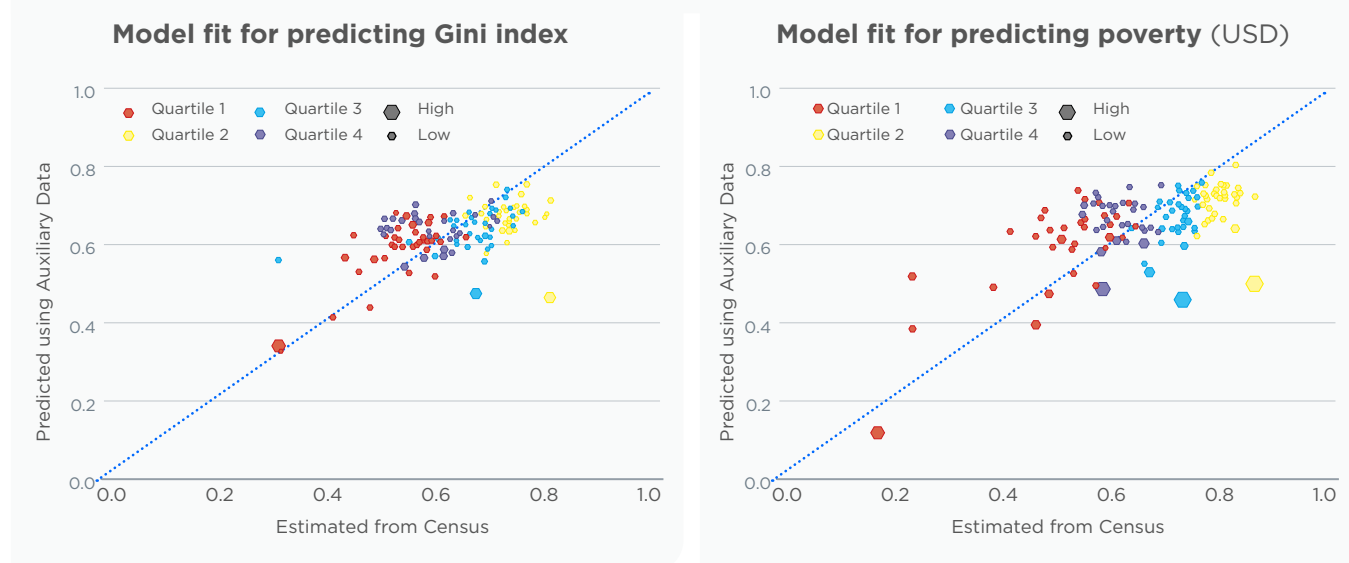
Model	Statistics	Average Income	Poverty	Extreme Poverty	Poverty (in USD)	Standards of Living Deprivation
	Min	1375.20	0.11	0.06	0.10	0.26
	Max	10012.20	0.88	0.79	0.87	1.00
Multiview GP	RMSE	785.97 (269.47)	0.09 (0.03)	0.13 (0.05)	0.10 (0.03)	0.07 (0.07)
Linear	RMSE	815.12 (199.84)	0.10 (0.03)	0.13 (0.05)	0.10 (0.03)	0.04 (0.02)
Image Only	RMSE	1159.38 (859.27)	0.13 (0.07)	0.16 (0.09)	0.13 (0.08)	0.13 (0.15)
Mobile Only	RMSE	742.21 (205.17)	0.09 (0.02)	0.12 (0.03)	0.09 (0.02)	0.04 (0.02)
Model	Statistics	FGT ₁ HTG	FGT ₂ HTG	Gini	FGT ₁ USD	FGT ₂ USD
	Min	14.73	217.11	0.32	0.35	0.12
	Max	87.83	7713.61	0.82	2.09	4.38
Multiview GP	RMSE	7.98 (2.94)	1082.58 (350.57)	0.07 (0.03)	0.19 (0.07)	0.62 (0.21)
Linear	RMSE	9.55 (4.65)	1295.78 (559.63)	0.09 (0.04)	0.23 (0.11)	0.74 (0.32)
Image Only	RMSE	7.85 (3.06)	1109.18 (392.08)	0.08 (0.03)	0.19 (0.07)	0.63 (0.22)
Mobile Only	RMSE	9.21 (3.50)	1267.92 (422.18)	0.08 (0.03)	0.22 (0.08)	0.72 (0.24)

Note: Mean values and the corresponding standard deviation (in parenthesis) over 10 cross-validation runs are reported.

Our model does consistently better than the linear model, across all the targets. Figure 11 highlights the performance of our model, which maps the non-linear relationships between the auxiliary data sources and the targets. Our model also does better

than using either of the datasets separately (mobile or imagery). For some of the targets, like income poverty, combining the different sources of data is beneficial.

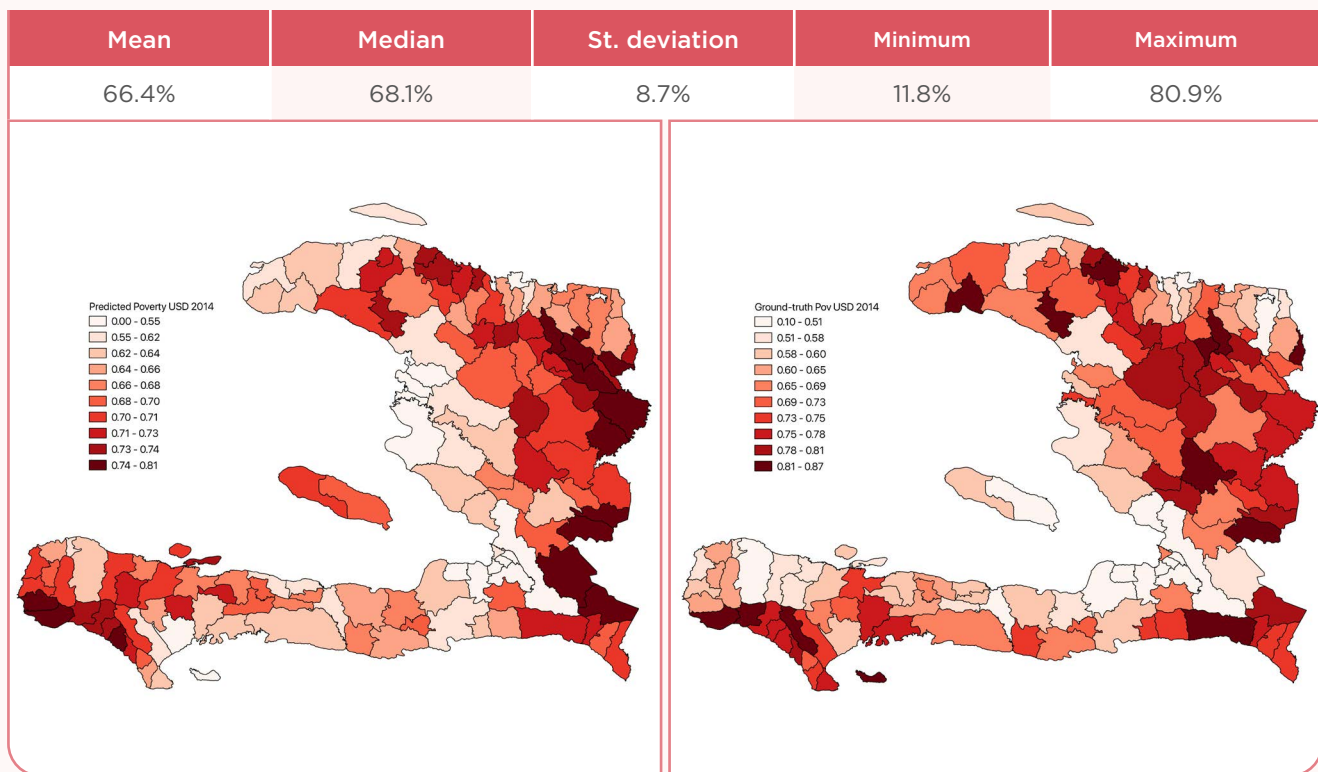
Figure 11: From Left: a) Denotes the comparison of actual and predicted Gini index values for all communes for 2014, b) Denotes the comparison of actual and predicted poverty (USD) values for all communes for 2014



Figures 12, 13, 14, and 15 display the cross-validated values of different social indicators of interest compared with the ground data (SAEs estimated from the census and ECVMAS). There are regions (like those around Port-au-Prince) where our model predicts values close to the ones derived from census/surveys (the SAE, or the ground-truth). However, there are also regions that feature marked differences between the predictions and the ground truth, which calls for further research²⁵. Nonetheless, on average, our predicted *commune* estimates had a different of only 1 percentage point difference with the ground-truth data for the income poverty and standards of living deprivation, and a 0.5 percentage point difference with the ground-truth, in the case of the income inequality estimates.

Figure 12 (left pane) shows the predicted income poverty at the *commune* level. The average *commune* has a poverty level exceeding 66%. The *commune* with the lowest level of predicted income poverty was Port-au-Prince in the Ouest department (with 11.8% of income poverty), while the *commune* with the highest predicted level of poverty was Cornillon/Grand Bois in the Ouest department (80.9%). Only eight out of the 138²⁶ examined *communes* for the income poverty prediction has poverty levels below 50%. High poverty is concentrated in regions within departments of Nord-Est and Centre in 2014. However, the *section communale* maps (which will be later discussed, starting on figure 16) confirm that there is heterogeneity in the spatial dispersion of poverty and associated deprivations²⁷.

