

# Mapping on a Budget: Optimizing Spatial Data Collection for ML

Livia Betti<sup>1</sup>, Farooq Sanni<sup>2,3</sup>, Gnouyaro Z. Sogoyou<sup>2,3</sup>, Togbe Agbagla<sup>2</sup>, Cullen Molitor<sup>3,4</sup>,  
Tamma Carleton<sup>3,4</sup>, Esther Rolf<sup>1</sup>

<sup>1</sup>University of Colorado Boulder

<sup>2</sup>Togo Data Lab

<sup>3</sup>Center for Effective Global Action

<sup>4</sup>University of California, Berkeley

livia.betti@colorado.edu

## Abstract

In applications across agriculture, ecology, and human development, machine learning with satellite imagery (SatML) is limited by the sparsity of labeled training data. While satellite data cover the globe, labeled training datasets for SatML are often small, spatially clustered, and collected for other purposes (e.g., administrative surveys or field measurements). Despite the pervasiveness of this issue in practice, past SatML research has largely focused on new model architectures and training algorithms to handle scarce training data, rather than modeling data conditions directly. This leaves scientists and policymakers who wish to use SatML for large-scale monitoring uncertain about whether and how to collect additional data to maximize performance. Here, we present the first problem formulation for the optimization of spatial training data in the presence of heterogeneous data collection costs and realistic budget constraints, as well as novel methods for addressing this problem. In experiments simulating different problem settings across three continents and four tasks, our strategies reveal substantial gains from sample optimization. Further experiments delineate settings for which optimized sampling is particularly effective. The problem formulation and methods we introduce are designed to generalize across application domains for SatML; we put special emphasis on a specific problem setting where our coauthors can immediately use our findings to augment clustered agricultural surveys for SatML monitoring in Togo.

## Introduction

Machine learning with satellite imagery (SatML) is increasingly used to fill crucial gaps in the coverage of existing ground-based data sampling efforts, resulting in global- and national-scale maps of poverty (Jean et al. 2016), deforestation (Hansen et al. 2013), cropland (Potapov et al. 2022), and many other outcomes. These maps unlock policy solutions on a wide range of urgent issues, informing design of conservation programs (Heilmayr et al. 2020), targeting of aid to those most in need (Smythe and Blumenstock 2022), and monitoring illegal activity (McDonald et al. 2021). And yet, SatML model performance is often modest, and improvements in performance are critical to avoiding policy design failures when using predictive models in practice.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: **Quality and cost of a training dataset depend on its spatial distribution.** This illustrative figure shows the performance ( $R^2$ ) of a model trained on datasets of sample size of 300 collected according to different sampling methods, for the task of predicting population density in the US.

To date, progress in SatML has been fueled by significant methodological advances in ML model architectures and satellite sensor technologies. In contrast, very little research attention has been paid to methods for training data collection, despite growing consensus that the quantity and spatial distribution of ground-referenced training labels is a consistent barrier to performance of these models (Burke et al. 2021; Rolf et al. 2024). For example, spatially clustered training data limit the generalizability of SatML maps over regions with large data gaps (Christensen et al. 2025). As governments and firms increasingly deploy SatML to monitor outcomes critical to human and environmental well-being, no evidence-based guidance exists on how to design sampling methodologies for the datasets used to train these models. Even more importantly, no prior work evaluates how the structure of those datasets interacts with the binding budget constraints that predominate in the data-poor environments where SatML is most impactful. Thus, while the quality and cost of training datasets depend critically on their spatial distribution (Figure 1), policymakers and researchers have no guidance on how to optimally conduct data collection for SatML applications.

To address this gap, **we formalize a novel problem statement for optimizing spatial collection of training data**, which captures two key features of SatML settings:

(1) Available labeled data are not intended for down-

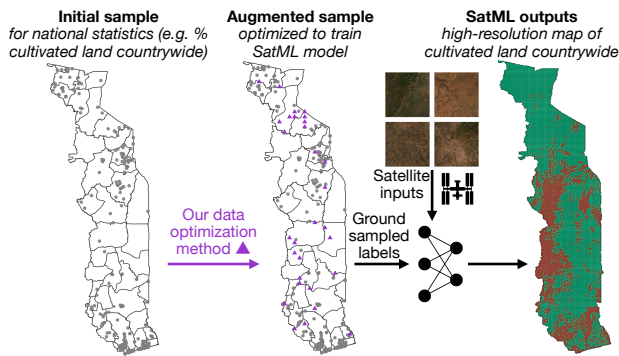


Figure 2: **Motivating problem setting in Togo:** Togo’s Ministry of Agriculture (MinAg) utilizes statistics on pest damage, crop yield, and cropland encroachment into protected areas to improve livelihoods for the 70% of the country’s population employed in farming. Currently, MinAg monitors this important economic sector with annual non-random agricultural surveys covering just 2,000 households due to the high costs of ground data collection. MinAg now aims to use these data to train a SatML model to map agricultural variables at high-resolution nationally. However, these agricultural surveys are not designed for SatML, limiting model performance. This setting reflects similar aims to fill data gaps with SatML in many low-income regions of the globe (Burke et al. 2021). In Togo’s case, the MinAg’s next survey will include a modest data augmentation budget aimed to optimize SatML performance. *Which extra households should be sampled to maximize performance of the SatML model?* This question is particularly challenging as the cost of additional samples varies across geography, with high costs of accessing remote regions with limited public infrastructure.

*stream use in SatML training, hence they exhibit strong spatial biases or gaps.* Crowd-sourced or web-scraped data represent convenience samples, as observations tend to be more concentrated in populated or easily accessible areas. Even demographic and economic surveys carefully designed to be nationally or subnationally representative (Corsi et al. 2012) often employ *cluster sample* methodologies that result in sparse pockets of labeled data (Figure 2).

(2) *Budget constraints and travel costs limit new data collection.* In ground-referenced data collection, travel costs can be a significant factor, and some regions may be infeasible to sample altogether due to terrain or safety concerns.

Using this formulation, we **develop the first methods to optimize spatial training datasets for SatML prediction**, accounting for heterogeneous physical data collection costs and binding budget constraints. We are motivated in large part by a specific problem in Togo, where our co-authors need to know how to spend a modest budget to augment a planned survey for training a model to generalize across the entire country (Figure 2). In extensive simulation experiments across sampling settings in three continents, we find that our method outputs training datasets that can out-perform common sense baselines and status-quo sam-

pling. This leads to either enhanced performance under a fixed sampling budget, or allows the user to keep performance standards high under increasingly stringent budget constraints.

With our problem formulation and novel method in hand, we can answer fundamental research questions on the nature of this problem, geared toward pressing questions that practitioners are already facing. Specifically, we design experiments to study:

- What types of representation are important to training generalizable models with spatial training data, and how do we optimize them in the face of budget constraints?
- In what real-world settings will additional data collection be most effective for increasing model performance?

We find that optimizing dataset design can have a significant impact on model performance, even when data collection budgets are tight. We hope this work provides practical tools for practitioners using SatML to fill critical data gaps and lays a foundation for future methodological work bridging the theory and practice of spatial dataset design in ML.

## Background & Related Work

SatML models are increasingly used for high-stakes applications such as agricultural monitoring, disaster relief, and poverty mapping, yet rely on training labels that are unevenly distributed across space (Aiken et al. 2022; Soman et al. 2023; Rolf et al. 2024). Lack of representative, high-quality labels has been a consistent bottleneck in performance (Cai et al. 2022): in a review of the use of satellite imagery and ML to fill gaps in economic surveys and agricultural and population censuses across the world, Burke et al. (2021) emphasize that the “largest constraint to model development is now training data rather than imagery.”

Approaches for collecting high-quality training data in SatML have cast this as an active learning problem, in which training datasets are refined through multiple rounds of label collection (Tuia et al. 2009; Desai and Ghose 2022; Kellenberger et al. 2019; Lin et al. 2020; Rodríguez et al. 2021; Stumpf et al. 2013; Ghozatlou et al. 2024). In a highly related study (but not using satellite imagery), Soman et al. (2023) evaluate active learning methods for the task of poverty estimation from call data records and phone surveys in Togo. Active learning is well-suited for remote satellite image annotation or phone surveys, where labeling is fast, iterative, and cost-independent across images. In contrast, ground-referenced data collection requires careful pre-planning and is usually executed in a single round, making active learning methods less practical.

Our problem formulation for one-round dataset augmentation is closer to traditional optimal experiment design formulations (Pukelsheim 2006), where data points (experiments) are chosen to satisfy desired optimality conditions. Related data valuation methods seek to measure the marginal utility of data points (Ghorbani and Zou 2019; Wang and Jia 2023), although these approaches typically focus on retroactively quantifying utility. The approaches we develop are inspired by recent work showing that dataset

composition is a consistent determinant of ML model performance (Rolf et al. 2021b; Hashimoto 2021). These works indicate that obtaining training data from different demographic groups within a population or sources of training data can lead to large differences in accuracy.

Here, we adapt these insights to spatial data, where groups can be defined in multiple ways – and where costs of labeling data can vary by group and location – to generate actionable conclusions regarding optimal sample design in SatML. In the spatial setting, Gaussian processes and Universal Kriging have been used for active sampling by quantifying uncertainty over space (Paley and Wolek 2020; Müller 2007). To our knowledge, this line of work has not addressed data collection for SatML specifically, and does not model the variable cost structure of physical data collection.

## Optimizing Spatial Data Composition under Physical Travel Constraints

Our goal is to maximize map accuracy of SatML models over specified geographic extents by selecting high-quality training data under physical sampling constraints. As there are many established SatML methods that can be used to associate collected labels with satellite imagery, we focus our optimization on the spatial distribution of training data (Figure 2, right). Motivated by real-world settings, we model this as a *data augmentation* problem: starting with some set of locations where labels are already sampled or will be sampled (e.g., a cluster sample from an annual survey), we aim to collect additional data to optimally extend that sample, under tight budget constraints. To keep our problem formulation as general as possible, we assume that the augmented sample must be designed only with knowledge of the spatial distribution of the original sample – this also reflects our motivating problem setting in Figure 2, where the entire survey must be designed at once, before any labels are collected.

### Problem Formulation

The SatML models we use are trained on image feature inputs  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$  and labels  $Y = \{y_1, \dots, y_n\}$ , then used to make predictions across the target region, composed of prediction units  $X_{\text{tgt}}$ , which is typically a grid of satellite tiles or small administrative units covering an entire country. We differentiate between prediction units  $X$  (satellite tiles), and *sampling units*, denoted by  $S = \{s_1, \dots, s_m\}$ , which represent the granularity at which data collection decisions are made. Sampling units can be any geographic units, for example census tracts, enumeration areas, or spatial clusters of households. We allow partial sampling within units (e.g., some households within a village), as is common for surveys defined over clusters or administrative areas. We can assume that each  $\mathbf{x}_i \in X$  corresponds to only one  $s$  and that the spatial correspondence between any  $s$  and its  $\mathbf{x}_i$  is known.