Figure 12: 2014 Dynamic Maps of income poverty (in USD) at the *commune* level.
The maps display the quantiles of the predicted (left)
vs. the ground-truth data (right)



Note: Please notice that the values for each color range slightly vary between the predicted and ground-data maps.

²⁵ Tables A.2, A.3, and A.4 in the annex list the *communes* with more than 10 percentage point difference between the ground truth and the predicted values.

²⁶ While the model was estimated for 140 *communes*, it failed to estimate the income poverty values for Delmas and Tabarre.

²⁷ In fact, within the least poor department, the Ouest department, there are several microregions registering income poverty exceeding 50%.

Figure 13 shows the predicted and ground truth gini coefficients for Haitian *communes*. The average *commune* has a predicted gini coefficient of 0.639. Similar to the case of the SAE, our predicted data shows that only six *communes* have a gini coefficient of under 0.50, most of these in the

Ouest department. The *commune* with the lowest predicted gini coefficient is Delmas, while the one with the highest predicted gini coefficient is Cornillon/Grand Bois, both of these in the Ouest department.

Figure 13: 2014 Dynamic Maps of Gini Index at *commune* level. The maps display the quantiles of the predicted (left) vs. the ground-truth data (right).

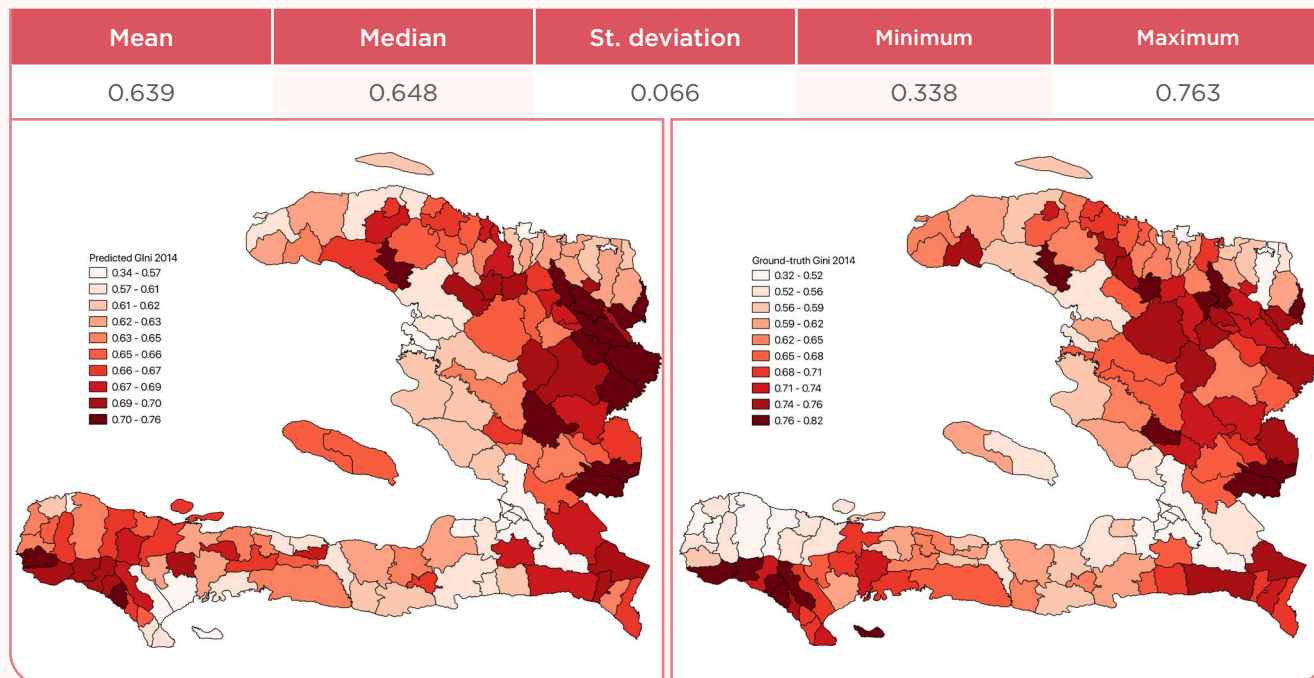


Figure 14: 2014 Dynamic Maps of FGT₂ Index (in USD) at *commune* level. The maps display the quantiles of the predicted (left) vs. the ground-truth data (right).

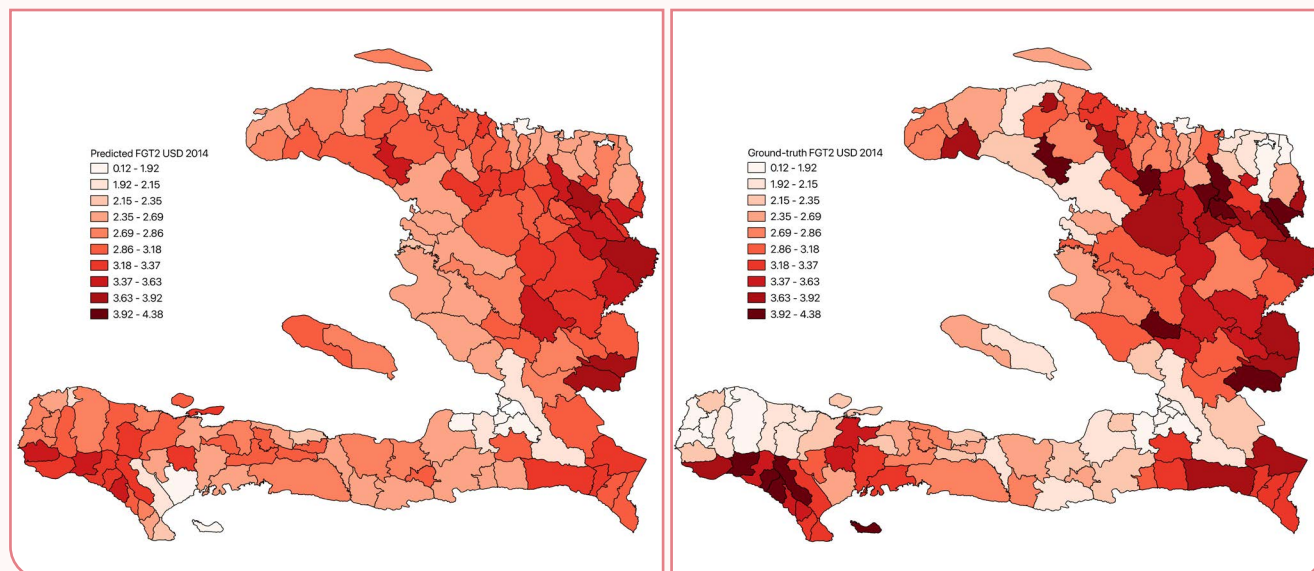
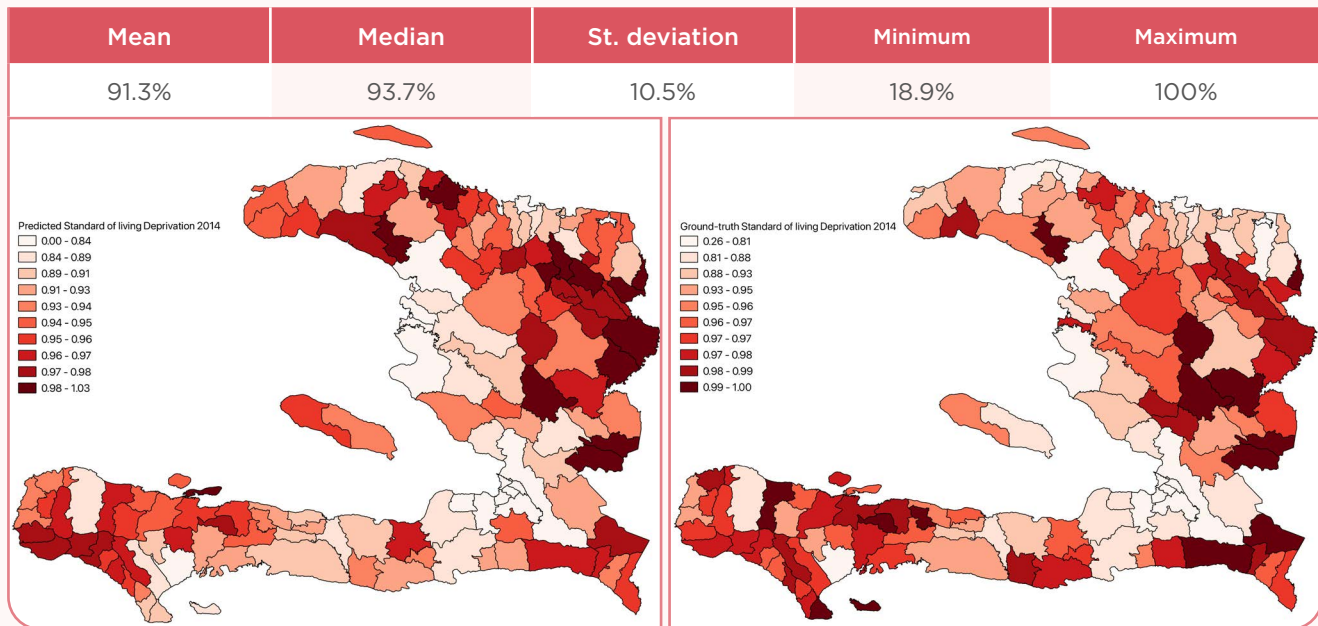


Figure 15 shows the predicted and ground-truth values for the Standards of Living deprivation at the *commune* level. Just as with the SAE, most *communes* register high levels of standards of living deprivation. The average *commune* is predicted

to have a 91.3% deprivation. The least deprived *commune* is Port-au-Prince²⁸ (18.9%), while the most deprived one is Vallières in the Nord-Est department (100%).

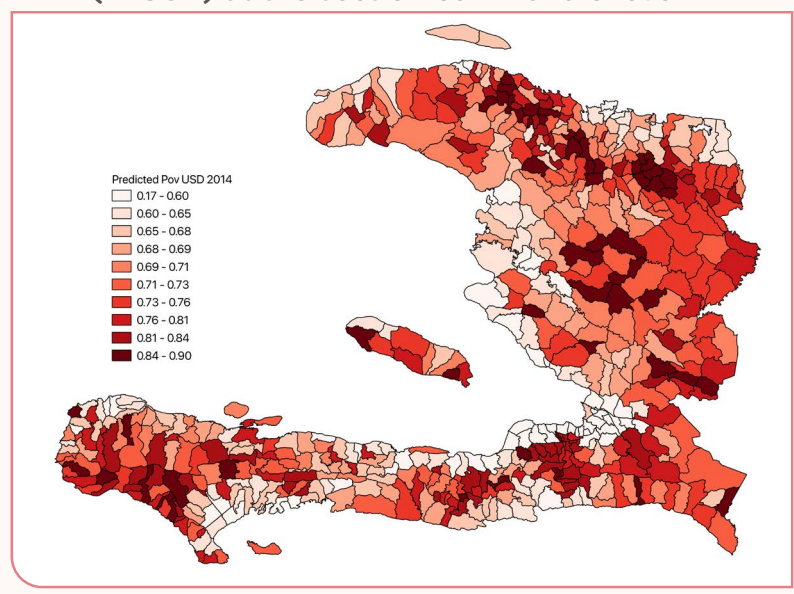
Figure 15: 2014 Dynamic Maps of Standards of Living Deprivation at *commune* level. The maps display the quantiles of the predicted (left) vs. the ground-truth data (right).



Maps at the *section communale* level

Figures 16-19 show the maps for the same four social indicators at the *section communale* level. These maps highlight the heterogeneity of socioeconomic deprivations that get masked when taking spatially aggregated statistics. However, as mentioned before, given that maps at this level of disaggregation have not been formally validated, these should be interpreted with caution²⁹ and are only presented to provide evidence of the heterogeneity within most *communes*.

Figure 16: 2014 Dynamic Map of Poverty (in USD) at the *section communale* level



²⁸ Notice however that Port-au-Prince is among the *communes* with the highest difference between the ground truth and estimates. The second least deprived *commune* is Pétionville (46.1%), followed by Croix-des-Bouquets (50%), both in Ouest.

²⁹ For this reason, no interpretation of the values or identification of the most affected *sections communales* will be provided

Figure 17: 2014 Dynamic Map of Gini Index at the *section communale* level

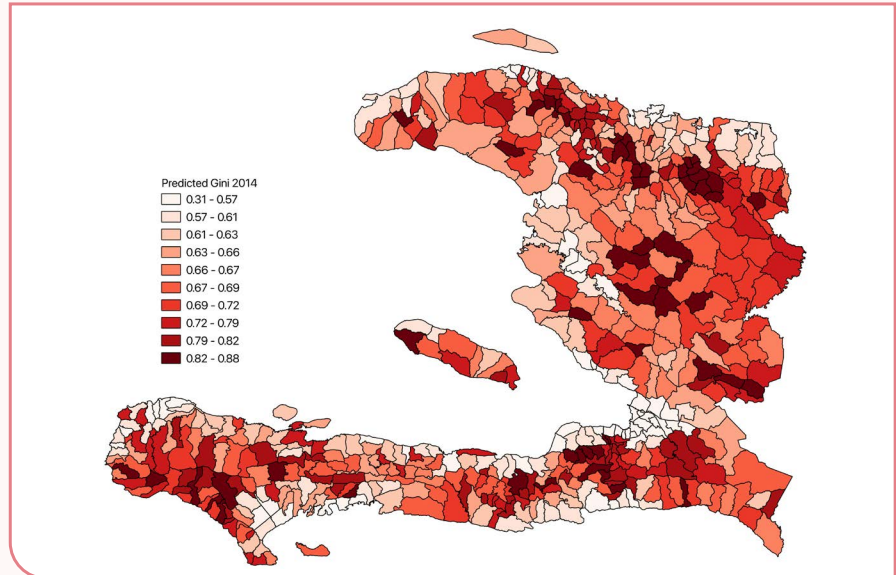


Figure 18: 2014 Dynamic Map of FGT₂ index (in USD) at the *section communale* level

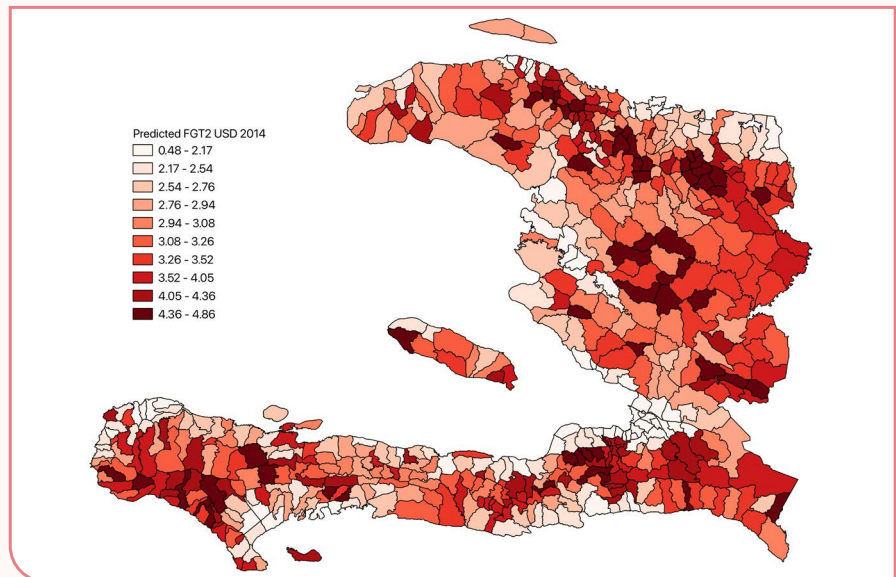
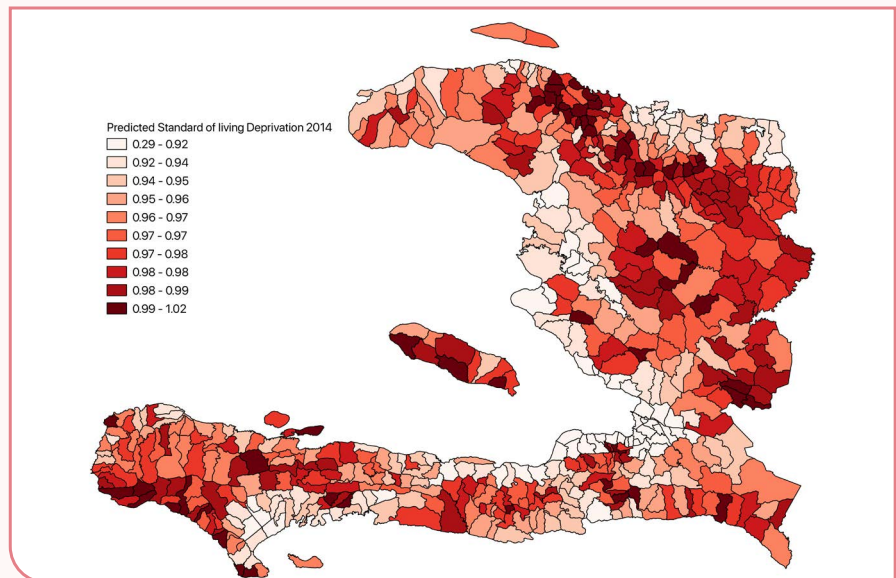


Figure 19: 2014 Dynamic Map of Standard of Living Deprivations at the *section communale* level