We denote the set of sampling units where data is already available or planned to be collected as  $S_0$ . We will augment this sample by selecting an additional set  $S_L$  to be labeled, drawing from the “source” set of possible sampling units  $S_{\text{src}}$ . The total labeled training set is then all  $\{\mathbf{x}_i, y_i\}$  pairs that correspond spatially with  $S_0 \cup S_L$ .

Our goal is to pick  $S_L$  to minimize the expected prediction loss of the trained SatML model averaged across the target region  $X_{\text{tgt}}$ , under a cost budget  $B$ . The objective is:

$$\begin{aligned} \arg \min_{S_L \subset S_{\text{src}} \setminus S_0} \mathbb{E} \left[ \sum_{\mathbf{x}_i \in X_{\text{tgt}}} \ell(\hat{f}(\mathbf{x}_i), y_i) \right] \\ \text{s.t. } \hat{f} \sim M(S_0 \cup S_L) \quad \text{and} \quad c(S_0 \cup S_L) \leq B \end{aligned} \quad (1)$$

where  $\ell : Y \times Y \rightarrow \mathbb{R}^+$  is the loss function,  $\hat{f}$  is trained using  $M$  on  $\{(\mathbf{x}_i, y_i) | \mathbf{x}_i \in S_0 \cup S_L\}$ , and  $c : 2^S \rightarrow \mathbb{R}^+$  assigns a cost to each subset of spatial units. In practice, sampling costs will depend greatly on the specific problem setting, so we leave this function in a general form.

Equation (1) models the SatML predictor  $\hat{f} : X \rightarrow Y$  as an (unknown) stochastic function  $M$  of its training data. This follows past work (Rolf et al. 2021b), where the authors estimate the expected loss of a trained predictor as a parametric function of (non-spatial) properties of the training data. In practice, eq. (1) is impossible to solve as  $M$  is unknown. Thus, in order to optimize dataset design in practice, we also need to estimate the relationship between a training set and the expected loss of a model trained on that sample – in deviation from past work, we propose to use utility functions as a proxy estimate of model performance.

### Method: Optimizing Proxy Quality Measures

Our proposed method optimizes a proxy utility function  $U$  defined over sets of sampling units:

$$\arg \max_{S_L \subset S_{\text{src}} \setminus S_0} U(S_0 \cup S_L) \quad \text{subject to} \quad c(S_0 \cup S_L) \leq B. \quad (2)$$

A good utility function will serve as a proxy for training set quality, such that maximizers of Equation (2) also incur low values of the expected loss in Equation (1). Since we can design the utility function  $U$  we use in our method, we select functions that make Equation (2) tractable to optimize.

We propose and test several concave utility functions, with the goal of understanding which proxy measures of quality best approximate the solution to eq. (1) under heterogeneous budgets  $B$  and cost functions  $c(\cdot)$ . Motivated by past work in data selection (Rolf et al. 2021b), we propose utility functions that capture two main factors of dataset composition: size and representativeness. We characterize representativeness through groups determined by geography, image embeddings, and other auxiliary information. We represent each of these as a function of a sample inclusion vector  $\mathbf{s} \in \{0, 1\}^{|S_{\text{src}}|}$ , where  $\mathbf{s}_i$  denotes whether sampling instance  $s_i$  contributes to the labeled sample  $S_0 \cup S_L$ . When solving the optimization in practice, we allow continuous values  $\mathbf{s} \in [0, 1]^N$ , which we treat as sampling probabilities.

Our first utility function reflects the common understanding that data quantity is important for model performance, and therefore is designed to maximize **dataset size**, where  $U(S) = |S|$ . Optimizing for a dataset size objective results in greedy sampling of the least expensive spatial units first until the budget is reached. We note that optimizing this is expected to bias the dataset toward the low cost regions.

Prior work has demonstrated that training set quantity is not the only determinant of model performance. Inspired by the group-specific scaling laws from Rolf et al. (2021b), we design **group-based representation** utility functions based on the representation from different types of groups  $g$  (e.g., spatial groups and image embedding groups). For a sample  $S$ , let  $n_g$  be the number of training instances belonging to each group  $g$  in a finite set of groups  $\mathcal{G}$ . Let  $\gamma_g$  denote the proportion of group  $g$  in the population. The (weighted) group-based representativeness utility function encourages balanced group coverage and is defined as

$$U(S) = -\lambda \sum_{g \in \mathcal{G}} \gamma_g n_g^{-1/2} - (1 - \lambda) n^{-1/2}.$$

The parameter  $\lambda \in (0, 1)$  balances group-based representativeness and overall dataset size, and follows from the group risk decomposition in Rolf et al. (2021b). When  $\lambda = 0$ , this objective reduces to the same dataset size objective.

To apply the concept of group-based representation to our spatial data optimization settings, we propose two broad strategies for defining groups. The first use existing administrative boundaries (e.g. states) as the set of groups  $\mathcal{G}$ . The second uses the distribution of images or other geographic data (e.g. land cover maps) to determine groups. In contrast to the size-based utility function, group-based representation objectives additionally score candidate training sets based on how well they represented the region at large. In Figure 3 (see Experiments section), we find all of our proposed utility functions are relatively good measures for ranking training datasets by performance of the resulting models.

## Experimental Settings

We use three datasets, each from different continents, representing agricultural and socioeconomic prediction tasks emblematic of SatML applications. All datasets have good spatial coverage over a country, allowing us to simulate sampling training datasets according to different methods. In our experiments, we randomly allocate 80% of the available data as the source set from which we draw different training set samples, and use the remaining 20% of the data as a test set to evaluate model performance across the country.