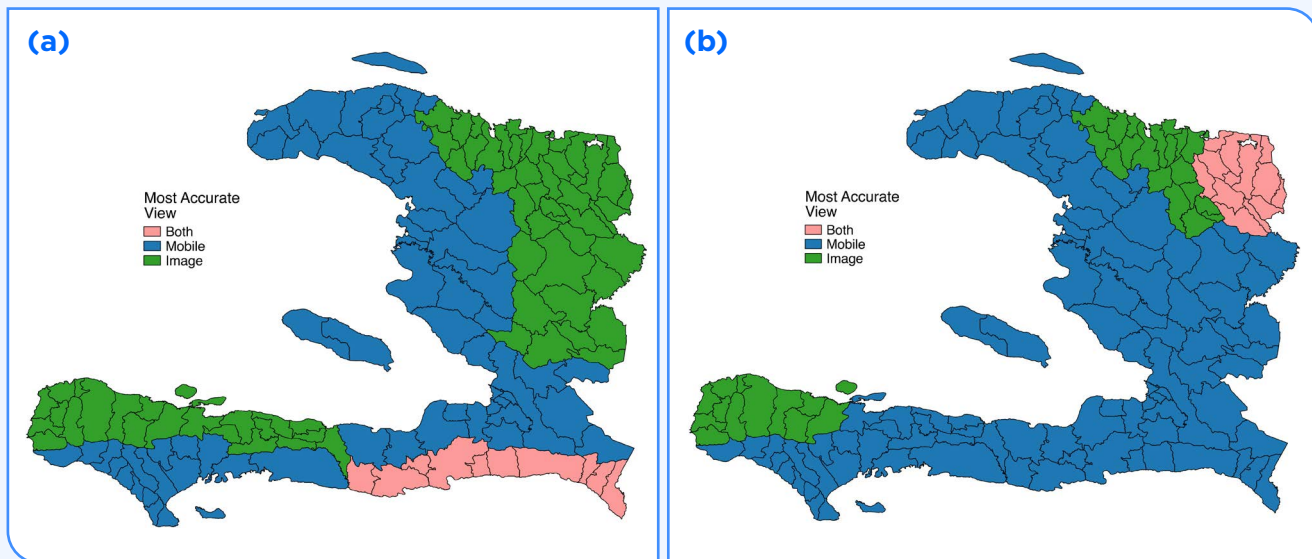


A note on dataset uncertainty

Figure 20 show the uncertainty associated with each dataset, which can assist policymakers when deciding which dataset is better in making poverty predictions for a given microregion. We see that mobile phone metadata holds better predictive

power, which is also reported in Pokhriyal and Jacques (2017). We also noticed a strong spatial correlation, which is understandable as poverty and its associated deprivations are spatially correlated.

Figure 20: Uncertainty associated with each data source evidenced by the most accurate one for income poverty (in USD) predictions (left) and standards of living deprivations (right) for 2014 (*commune* level).



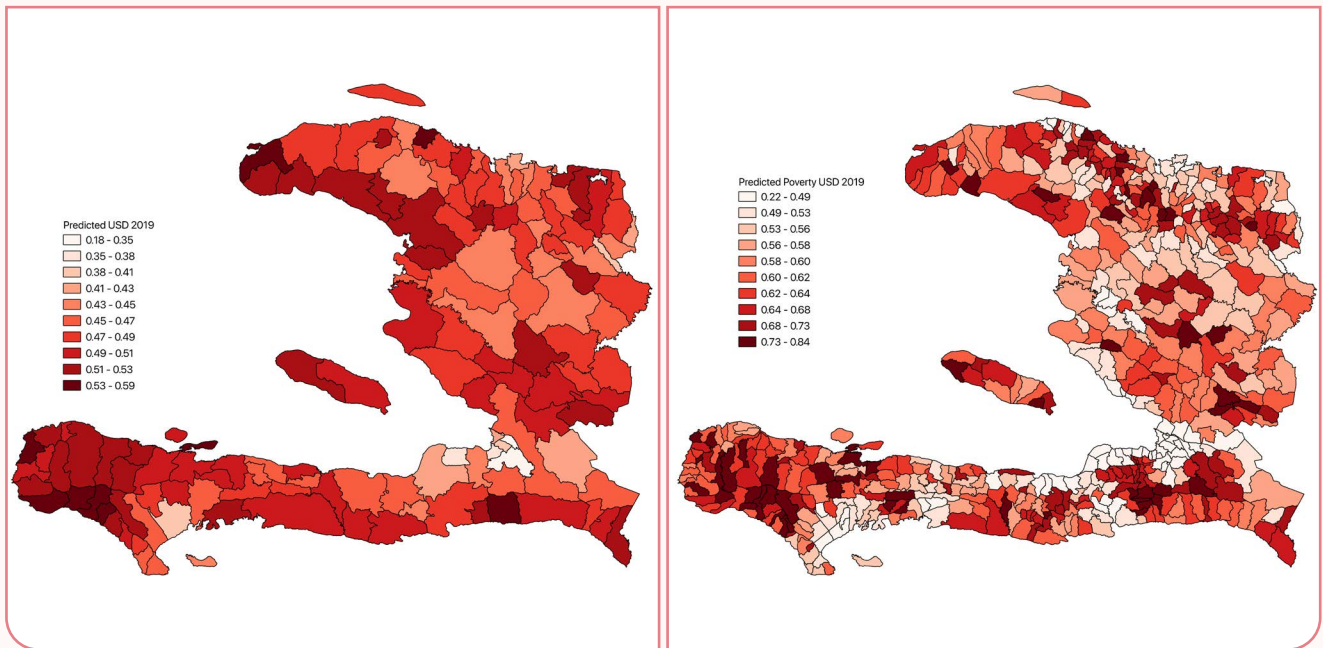
Forecasting 2019 deprivations with 2019 satellite imagery

Since both the satellite and mobile phone metadata are regularly available, the resulting model could be updated several times in-between the household surveys. To test the forecasting ability, we check how a model that is trained on learning the relationship between geospatial covariates and poverty for a specific point in time (say 2014) performs when newer data (say 2019) is inputted. Given that we did not have access to more recent mobile data, the 2019 maps were created exclusively with satellite imagery (specifically, using Google API satellite images for Haiti from 2019) (Gorelick et al., 2017). Each of the images was available at zoom level 16, which is equivalent to approximately 1 x 1 sq. km on the ground. Just like with the 2014 images, we used the night-light-based feature extractor to get features corresponding to each satellite imagery

and aggregated them to *commune* level, as well as the *section communale* level.

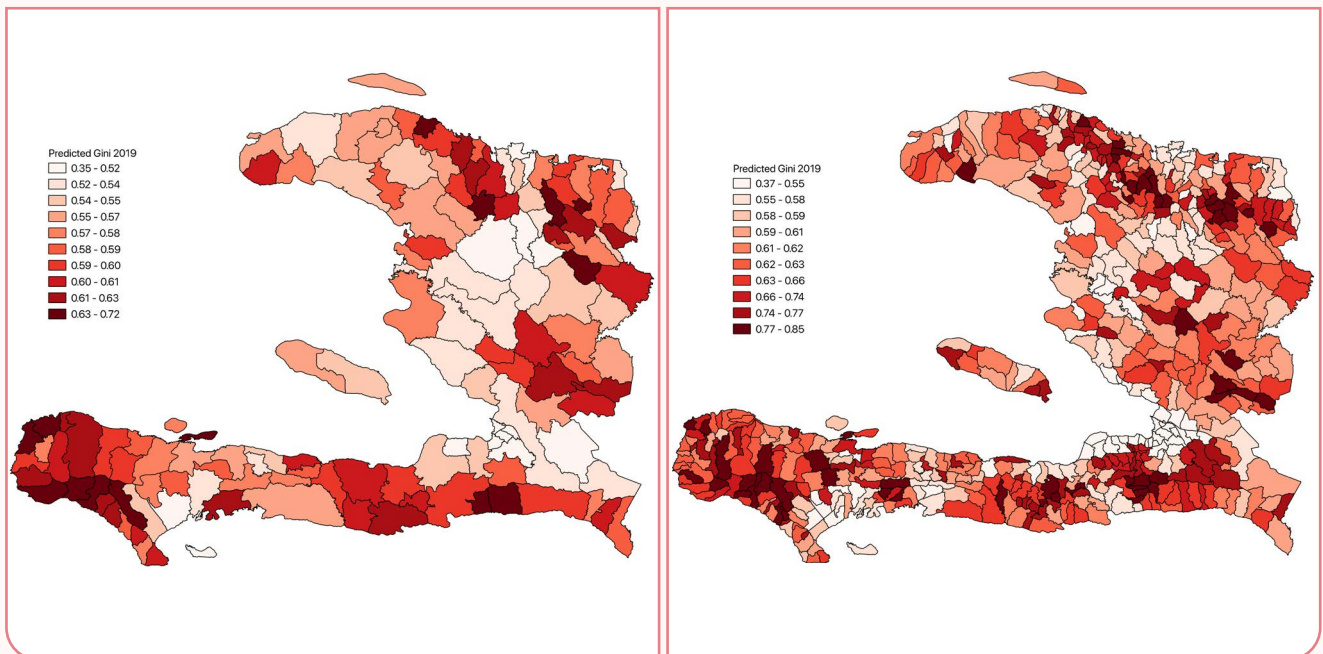
We used our model that is trained on 2014 satellite imagery at *commune* level. To predict the poverty and associated deprivations at the *commune* level in 2019, we provide our model with the 2019 satellite features aggregated at *commune* level during testing time. Similarly, to predict poverty and associated deprivations at the *section communale* level, we test/evaluate our model with aggregated *section communale* data for 2019. Figures 21-24 show the resulting maps at both microregional levels. Once again, notice the heterogeneity within every *commune*.

Figure 21: Income poverty (in USD) at the *commune* level (left) and *section communale* level (right) in 2019.



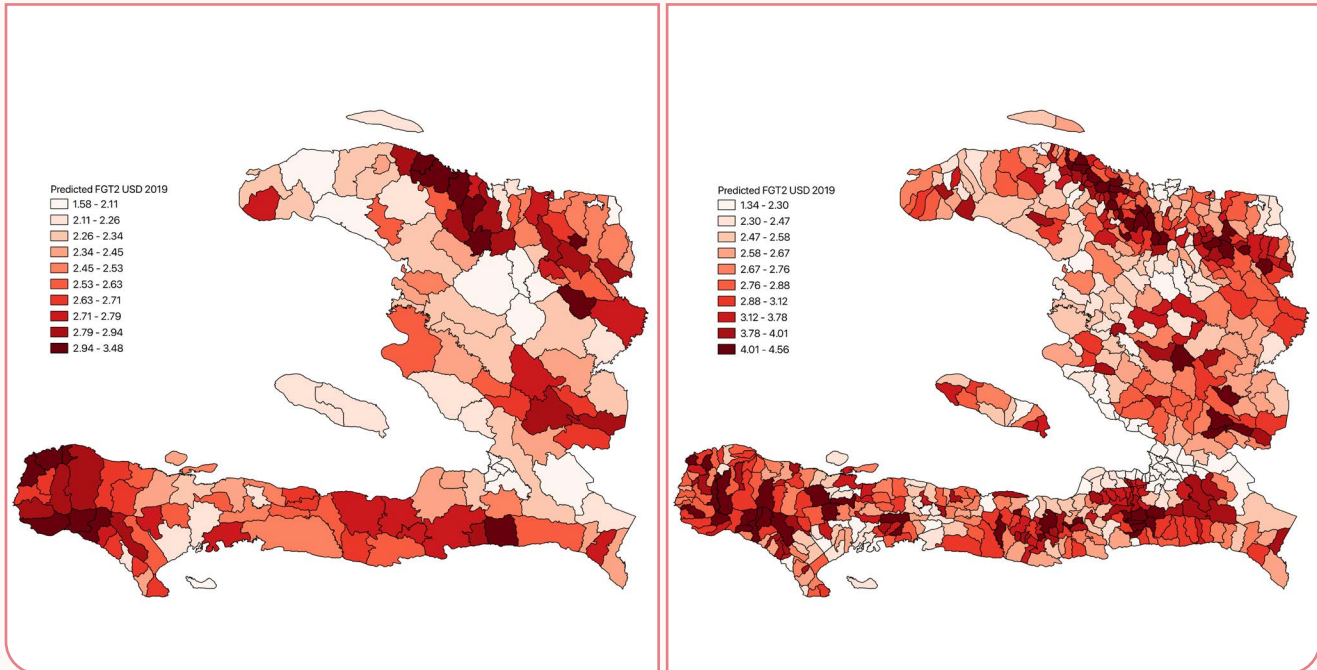
Note: notice that the colors and quantiles vary between maps.

Figure 22: Gini index at the *commune* level (left) and *section communale* level (right) in 2019.



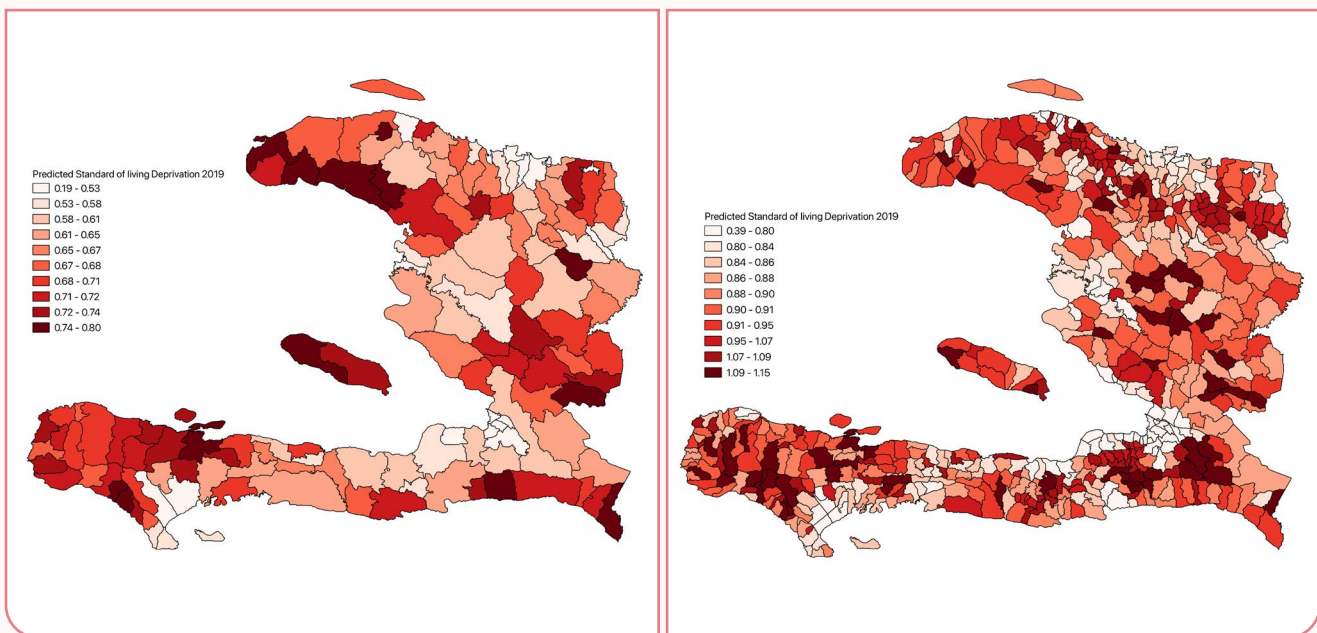
Note: notice that the colors and ranges vary between maps.

Figure 23: FGT₂ (in USD) at the *commune* level (left) and *section communale* level (right) in 2019.



Note: notice that the colors and ranges vary between maps.

Figure 24: Standards of living deprivations at the *commune* level (left) and *section communale* level (right) in 2019.



Note: notice that the colors and ranges vary between maps.

It is important to note that the forecasted maps from 2019 must be interpreted with caution, as these have not been validated by SAEs and were only elaborated with satellite imagery (not the mobile data). As a reminder, the combination of both datasets seems to perform best than images alone.

Nonetheless, despite the lack of validation for these last maps, several general trends can be observed when comparing both groups of maps:

1. Better-performing *communes* were usually located in the Ouest Department in both 2014 and 2019. The following *communes* consistently perform better than the rest in terms of poverty, inequality and social deprivations: Pétion-Ville, Port-au-Prince, Delmas, Cap-Haitien and Tabarre.

2. Nord-Ouest *communes* increasingly became poorer and more deprived in 2019 relative to 2014, when the worst-performing *communes* were mainly located in Nord-Est and Centre.

This is consistent with the latest findings of the World Food Programme's 2020 Global Report on Food Crisis, which shows that the population in the Nord-Ouest department is the one undergoing the most urgent food insecurity levels (IPC Phase 4).

3. Some *communes* in the Sud Department became more deprived in 2019 relative 2014.

This may be due to the impact that Hurricane Matthew had on these *communes* in late 2016.



5

Policies to tackle social deprivations in the COVID-19 era and the recovery period

The preceding sections have clearly demonstrated that a large portion of the Haitian population was in a vulnerable position prior to the arrival of the COVID-19 pandemic. The crisis associated with this pandemic will further deepen social disparities in Haiti if no comprehensive social protection plan is introduced. According to the IMF (Ghayad et al., 2019), to promote inclusive growth, a country has three tools at its disposal: progressive income taxation, health and education spending, and a social safety net. Nonetheless, Haiti's tax collection is low to begin with: in FY2019, it was estimated to

be 10.9% of GDP³⁰, and the Ministry of Economy and Finance expects it to be in the order of 11.6% of GDP in FY2020. Haiti's tax system also tends to be regressive, relying heavily on indirect taxes. In FY2019, income taxes amounted to only 26% of total tax revenues, while the TCA (sales tax) and custom taxes were equivalent to 28% and 26.3%, respectively³¹.

Further, Haiti's health and education spending are relatively low: according to Ghayad et al., 2019, Haiti's spending in these two categories summed up to less than 4% of GDP in FY2018, well below the regional average (over 8% of GDP). With the arrival of COVID-19, the Haitian government increased its health spending allocation by a factor of four, bringing the budget allocation of the Ministry of Public Health and Population (MSPP) to 10.9% of the total FY2020 budget.

In terms of social safety nets, several efforts toward creating social programs have been rolled out over the past five years, yet the social protection system has remained notably fragmented, and have consisted mainly of tuition assistance and nutritional support programs funded by various development and humanitarian organization (Ghayad et al., 2019). In 2014, an Action Plan for Reduction of Poverty (PAARP) was launched, which introduced a set of social programs under a social assistance umbrella called *Ede Pep*³², and included a national fee waiver program for basic education (PSUGO). In 2016, a Sectorial Table of Social Protection was introduced by the Ministry of Social Affairs and Labor (MAST),

30 This figure has been falling from its value in the previous years, as it was, on average, 13.4% of GDP between FY2015-2018.

31 One could argue that this could be related to the socioeconomic crisis in FY2019, yet this regressivity has also been present over (at least) the ten preceding years.