**Togo Soil Fertility** This dataset is closest to our motivating example and consists of 23,683 instances covering a  $1\text{km}^2$  grid of Togo, excluding protected areas, forests, and hydrographic zones. Each point is associated with a Sentinel-2 satellite image and labeled with water pH from a ground survey. Note that the data structure differs from the motivating example in Figure 2, as here we require sufficiently dense labeled data to evaluate the performance of various sampling regimes (such evaluation is impossible with the standard, clustered, 2,000-household agricultural survey). The original dataset lacks temporal labels, so we trained our models using features derived from 10m resolution 10 band Sentinel-2 images across two temporal windows spanning six years. We used the best-performing features (using imagery from January to June 2021) for all our analyses.

Sampling Type	Size	Rep (ours)		
		Admin	Image	NLCD
Cluster	0.740	0.740	0.744	0.736
Convenience/Urban	0.689	0.518	0.685	0.701
Random	0.959	0.716	0.955	0.949
Overall	0.381	0.347	0.510	0.540

Table 1: **Utility functions can estimate the performance of training datasets before labels are collected.** Spearman rank correlation ( $\rho$ ) between utility metrics and  $R^2$  scores of models trained on subsets of the USA Population dataset.  $\rho$  values are provided both within sampling types, and aggregated across sampled sets of all types (“Overall”).




Figure 3: **Illustrative examples of  $R^2$  score vs. utility metric**, corresponding to dataset samples used in Table 1.

**India SECC Consumption** We use the per-capita consumption measurements from the 2012 Indian Socio Economic and Caste Census (SECC) to represent a developing country administrative data use case. The data were originally packaged as the Socioeconomic High-resolution Rural-Urban Geographic Dataset for India (SHRUG) v2 (an updated version of Asher et al. (2021)) at the village/town level. We use the dataset from Aiken, Rolf, and Blumenstock (2023), which aggregates small administrative units together, resulting in 63,356 distinct units for training and evaluation. These units are matched to pre-computed MOSAIKS features derived from 4m resolution Planet Labs, Inc. imagery (accessed at <https://www.mosaiks.org/>).

**USAVars Tree Cover, Population Density** To obtain high-quality, densely-sampled data from a more data-rich environment, we use the USAVars dataset from Rolf et al. (2021a). This dataset contains 97,876 labeled  $1\text{km}^2$  grid cells, distributed roughly uniformly at random across the contiguous US. Labels are matched with features derived from NAIP aerial imagery ( $1\text{km}^2$  crops of 1m resolution imagery resampled to 4m/pixel). We use two different variables from USAVars – percent tree cover and (log transformed) population density – both of which are available on the torchgeo Python package (Stewart et al. 2022).

**MOSAICS Image feature embeddings.** For all experiments, we use unsupervised random convolutional feature (RCF) embeddings with a trained ridge regression prediction head (MOSAICS; Rolf et al. (2021a)) as our SatML

method, fit using 5-fold cross-validation over the train set, with regularization parameter selected from 10 log-spaced values from  $10^{-5}$  to  $10^5$ . MOSAIKS is designed to be computationally accessible, and thus represents the type of SatML model likely to be used in developing countries and resource-constrained settings. MOSAIKS has been shown to have high predictive skill (Brown et al. 2025; Corley et al. 2024). This fits our use case of keeping the SatML model fixed while experimenting with different data conditions. Since MOSAIKS only requires retraining a linear model, we are able to experiment with many more data conditions and random seeds than with a fully trained neural network. For the India dataset, we use the same precomputed MOSAIKS features as Aiken, Rolf, and Blumenstock (2023); for the Togo and USAVars datasets, we generate new RCF features using the torchgeo (Stewart et al. 2022) RCF implementation with empirical patches of kernel size 4 and bias  $-1$ .

### Initial Samples and Augmentation Costs

Our experimental simulation protocol is designed to closely mimic existing survey protocols, particularly those conducted in developing countries, such as the Demographic and Health Surveys (DHS) Program. To draw the initial samples for each country, we sample clusters with probability proportional to their size from a subset of  $N$  strata across the country; within each cluster we label  $k$  points chosen uniformly at random. For clusters with fewer than  $k$  points, we sample all points. To augment the initial set, we sample at the cluster level and label  $k$  points uniformly at random. We fix the choice of  $N$  strata but resample the clusters from within those  $N$  strata each time, so our results reflect variability due to different possible initial samples after strata are determined. To reflect spatially heterogeneous costs, we assign cost  $c_1$  to clusters that fall within the  $N$  strata in the initial sample and  $c_2$  to clusters in other strata, where  $c_2 > c_1$  reflects higher costs of accessing unseen strata, due to required travel and training new enumerators.

Due to the differing sizes of each country and the different structures of our three datasets, we tailor the parameters of our simulations to be meaningful for each country. For the USA setting, the strata are the 48 contiguous U.S. states and clusters are counties within each state,  $N = 5$ , and  $k = 10$ . For the India setting, the strata are 28 states and clusters are districts within each state,  $N = 10$ , and  $k = 20$ . For the Togo setting, the strata are Togo’s 5 regions and the clusters are cantons within each region,  $N = 2$ , and  $k = 25$ .

### Group-based Representation

Across the three countries, we define two types of spatial groups: groups determined by administrative units, and image-based groups, determined by  $k$ -means clustering in high-dimensional image embedding space. We refer to corresponding utility functions as **Admin-Rep** and **Image-Rep**, respectively. In the USA and India settings, states serve as the administrative units in Admin-Rep, and images are clustered into 8 groups for Image-Rep. For Togo, we use regions as the administrative units and form 3 image groups. In the USA setting, we include an additional grouping using data from the National Land Cover Database (NLCD).