32 In total, EDE PEP includes 11 social programs to help mothers of schoolchildren and university students, vulnerable and food insecure households, and farmers and households affected by natural disasters.

which consisted of consultations with stakeholders nationwide in order to draft a national social protection policy. In addition, the National Strategic Development Plan (PDSH) — which intended to transform Haiti into an emerging country by 2030 — has a social reFOUNDATION pillar which contemplates modern health and education networks accessible to all Haitians, as well as policies of social inclusion and gender equality.

More recently, according to the 2020 Article IV Staff Report of the IMF with Haiti, a National Plan for Social Protection and Progress (PNPPS) was being finalized. In this staff report, the IMF recommends that the PNPPS reduces the fragmentation and overlap of existing programs, and that a limited number of unconditional, quasi universal cash transfer programs for vulnerable groups is established. Cash transfer programs that are simple in design have proven to be effective in other low-income countries and have the potential to be introduced in a quicker fashion than conditional cash transfers. Further, Banerjee et al., 2019 demonstrated that more refined targeting programs tend to be less effective in countries unable to identify beneficiary or lacking capacity to implement such targeted programs, thus providing evidence in favor of simplified, unconditional cash transfers in countries like Haiti.

With the arrival of COVID-19, the launching of a set of unconditioned cash transfer programs under the PNPPS umbrella is more essential than ever to provide quick relief to the most vulnerable segments of the population. As part of the measures to mitigate the impact of COVID-19 on the most vulnerable, the Haitian government distributed unconditional cash transfers³³ (via a mobile payment system) and food kits to 1.5 million vulnerable households. Nonetheless, the distribution of the cash transfers faced several challenges, including the lack of a comprehensive and accurate list of the beneficiaries' phone numbers, the fact that many target beneficiaries' telecom provider was not linked to the mobile payment system³⁴, and the fact that a number of beneficiaries had no active mobile

wallet. These challenges must be tackled in order to ensure the effective distribution of the PNPPS cash transfers.

Further, the following actions should be considered when preparing to roll out the PNPPS: (i) the expansion and update of the Information System (SIMAST) of the MAST so that there is a comprehensive source of information that can be consulted to locate beneficiaries. Our poverty maps could be used as a first step to try to locate beneficiaries at a more disaggregate level (ii) the identification of clear and reliable means of distribution or payments depending on the type of benefit, beneficiaries and their location (e.g. mobile payments might work in the city, but in the remote areas, for instance, it would be necessary to distribute cash through local credit unions); and (iii) a clear articulation of the role of the various actors for each step of the intervention, from the definition of the benefits to the deployment of these, as Haiti currently relies on a variety of actors, including financial operators and NGOs.

Finally, it will be important to consider making complementary investments such as transport infrastructure, in order to reduce the high cost of mobility for Haitians living in remote areas. This will be key in places such as the Nord-Ouest department, which, as shown in figure 9 above, is practically disconnected from the rest of the country. According to Hausmann (2015), there are enormous productivity differences among regions within several countries. These differences are explained by the fact that some regions within a country are better endowed with inputs for productivity to take place (such as electricity, more readily available forms of transportation, and the convergence of different talents). Therefore, the poor are potentially being excluded from the higher productivity areas of the country. Investing on quality infrastructure will allow for Haitians to connect to the rest of the productive sectors, will result in productivity gains, and could ultimately be associated with a more inclusive growth.

³³ Equivalent to about US\$30.

³⁴ Only one of the two mobile carriers (Digicel) is connected to the mobile payment system.

6 Conclusion

Given the high costs associated with collecting household survey data, reporting poverty and other social indicators on a regular basis is a challenge in several developing countries. The present report adds to the poverty literature on Haiti by disaggregating 2012 estimates for income poverty, income inequality, and standard of living deprivations at the *commune* level (previously, and to the best of our knowledge, these were only available at the departmental level). These estimates were then used to validate the prediction of these values for 2014, which were estimated using features extracted from aerial imagery and anonymized call detailed records straddling 2016, 2017, and 2018 (thus making the 2014 values dynamic in nature).

The 2012 disaggregate estimates (called Small Area Estimates or SAEs throughout the report) were obtained using a model based using the Gradient Boosting Machine technique. The national poverty rate estimated from this model is 57.1%, which is very close to the 2012 official national poverty rate using department level data (58.5%). The SAEs allowed for us to determine that one out of four Haitian living in poverty resided in ten *communes* (out of the 140 examined *communes*): Gonaïves, Cité Soleil, Port-au-Prince, Saint-Marc, Cap-Haïtien, Carrefour, Dessalines, Petite Riviere, Saint-Michel de l'Attalaye and Port-de-Paix. Further, there is a concentration of *communes* with very high levels of poverty in the eastern side of Artibonite, the southernmost *communes* in the Sud and Sud-Est departments, and in the Centre department. Moreover, Among the 10 *communes* with the highest poverty rate in 2012, six were located in the northern region. The *communes* with the most intense poverty levels in 2012 (measured by the FGT2 index) were located in the Sud and Artibonite departments. The *communes* with less poverty intensity were those located in the Ouest department (Port-au-Prince, Pétion-Ville, Carrefour, Delmas) and Cap-Haïtien (Nord).

Income inequality is rampant in Haiti, with only three out of the 140 *communes* examined in this exercise with a gini coefficient of less than 0.50. Furthermore, in terms of standards of living deprivations, most *communes* registered high levels of deprivation (over 90%, on average). The *communes* with the least deprivations were mostly located in the Ouest department (Delmas, Port-au-Prince, Carrefour, Pétion-Ville, and Tabarre).

The 2014 dynamic maps and the 2019 predictions were estimated using a model based on uncertainty-weighted multi-view Gaussian process regressions. These predictions were presented at the *commune*

and the *section communale* levels. The results from this process are consistent with the SAEs, although only the 2014 dynamic maps at the *commune* level were validated with the ground-truth data. Similar to the SAEs, the 2014 validated estimates show that only eight out of the *communes* for which predictions were available had poverty levels below 50%, and that only six out of the 140 *communes* for which predictions were available had gini coefficients under 0.50. Further, high standards of living deprivations continued to be a trend present in the 2014 dynamic predictions.

The *section communale* maps, which should be interpreted with caution due to the lack of validation, show high heterogeneity of income poverty, income inequality, and standard of living deprivations within each *commune*. Further, when comparing the 2014 dynamic maps with the 2019 maps (the latter which also lack validation), the following general trends were observed: a) Better-performing *communes* were usually located in the Ouest Department in both 2014 and 2019. These *communes* were: Pétion-Ville, Port-au-Prince, Delmas, Cap-Haïtien and Tabarre ; b) Nord-Ouest *communes* increasingly became poorer and more deprived than the rest in 2019 relative to 2014, when the worst-performing *communes* were mainly located in Nord-Est and Centre; and c) some *communes* in the Sud Department became more deprived than the rest in 2019 relative 2014. This may be due to the impact that Hurricane Matthew had on these *communes* in late 2016.

Our maps provide evidence that Haiti needs a comprehensive plan for inclusive growth, but this will require that more resources be mobilized toward social spending and that the Haitian tax system is made more progressive. More importantly, in the context of COVID-19, the rolling out of a set of simple, unconditioned cash transfers (via the PNPPS) is more urgent than ever in order to mitigate the disproportionate effect that the pandemic will have on the most vulnerable. Nonetheless, several logistical challenges must be tackled first to guarantee the timely delivery of these transfers. The poverty and inequality maps presented in this report represent a powerful tool for the identification of the most vulnerable at a more disaggregate level and provide the framework for further update with more recent data and for other social protection applications. Finally, it will be important to consider important infrastructure investments in order to better connect some of the most remote areas in Haiti to more productive regions of the country. All else constant, such actions will result in productivity gains and thus, in a more inclusive growth in Haiti.



Bibliography

- Alderman, H., Babita, M., Demombynes, G., Makhatha, N. & Özler, B. (2002). How low can you go? Combining census and survey data for mapping poverty in South Africa. *J. Afr. Econom.*, 11(2), 169– 200.
- Bahn, V. and B. J. McGill. Testing the predictive performance of distribution models. *Oikos*, 122(3): 321–331, 2013.
- Banerjee, Abhijit, Paul Niehaus, and Tavneet Suri. 2019. Universal basic income in the developing world. *Annual Review of Economics* 11.
- BBS & UNWFP. Local Estimation of Poverty and Malnutrition in Bangladesh.* Dhaka: The Bangladesh Bureau of Statistics and The United Nations World Food Programme. 2004.
- Blondel, V.D., M. Esch, C. Chan, F. Cl´erot, P. Deville, E. Huens, F. Morlot, Z. Smoreda, and C. Ziemlicki. Data for development: the d4d challenge on mobile phone data. *arXiv preprint arXiv:1210.0137*, 2012.
- Blumenstock, J., G. Cadamuro, and R. On. Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264):1073–1076, 2015.
- Chambers, R. & Tzavidis, N. (2006). M-quantile models for small area estimation. *Biometrika*, 93(2), 255– 268.
- Christie, G., N. Fendley, J. Wilson, and R. Mukherjee. Functional map of the world. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- Cressie, N. The origins of kriging. *Mathematical Geology*, 22(3), 1990.
- Das, S. and Haslett, S., 2019. A Comparison of Methods for Poverty Estimation in Developing Countries. *International Statistical Review*, 87(2), pp.368-392.
- Demombynes, G., Lanjouw, J.O., Lanjouw, P. and Elbers (2006) “How Good a Map? Putting Small Area Estimation to the Test”, Policy Research Working Paper No. 4155, The World Bank.
- Deville, P., C. Linard, S. Martin, M. Gilbert, F. R. Stevens, A. E. Gaughan, V. D. Blondel, and A. J. Tatem. Dynamic population mapping using mobile phone data. *Proceedings of the National Academy of Sciences*, 111(45):15888–15893, 2014.
- Elbers, C. & Lanjouw, J. O. & Lanjouw, P., (2002). “Micro-level estimation of welfare,” Policy Research Working Paper Series 2911, The World Bank.
- Elbers, C., Lanjouw, J.O., Lanjouw, P. & Leite, P.G. (2004). Poverty and inequality in Brazil: new estimates from combined PPV-PNAD. *In Inequality and Economic Development in Brazil*, pp. 81–104. Washington, D.C.: World Bank.
- Elbers, Chris, Peter F. Lanjouw, and Phillippe G. Leite. “Brazil within Brazil: Testing the poverty map methodology in Minas Gerais.” World Bank Policy Research Working Paper Series, Vol (2008).

- Elvidge, C. D., Sutton, P. C., Ghosh, T., Tuttle, B. T., Baugh, K. E., Bhaduri, B., & Bright, E. (2009). A global poverty map derived from satellite data. *Computers & Geosciences*, 35(8), 1652-1660.
- Engstrom, R.; J. Hersh, and D. Newhouse. Poverty from space: using high-resolution satellite imagery for estimating economic well-being. Policy Research working paper, World Bank, 2017.
- Fernandez-Taranco, O., D. Keita, P. Hiebra, J. Le Nay, A. Ambroise, P. Rouzier, C. Cadet. Haiti and Governance for Human Development. 2013. United Nations Development Programme.
- Flowminder final consulting report: Caracol Industrial Park and National Road Network. An analysis of commuting and migration patterns (Mimeo)
- Fujii, T. (2004). *Commune-Level Estimation of Poverty Measures and its Application in Cambodia*, Vol. 2004/48 Helsinki: UNU-WIDER, United Nations University (UNU).
- Giusti, C., Marchetti, S., Pratesi, M. & Tzavidis, N. (2011). Small area methodologies for poverty estimation: an application to Italian data. In *International Statistical Institute: Proceedings of the 58th World Statistical Congress*, Dublin, pp. 6234- 6237. Haslett, S. & Jones, G. Estimation of Local Poverty in the Philippines. Philippines National Statistics Co-ordination Board / World Bank, November 2005.
- Ghayad, R., F. Lambert, M. Rousset, and M. Bellon (2019). Haiti Selected Issues: IMF Country Report No. 20/122; December 23, 2019
- Gorelick, N. et al. Google earth engine: Planetary-scale geospatial analysis for everyone. in *Remote Sensing of Environment* (Elsevier, 2017).
- Haslett, S. & Jones, G. *Small Area Estimation of Poverty, Caloric Intake and Malnutrition in Nepal*. Kathmandu: Nepal Central Bureau of Statistics/World Food Programme, United Nations/World Bank. 2006.
- Haslett, S., Jones, G. & Isidro, M. (2014). *Small-area Estimation of Child Undernutrition in Bangladesh*. Dhaka: Bangladesh Bureau of Statistics, United Nations World Food Programme and International Fund for Agricultural Development. ISBN 978-984-33-9085-1.
- Hausmann, R., 2015. What Should We Do About Inequality?. Copy at <http://www.tinyurl.com/y6byadd3>
- Head, A., M. Manguin, N. Tran, and J. E. Blumenstock. Can human development be measured with satellite imagery? In *Proceedings of the Ninth International Conference on Information and Communication Technologies and Development*, ICTD '17, 2017.
- Healy, A.J., Hitsuchon, S. and Vajaragupta, Y., 2003. Spatially disaggregated estimates of poverty and inequality in Thailand. *Massachusetts Institute of Technology and Thailand Development Research Institute*.
- Hersh, J., R. Engstrom, M. Mann, A. Mejia and L. Martin. Mapping Income Poverty in Belize Using Satellite Features and Machine Learning. Inter-American Development Bank. 2020.
- Huang, G., Z. Liu, and K. Q. Weinberger. Densely connected convolutional networks. *CoRR*, 2016.
- Jean, N., M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790-794, 2016.
- Kilic, Talip; Serajuddin, Umar; Uematsu, Hiroki; Yoshida, Nobuo. 2017. *Costing household surveys for monitoring progress toward ending extreme poverty and boosting shared prosperity (English)*. Policy Research working paper; no. WPS 7951; LSMS. Washington, D.C.: World Bank Group. <http://documents.worldbank.org/curated/en/260501485264312208/Costing-household-surveys-for-monitoring-progress-toward-ending-extreme-poverty-and-boosting-shared-prosperity>

- Marchetti, S., Beresewicz, M., Salvati, N.S. & Wawrowski, L. (2018). The use of a three-level M-quantile model to map poverty at local administrative unit 1 in Poland. *J. R. Stat. Soc.: Ser. A (Stat. Soc.)*, 181(3), 1077– 1104.
- Molina, I. & Rao, J.N. (2010). Small area estimation of poverty indicators. *Can. J. Stat.*, 38(3), 369– 385.
- Pandey, S. M., T. Agarwal, and N. C. Krishnan. Multi-task deep learning for predicting poverty from satellite images. In IAAI, 2018.
- Pokhriyal, N. and W. Dong. Virtual network and poverty analysis in Senegal. *D4D Challenge Senegal Scientific Papers, Netmob*, 2015.
- Pokhriyal N., and D.C. Jacques. Combining disparate data sources for improved poverty prediction and mapping. Proceedings of the National Academy of Sciences, 2017.
- Rasmussen, C.E. and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Rodríguez Castelán, C., Ingmar Weber, Damien Jacques, and Trevor Monroe. “Making a better poverty map”. World Bank. 2019. <https://blogs.worldbank.org/opendata/making-better-poverty-map>.
- Santos, M. & Villatoro, P. (2018): A Multidimensional Poverty Index for Latin America. Review of Income and Wealth, Vol. 64, Issue 1, pp. 52-82, 2018
- Solt, Frederick. 2019. “Measuring Income Inequality Across Countries and Over Time: The Standardized World Income Inequality Database.” SWIID Version 8.2, November 2019.
- Soto, V., V. Frias-Martinez, J. Virseda, and E. Frias-Martinez. Prediction of socioeconomic levels using cell phone records. In *Proceedings of the 19th International Conference on User Modeling, Adaption and Personalization*, pages 377–388. Springer, 2011.
- Steele, J.E., P. R. Sundsøy, C. Pezzulo, V. A. Alegana, T. J. Bird, J. Blumenstock, J. Bjelland, K. Engø-Monsen, Y.-A. de Montjoye, A. M. Iqbal, K. N. Hadiuzzaman, X. Lu, E. Wetter, A. J. Tatem, and L. Bengtsson. Mapping poverty using mobile phone and satellite data. *Journal of The Royal Society Interface*, 14(127), 2017.
- Tang, B. V., Y. Sun, Y. Liu, and D. S. Matteson. Dynamic poverty prediction with vegetation index. Workshop on Modeling and Decision-Making in the Spatiotemporal Domain, 32nd Conference on Neural Information Processing Systems. 2018.
- Tzavidis, N., Salvati, N., Pratesi, M. & Chambers, R. (2008). M-quantile models with application to poverty mapping. *Stat. Methods Appl.*, 17(3), 393– 411.
- Watmough, G. R., C. L. J. Marcinko, C. Sullivan, K. Tschirhart, P. K. Mutuo, C. A. Palm, and J.-C. Svenning. Socioecologically informed use of remote sensing data to predict rural household poverty. *Proceedings of the National Academy of Sciences*, 2019.
- World Bank Group. Investing in people to fight poverty in Haiti. 2014.
- World Food Programme. 2020 - Global Report on Food Crises. <https://www.wfp.org/publications/2020-global-report-food-crises>
- United Nations Development Programme. 2019. “Global Multidimensional Poverty Index 2019: Illuminating Inequalities”. <http://hdr.undp.org/en/2019-MPI>
- Xu, C., D. Tao, and C. Xu. A survey on multi-view learning. *arXiv*, abs/1304.5634, 2013



Appendix

Table A.1

List of Common Variables	
Home ID	Access to public water service network
Haitian Department (categorical variable ranging from 1 to 10)	Access to treated water
Total household income	Water is drawn from wells
Final official weight	Water is obtained from rain collection
Family income equivalent home	Water is obtained from other sources
Total income of the household with imputations	Water is obtained from river
Natural logarithm of the total income of the household with imputations	Dependency ratio 1. Total people of non-working age as a percent of the total people of working age
Family income equivalent with imputations	Dependency ratio 2. People who do not work as a percent of the people of working age who are currently working
Natural logarithm of household income equivalent charges	Number of women at home
Persons per household	Number of children at home
Overcrowding. Persons per household on all rooms of a dwelling	Number of people working in the public sector
Thatched roof	Number of people working
Cement roof	Number of people working in a cooperative
Plastic roof	Number of people working in an NGO
Tin Roof	Number of people working in other jobs
Tile roof	Number of teenagers at home
Other type of roof	Number of children at home
Wooden walls	Number of people with elementary school education
Mud walls	Number of people who are employers at their jobs
Concrete walls	Number of people who are self-employed
Plate walls	Number of people who are home helpers
Cardboard walls	Number of people who are interns
Brick walls	Number of people who are on a payroll
Glisse walls	Number of people with high school level education
Other type of walls	Number of people without education
Wooden floor	Number of people with higher education