At a resolution of 30m, each pixel in the 1km<sup>2</sup> NAIP image is assigned a land cover class. The resulting land cover proportion vectors are clustered using  $k$ -means in embedding space, resulting in 8 NLCD groups; we refer to the corresponding utility function as **NLCD-Rep**. Both NLCD-Rep and Image-Rep construct groups that are clustered in data space, though not necessarily spatially contiguous.

We experiment with  $\lambda = 0.5$  and  $\lambda = 1$ , and report  $\lambda = 0.5$ . We solve this optimization using the MOSEK solver via cvxpy (Diamond and Boyd 2016; Agrawal et al. 2018).

## Experiments

### Estimating the Quality of Spatial Training Data

Before we optimize spatial data collection in the face of physical data collection costs, we first conduct an experiment to isolate the capabilities of our utility functions. As we want to use these functions for optimization objectives, we are specifically interested in how well each utility function can rank different spatial training datasets in terms of the performance of a resulting trained SatML model.

In Table 1 and Figure 3, we analyze the  $R^2$  performance of a SatML model trained on different datasets in relation to the utility of each dataset. The setup is the same as in Figure 1: we take sub-samples of the USAVars Population training set, at increments of 100 points from 100 to 1000 using three sampling methods: cluster sampling with 10 points per cluster, convenience sampling, in which samples are selected with probability proportional to distance from the top 20 urban areas in the USA, and random sampling over space.

All of our proposed utility functions demonstrate skill at ranking different training datasets in terms of the downstream performance, as indicated by positive Spearman  $\rho$  rank correlation coefficient values. Among these, the group-based utility functions Image-Rep and NLCD-Rep most closely align (in rank) to  $R^2$  performance. All utility functions capture meaningful rankings within each specific sample type (Random, Cluster, Convenience), as indicated by similar Spearman’s  $\rho$  values (with  $\rho > 0.5$ ). However, there is more variability in predictive value when such utility functions are used to rank all the training sets in aggregate, with Image and NLCD-based representation outperforming size and administrative boundary representation metrics.

### Evaluating Methods for Spatial Data Collection

Next, we assess progress toward our main goal – optimizing these utility functions to increase model performance while adhering to spatial cost constraints. To evaluate our framework in these augmentation experiments, we use the cluster sampling method discussed in the previous section to generate initial samples of a fixed size that yield a positive (sometimes small)  $R^2$  score. We construct these initial samples to be large enough to accurately compute standard metrics used in policy or reporting, though not necessarily sufficient for high model performance. We then simulate an additional round of data collection using small budgets.

Table 2 shows the resulting model performance when using Admin-Rep, Image-Rep, and NLCD-Rep, compared against 3 baselines: default cluster sampling, which samples

Country & Task (Initial sample $R^2$ )	method used to augment initial sample						
	Budget	Default	Greedy (Size)	Random	Rep. (ours) admin	Rep. (ours) image	Rep. (ours) NLCD
<b>Togo Soil Fertility</b> (0.09 ± 0.12)	100	0.10 ± 0.13	0.08 ± 0.14	0.13 ± 0.07	0.08 ± 0.15	<b>0.17 ± 0.04</b>	n/a
	200	0.15 ± 0.05	0.06 ± 0.17	0.12 ± 0.11	0.15 ± 0.07	<b>0.18 ± 0.04</b>	n/a
	500	0.15 ± 0.06	0.06 ± 0.11	0.17 ± 0.03	0.20 ± 0.04	<b>0.21 ± 0.02</b>	n/a
	1000	0.15 ± 0.07	0.13 ± 0.05	0.20 ± 0.05	<b>0.23 ± 0.02</b>	<b>0.23 ± 0.01</b>	n/a
<b>India Consumption</b> (0.13 ± 0.06)	500	0.14 ± 0.04	0.14 ± 0.06	<b>0.18 ± 0.09</b>	<b>0.18 ± 0.08</b>	0.17 ± 0.04	n/a
	1000	0.16 ± 0.05	0.18 ± 0.04	<b>0.26 ± 0.05</b>	0.19 ± 0.27	0.18 ± 0.06	n/a
	2000	0.18 ± 0.04	0.16 ± 0.10	0.29 ± 0.06	<b>0.33 ± 0.02</b>	0.28 ± 0.02	n/a
	5000	<i>infeasible</i>	0.31 ± 0.07	0.37 ± 0.02	0.38 ± 0.03	<b>0.39 ± 0.02</b>	n/a
<b>USA Population</b> (0.0 ± 0.26)	50	<b>0.25 ± 0.19</b>	-0.21 ± 0.60	0.20 ± 0.16	0.10 ± 0.34	0.10 ± 0.30	0.13 ± 0.26
	100	0.31 ± 0.11	0.25 ± 0.17	0.23 ± 0.12	0.25 ± 0.12	<b>0.35 ± 0.07</b>	0.28 ± 0.11
	500	0.45 ± 0.03	0.44 ± 0.04	0.45 ± 0.05	<b>0.47 ± 0.03</b>	0.45 ± 0.03	0.46 ± 0.01
	1000	<i>infeasible</i>	0.48 ± 0.04	0.50 ± 0.03	<b>0.52 ± 0.02</b>	0.49 ± 0.04	<b>0.52 ± 0.01</b>
<b>USA Treecover</b> (0.49 ± 0.13)	50	0.56 ± 0.10	0.62 ± 0.10	0.60 ± 0.10	0.55 ± 0.12	0.57 ± 0.06	<b>0.68 ± 0.05</b>
	100	0.63 ± 0.07	0.63 ± 0.07	0.65 ± 0.08	<b>0.68 ± 0.05</b>	0.66 ± 0.06	0.66 ± 0.09
	500	0.75 ± 0.04	0.75 ± 0.03	0.72 ± 0.05	0.77 ± 0.01	0.76 ± 0.02	<b>0.78 ± 0.02</b>
	1000	<i>infeasible</i>	0.79 ± 0.01	0.78 ± 0.04	<b>0.81 ± 0.01</b>	<b>0.81 ± 0.01</b>	<b>0.81 ± 0.01</b>