Earthen floor	Number of people who are literate
Concrete floor	Number of people who are studying
tiled floor	Number of people over 65 who do not work
Ceramic floor	Number of household members over 15 who do not work
Other type of floor	Average age of home occupants
Kay house	The head of the household in women
slum (house built with materials from other buildings)	There is a spouse at home
Ajoupa (hut)	Head of the household has primary education
low-rise house	Head of the household has tertiary education
One-story house	Head of the household has secondary education
Apartment	Spouse has primary education
Villa house	Spouse has tertiary education
Other type of dwelling	Spouse has secondary school education
Access to public waste disposal collection services	Age of the head of household
Waste disposal through a private service	Head of the household works as a civil servant
Waste disposal in vacant lot	Head of the household works as a domestic worker
Waste disposal in ravines	Head of the household works in a cooperative
Waste disposal into sewers	Head of the household works in an NGO
Waste disposal on public roads	Head of the household works in another place
Waste disposal into the sea	Spouse works as a civil servant
Waste incineration	Spouse works as a domestic worker
Cooks with wood	Spouse works in a cooperative
Cooks with propane	Spouse works in an NGO
Cooks with electric stove	Spouse works in another place
Cooks with kerosene	Head of the household is a salaried worker
Cooks with charcoal	Head of the household is an employer
Cooks with solar oven	Head of household is self-employed
Cooks with other methods	Head of the household is a family assistant
Home owner	Head of the household is intern
Tenant	Spouse is a salaried worker
Home is a farm	Spouse is an employer
Inhabits a home for free	Spouse is self-employed
De-facto occupant	Spouse is assistant and self-employed
Other type of house occupant	Spouse is an intern
Number of rooms per dwelling	Type of locality
Unknown ID	Weight
Individual weighting per home. Under 12 worth .5 and 13 to 18 are worth .75	Post stratification
The water comes from private source	

BOX A.1: An explanation of Call Detailed Records (CDR)

The data used for the estimation of our maps was provided by Flowminder, a non-profit organization that collects, aggregates, integrates and analyzes anonymous mobile operator data, satellite and household survey data across several countries. This organization was hired to elaborate a report for the IDB on commuting and migration patterns of the Haitian population using Digicel call detailed records (CDR).

According to Flowminder, CDR data is generated by telecommunications providers and are the basis for the customer billing process. An operator's network consists of a group of base stations which route signals from an initiating mobile device to a receiving mobile device. A single base station can consist of multiple cells (receivers) operating at different mobile technology generations (For instance: 3G and 4G). A base station's geographic location is usually fixed, but on some occasions, specially designed, mobile base stations – also known as cells on wheels – can be used to provide extra capacity at particular locations. A base station can only receive signals from mobile devices in a constrained region around its physical location. The spatial density of base stations can vary widely. In Port-Au-Prince, this density can reach up to six base stations per square kilometer, while in rural areas, this density can be significantly lower with a single base station covering many square kilometers.

A mobile device is uniquely identified by its Mobile Station International Subscriber Directory Number (MSISDN) and International Mobile Subscriber Identifier (IMSI). In order to bill its customers, a telecommunication operator maintains a database of CDR. Every voice call results in the following information stored in the database:

- Start time of the call
- Duration of the call
- MSISDN of initiating party
- Location ID of cell used by initiating party
- MSISDN of receiving party
- Location ID of cell used by receiving party.

Sometimes, additional information such as the handset type is included. Due to the highly private nature of CDR data, Flowminder employed a series of anonymization techniques to reduce the risk of individual or small group disclosure. One of these measures was to “hash” unique identifiers (such as MSISDN). The hashing process was undertaken by telecom staff before Flowminder had access to the data. The infrastructure used for analysis was securely maintained behind telecommunication operator firewalls to reduce the risk of unauthorized access to the raw data.

Table A.2: Communes with more than 10 percentage point difference between the ground truth and the estimates (income poverty)

Departement	Commune	Difference
Sud	Île à Vache	0.37
Artibonite	Grande Saline	0.27
Nord-Ouest	Baie de Henne	0.19
Artibonite	L'Estère	0.15
Artibonite	Terre Neuve	0.14
Sud	Torbeck	0.14
Artibonite	La Chapelle	0.14
Sud	Saint louis du Sud	0.14
Sud	Chantal	0.13
Centre	Saut d'Eau	0.12
Nord	Plaisance	0.12
Sud	Saint Jean du Sud	0.12
Ouest	Cité Soleil	0.12
Artibonite	Saint-Michel de l'Attalaye	0.12
Sud	Cavaillon	0.11
Nord	Pignon	0.11
Nord	Bahon	0.11
Nord-Est	Capotille	0.11

Table A.3: Communes with more than 10 percentage point difference between the ground truth and the estimates (gini coefficient)

Departement	Commune	Difference
Sud	Île à Vache	0.35
Artibonite	Grande Saline	0.20
Sud	Torbeck	0.14
Sud	Chantal	0.13
Nord	Bahon	0.13
Artibonite	La Chapelle	0.13
Nord	Plaisance	0.13
Sud	Saint louis du Sud	0.12
Nord	Pignon	0.11
Sud	Port-Salut	0.11
Sud	Saint Jean du Sud	0.11

Table A.4: Communes with more than 10 percentage point difference between the ground truth and the estimates (standard of living deprivation)

Departement	Commune	Difference
Ouest	Delmas	0.26
Ouest	Port-au-Prince	0.22
Artibonite	Grande Saline	0.19
Ouest	Croix-Des-Bouquets	0.18
Ouest	Cité Soleil	0.12
Sud	Île à Vache	0.11

Details of the Methodology

The data generative process for our Gaussian Process based regression given as:

$$f(\mathbf{x}) \sim GP(\mathbf{m}(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$$

$$y_i \sim N(\beta^T \mathbf{x}_i + f(\mathbf{x}_i), \sigma_n^2), \forall i$$

A *GP* is a stochastic process, indexed by $\mathbf{x} \in \mathbb{R}^d$. Any finite sample generated from it is jointly multivariate normal (Rasmussen and Williams, 2006). $\mathbf{m}(\mathbf{x})$ is the mean of $f(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ is a kernel function that defines the covariance between any two evaluations of $f(\mathbf{x})$, i.e., $\mathbf{m}(\mathbf{x}) = E[f(\mathbf{x})]$, and $k(\mathbf{x}, \mathbf{x}') = E[(f(\mathbf{x}) - \mathbf{m}(\mathbf{x}))(f(\mathbf{x}') - \mathbf{m}(\mathbf{x}'))]$. For model simplicity, we assume that $\mathbf{m}(\mathbf{x}) = 0$, which is a standard practice in GP based methods.

Given a training set of examples, $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, the GP prior on $f(\cdot)$, and other terms in the equation, the posterior distribution of \mathbf{y}_* (for an unseen input vector, \mathbf{x}_*), is a Gaussian distribution, with the following mean and variance:

$$\bar{\mathbf{y}}_* = E[\mathbf{y}_*] = \beta^T \mathbf{x}_* + \mathbf{k}^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}$$

$$\sigma_*^2 := \text{var}[\mathbf{y}_*] = \mathbf{k}_* - \mathbf{k}^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k} + \sigma_n^2$$

Here, $\mathbf{y} = [y_1, y_2, \dots]^T$, and \mathbf{K} is a matrix which contains the kernel function evaluation on each pair of training inputs, i.e., $\mathbf{K}[i, j] = k(\mathbf{x}_i, \mathbf{x}_j)$, \mathbf{k} is a vector of the kernel computation between each training input and the test input, i.e., $\mathbf{k}[i] = k(\mathbf{x}_*, \mathbf{x}_i)$, $\mathbf{k}_* = k(\mathbf{x}_*, \mathbf{x}_*)$, and \mathbf{I} is an identity matrix.

Choice of kernel function

The role of the kernel function is to specify how the function values, $f(\mathbf{x})$ and $f(\mathbf{x}')$, vary as the function of their corresponding inputs, \mathbf{x} and \mathbf{x}' . We use the following kernel function:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right) \exp\left(-\frac{\|x_s - x'_s\|^2}{2l_s^2}\right)$$

where x_s and x'_s are the spatial coordinates (latitude, longitude) of the *commune* centers corresponding to \mathbf{x} and \mathbf{x}' , respectively. The first exponent term captures non-linear dependencies in the feature space. The second exponent term plays the same role, but in the geographic space, and it models the spatial autocorrelation as a continuous function, which is the same as Kriging, a widely used methods in geostatistics (Cressie, 1990). The parameter σ_f^2 is the variance of the stochastic process f , l is the process length-scale for the feature space part, and l_s is the process length-scale for the spatial part.

The quantities $\beta, l, l_s, \sigma_n^2$, and σ_f^2 are estimated by maximizing the marginalized log-likelihood of the training data.



Estimating and Forecasting Income Poverty and Inequality in Haiti

using Satellite Imagery
and Mobile Phone data

June 2020