Table 2: Prediction performance ( $R^2$  scores  $\pm$  one standard deviation for 10 random seeds), of SatML models trained with augmented samples. Here, budget represents the additional budget due only to augmentation. Initial samples of size 500 are chosen for Togo, size 2000 are chosen for India consumption, and size 100 are chosen for USA tasks. Since NLCD data is only available in the U.S., for Rep-NLCD is not applicable (n/a) for Togo and India. The default clustering performance is excluded for large budgets in India and the USA as these budgets cannot be realized with the cluster sampling procedure.

clusters in the already-sampled strata with probability proportional to size, greedy clusters<sup>1</sup>, which maximizes the number of clusters collected, and random clusters, which randomly selects clusters available for labeling regardless of cost. We observe that our cost-aware optimized sampling methods present competitive strategies compared to baseline methods – achieving the best average performance in nearly all settings. Overall, the optimized sampling framework is able to consistently improve over the baseline methods, with all three of our representation-based utility functions performing well compared to baselines. Notably, Image-Rep consistently achieves the highest performance on the Togo Soil Fertility dataset, though the winning representation-based method is more mixed in the other three tasks.

While performance differences due to different sampling methods are often within standard deviations of each other, it is encouraging that our methods perform consistently well and outperform baseline methods in almost all datasets and budgets. In fact, the variability in performance indicated by the large standard deviations underscores the sensitivity of model performance to sampling methods and thus the importance of studying dataset design.

Interestingly, Random Clusters emerges as a competitive baseline in Table 2. Particularly, for the India Consumption and Togo soil fertility tasks, Random Clusters tends to outperform Greedy Clusters and default Clusters, suggesting

<sup>1</sup>We implement Greedy Clusters using the size-based utility function to be consistent with the other optimization methods, yet still reflect the common heuristic of favoring large training datasets.

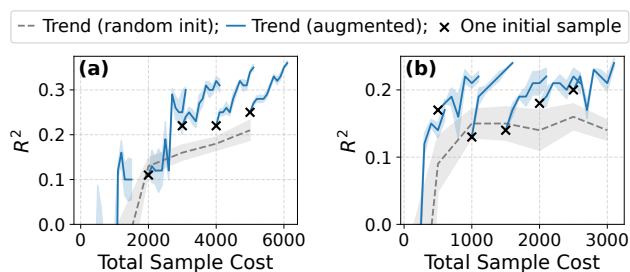
that there is fundamental value in using a sample that is well-dispersed over space, even at the cost of fewer samples.

### When is Optimized Data Collection Most Valuable?


Our next set of experiments investigate for what spatial problem settings is optimized sampling for augmentation worthwhile – expanding our analysis from the tables above to assess performance across varying cost structures (Figure 5), and different sizes of initial datasets (Figure 4). We focus this analysis on our two datasets reflecting more real-world scenarios, India consumption and Togo soil fertility.

In Figure 4, we compare optimized sampling to a status quo baseline, representing the possibility of governments or firms increasing sampling budgets, but allocating those budgets according to the same cluster strategy originally planned. Here, we model this as using larger budgets to allow more clusters to be sampled. We note that for large additional budgets, survey designers might add additional strata (increase  $N$ ), but in these experiments, we focus on low-budget augmentation. Across both the datasets in Figure 4, using our cost-aware augmentation methods (blue) results in performance gains over increasing the budget of default cluster sampling (grey) regardless of the initial sample size. Default cluster sampling is spatially constrained to pre-determined strata, limiting its maximum performance. The augmented samples in both datasets consistently achieve a higher performance ceiling, even at lower total cost.

In addition to varying the initial samples, we investigate



**Figure 4: Optimized sampling improves performance compared to default cluster sampling methods.** Performance of our best-performing cost-aware augmentation methods for India (Rep-Admin, panel a) and Togo (Rep-Image; panel b). Blue lines show augmentations to a single initial sample (black  $\times$ 's), and the first blue line, with no black  $\times$ , represents augmenting from an empty initial set. Gray dashed line and shaded areas show mean  $\pm 1$  standard error across 10 random initial samples. Augmentation methods are run for 5 random seeds.



**Figure 5: Performance gains due to optimized sampling depend on the cost structure.** Performance gain ( $\Delta R^2$  from initial sample) as we change the difference between in-region and out-of-region costs ( $c_2 - c_1$ ) for the soil fertility prediction task in Togo. We use an initial samples of size 500, in-region cost of  $c_1 = 25$ , and a budget of 500. We increase the out-of-region cost  $c_2$  between 25 and 50. Lines and shaded areas show average values  $\pm 1$  standard error.

the sensitivity of our methods to differences in spatial cost structures. In Figure 5, we vary the cost difference between in- and out-of-region cost for the Togo soil dataset. We find that performance gains ( $\Delta R^2$ ) decrease as the cost of difficult-to-sample regions increases. This makes sense as higher cost regions result in less data being collected for the same budget. Even so, Image-Rep (blue) consistently offers higher average performance gains irrespective of the cost difference, demonstrating that cost-aware optimized sampling is robust to varying cost discrepancies across samples.

## Discussion

The need for methods to aid in ground-based data collection for SatML is clear – to quote Burke et al. (2021), “expanding the quantity, and in particular, the quality of labels

will quickly accelerate progress in this field [of using satellite data for sustainable development] and will allow both researchers and practitioners to measure new outcomes and to accurately assess model performance.” In this work, we present the first approach to optimized sampling in realistic settings of preexisting clustered data and heterogeneous label collection costs. In experiments covering Togo, India, and the US, we found that our sampling optimization methods that encourage representation across spatial groups are almost always the top performing sampling method. The strong performance of our method, as well as the competitiveness of the baseline random clusters method, emphasizes the value of a representative or spatially dispersed sample, a paradigm that can be in conflict with the traditional goal in ML of collecting a large quantity of data.

By establishing the important influence of spatial training data distributions on SatML model performance, our study has important implications for future work, including the immediate use of our results in Togo (discussed in the following section) and promising areas for new research. Perhaps most prominently, there is the opportunity to improve on the optimization methods proposed here. Our results showed that the best way to define spatial groups may differ per country and prediction task; future work could work to provide a single best definition or delineate when different utility functions would be best. Studying data augmentation methods in settings where the initial sample has more extreme spatial biases (e.g. convenience sampling near population dense areas, which Figure 1 suggests is bad for SatML) would also be an interesting and valuable area of future work. Lastly, we hope to further investigate how our method can influence uncertainty quantification and generalizability of models.

**Takeaways for the Togo Problem Setting.** Our results have immediate implications for planned agricultural surveys in Togo, addressing our motivating problem setting outlined in Figure 2. Overall, **our findings imply an image-based representation sampling method is recommended**, as the samples generated according to this utility function consistently performed best in Togo (Table 2). While other methods are within the margin of error bounds of our best results in Table 2, we find that image-based sampling consistently performs well across different initial sampling budgets and data collection cost ratios (Figures 4 and 5), further increasing our confidence in recommending this method. Improving predictions with an optimally augmented sample (planned for 2026) will maximize the accuracy and thus social impact of the model trained on this dataset. Of course, while our results are promising, we recommend proceeding with care when putting this method into practice. First, we recommend visualizing the augmented samples and discussing their properties in consultation with the MinAg. Second, the annual survey collects data about different tasks than those used for our simulation study (due to data availability); thus, we recommend conducting preliminary analyses on related datasets if possible, to explore how the performance improvements recovered here are likely to translate.

## Acknowledgements

We would like to thank the Togolese Ministry of Agriculture, Village Hydraulics and Rural Development as well as the Togolese Institute of Agronomic Research, for providing us with the Togo Soil Fertility Data; Sean Luna McAdams for facilitating collaboration; and Emily Aiken for assisting with access to the India SECC Data. This work was funded by The McGovern Foundation and the Fund for Innovation in Development. In this work, we also used Jetstream2 at Indiana University through allocation CIS240692 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296. Finally, we would like to thank the reviewers for many helpful comments that contributed to the final version of this paper.

## References

- Agrawal, A.; Verschueren, R.; Diamond, S.; and Boyd, S. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1): 42–60.
- Aiken, E.; Bellue, S.; Karlan, D.; Udry, C.; and Blumenstock, J. E. 2022. Machine learning and phone data can improve targeting of humanitarian aid. *Nature*, 603(7903): 864–870.
- Aiken, E.; Rolf, E.; and Blumenstock, J. 2023. Fairness and representation in satellite-based poverty maps: evidence of urban-rural disparities and their impacts on downstream policy. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 5888–5896.
- Asher, S.; Lunt, T.; Matsuura, R.; and Novosad, P. 2021. Development research at high geographic resolution: an analysis of night-lights, firms, and poverty in India using the shrg open data platform. *The World Bank Economic Review*, 35(4): 845–871.
- Brown, C. F.; Kazmierski, M. R.; Pasquarella, V. J.; Rucklidge, W. J.; Samsikova, M.; Zhang, C.; Shelhamer, E.; Lahera, E.; Wiles, O.; Ilyushchenko, S.; Gorelick, N.; Zhang, L. L.; Alj, S.; Schechter, E.; Askay, S.; Guinan, O.; Moore, R.; Boukouvalas, A.; and Kohli, P. 2025. AlphaEarth Foundations: An embedding field model for accurate and efficient global mapping from sparse label data. arXiv:2507.22291.
- Burke, M.; Driscoll, A.; Lobell, D. B.; and Ermon, S. 2021. Using satellite imagery to understand and promote sustainable development. *Science*, 371(6535): eabe8628.
- Cai, W.; Encarnacion, R.; Chern, B.; Corbett-Davies, S.; Bogen, M.; Bergman, S.; and Goel, S. 2022. Adaptive sampling strategies to construct equitable training datasets. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1467–1478.
- Christensen, D.; Carleton, T.; Rolf, E.; Molitor, C.; Biswas, S.; Yan, K.; and Blair, G. 2025. Estimating the Footprint of Artisanal Mining in Africa. Technical report, National Bureau of Economic Research.
- Corley, I.; Robinson, C.; Dodhia, R.; Ferres, J. M. L.; and Najafirad, P. 2024. Revisiting Pre-trained Remote Sensing Model Benchmarks: Resizing and Normalization Matters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3162–3172.
- Corsi, D. J.; Neuman, M.; Finlay, J. E.; and Subramanian, S. 2012. Demographic and health surveys: a profile. *International journal of epidemiology*, 41(6): 1602–1613.
- Desai, S.; and Ghose, D. 2022. Active Learning for Improved Semi-Supervised Semantic Segmentation in Satellite Images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 553–563.
- Diamond, S.; and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83): 1–5.
- Ghorbani, A.; and Zou, J. 2019. Data Shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*, 2242–2251. PMLR.
- Ghozatlou, O.; Dacu, M.; Focsa, A.; Conde, M. H.; and Ullo, S. L. 2024. A review and a perspective of deep active learning for remote sensing image analysis: Enhanced adaptation to user conjecture. *IEEE Geoscience and Remote Sensing Magazine*, 12(3): 125–148.
- Hansen, M. C.; Potapov, P. V.; Moore, R.; Hancher, M.; Turubanova, S. A.; Tyukavina, A.; Thau, D.; Stehman, S. V.; Goetz, S. J.; Loveland, T. R.; et al. 2013. High-resolution global maps of 21st-century forest cover change. *science*, 342(6160): 850–853.
- Hashimoto, T. 2021. Model Performance Scaling with Multiple Data Sources. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 4107–4116. PMLR.
- Heilmayr, R.; Rausch, L. L.; Munger, J.; and Gibbs, H. K. 2020. Brazil’s Amazon soy moratorium reduced deforestation. *Nature Food*, 1(12): 801–810.
- Jean, N.; Burke, M.; Xie, M.; Davis, W. M.; Lobell, D. B.; and Ermon, S. 2016. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301): 790–794.
- Kellenberger, B.; Marcos, D.; Lobry, S.; and Tuia, D. 2019. Half a percent of labels is enough: Efficient animal detection in UAV imagery using deep CNNs and active learning. *IEEE Transactions on Geoscience and Remote Sensing*, 57(12): 9524–9533.
- Lin, Y.; Vosselman, G.; Cao, Y.; and Yang, M. Y. 2020. Active and incremental learning for semantic ALS point cloud segmentation. *ISPRS journal of photogrammetry and remote sensing*, 169: 73–92.
- McDonald, G. G.; Costello, C.; Bone, J.; Cabral, R. B.; Farabee, V.; Hochberg, T.; Kroodsmas, D.; Mangin, T.; Meng, K. C.; and Zahn, O. 2021. Satellites can reveal global extent of forced labor in the world’s fishing fleet. *Proceedings of the National Academy of Sciences*, 118(3): e2016238117.
- Müller, W. G. 2007. *Collecting spatial data: optimum design of experiments for random fields*. Springer.

Paley, D. A.; and Wolek, A. 2020. Mobile sensor networks and control: Adaptive sampling of spatiotemporal processes. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1): 91–114.

Potapov, P.; Turubanova, S.; Hansen, M. C.; Tyukavina, A.; Zalles, V.; Khan, A.; Song, X.-P.; Pickens, A.; Shen, Q.; and Cortez, J. 2022. Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. *Nature Food*, 3(1): 19–28.

Pukelsheim, F. 2006. *Optimal design of experiments*. SIAM.

Rodríguez, A. C.; D’Aronco, S.; Schindler, K.; and Wegner, J. D. 2021. Mapping oil palm density at country scale: An active learning approach. *Remote Sensing of Environment*, 261: 112479.

Rolf, E.; Klemmer, K.; Robinson, C.; and Kerner, H. 2024. Mission Critical—Satellite Data is a Distinct Modality in Machine Learning. In *International Conference of Machine Learning*.

Rolf, E.; Proctor, J.; Carleton, T.; Bolliger, I.; Shankar, V.; Ishihara, M.; Recht, B.; and Hsiang, S. 2021a. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature communications*, 12(1): 4392.

Rolf, E.; Worledge, T. T.; Recht, B.; and Jordan, M. 2021b. Representation matters: Assessing the importance of subgroup allocations in training data. In *International Conference on Machine Learning*, 9040–9051. PMLR.

Smythe, I. S.; and Blumenstock, J. E. 2022. Geographic microtargeting of social assistance with high-resolution poverty maps. *Proceedings of the National Academy of Sciences*, 119(32): e2120025119.

Soman, S.; Aiken, E.; Rolf, E.; and Blumenstock, J. 2023. Can Strategic Data Collection Improve the Performance of Poverty Prediction Models? International Conference on Learning Representations (ICLR 2023) Workshop on Practical Machine Learning for Developing Countries.

Stewart, A. J.; Robinson, C.; Corley, I. A.; Ortiz, A.; Ferrer, J. M. L.; and Banerjee, A. 2022. Torchgeo: deep learning with geospatial data. In *Proceedings of the 30th international conference on advances in geographic information systems*, 1–12.

Stumpf, A.; Lachiche, N.; Malet, J.-P.; Kerle, N.; and Puisant, A. 2013. Active learning in the spatial domain for remote sensing image classification. *IEEE transactions on geoscience and remote sensing*, 52(5): 2492–2507.

Tuia, D.; Ratle, F.; Pacifici, F.; Kanevski, M. F.; and Emery, W. J. 2009. Active Learning Methods for Remote Sensing Image Classification. *IEEE Transactions on Geoscience and Remote Sensing*, 47(7): 2218–2232. Conference Name: IEEE Transactions on Geoscience and Remote Sensing.

Wang, J. T.; and Jia, R. 2023. Data Banzhaf: A Robust Data Valuation Framework for Machine Learning. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, 6388–6421. PMLR. ISSN: 2640-3498.