

# SOMOSPIE: A modular SOil MOisture SPatial Inference Engine based on data-driven decisions

Danny Rorabaugh

*Electrical Engineering and Computer Science*  
*University of Tennessee, Knoxville*  
*dror@utk.edu*

Mario Guevara

*Plant and Soil Sciences*  
*University of Delaware, Newark*  
*mguevara@udel.edu*

Ricardo Llamas

*Plant and Soil Sciences*  
*University of Delaware, Newark*  
*rllamas@udel.edu*

Joy Kitson

*Computer and Information Sciences*  
*University of Delaware, Newark*  
*tkitson@udel.edu*

Rodrigo Vargas

*Plant and Soil Sciences*  
*University of Delaware, Newark*  
*rvargas@udel.edu*

Michela Taufer

*Electrical Engineering and Computer Science*  
*University of Tennessee, Knoxville*  
*taufer@utk.edu*

**Abstract**—The current availability of soil moisture data over large areas comes from satellite remote sensing technologies (i.e., radar-based systems), but these data have coarse resolution and often exhibit large spatial information gaps. Where data are too coarse or sparse for a given need (e.g., precision farming), one can leverage machine-learning techniques coupled with other sources of environmental information (e.g., topography) to generate gap-free information at a finer spatial resolution (i.e., increased granularity). To this end, we develop a spatial inference engine consisting of modular stages for processing spatial environmental data, generating predictions with machine-learning techniques, and analyzing these predictions. We demonstrate the functionality of this approach and the effects of data processing choices via multiple prediction maps over a United States ecological region with a highly diverse soil moisture profile (i.e., the Middle Atlantic Coastal Plains). The relevance of our work derives from a pressing need to improve the spatial representation of soil moisture for applications in environmental sciences (e.g., ecological niche modeling, carbon monitoring systems, and other Earth system models) and precision farming (e.g., optimizing irrigation practices and other land management decisions).

**Index Terms**—soil moisture, remote sensing, machine learning, data-driven decisions

## I. INTRODUCTION

Soil moisture is a critical variable that links climate dynamics with water and food security. It regulates land-atmosphere interactions (e.g., via evapotranspiration—the loss of water from evaporation and plant transpiration to the atmosphere), and it is directly linked with plant productivity and plant survival [1]. Information on soil moisture is important to design appropriate irrigation strategies to increase crop yield, and long-term soil moisture coupled with climate information provides insights into trends and potential agricultural thresholds and risks [2]–[4]. Thus, information on soil moisture is needed to assess the implications of environmental variability and is consequently a key factor to inform and enable precision farming. Currently, large areas of western states of

M. G. acknowledges support from a Conacyt fellowship; M. T. and R. V. acknowledge support from NSF (#1724843, #1854312); and R. V. acknowledges partial support from NSF (#1652594).

the conterminous United States (CONUS) are experiencing an exceptional drought, and most information on water limitation has been derived from changes in precipitation patterns [5]. This is just one example of where soil moisture information can be utilized in situations of critical importance.

The current availability in soil moisture data over large areas comes from remote sensing (i.e., satellites with radar sensors), which provides nearly global coverage of soil moisture at spatial resolution of tens of kilometers [6], [7]. Recent efforts have devoted to increase the spatial resolution of current estimates [8]. Other efforts have focused on harmonizing historical satellite soil moisture records for larger periods of time and from several information sources [9]. Satellite soil moisture data has two main shortcomings. First, although satellites can provide daily global information, they are limited to coarse spatial resolution (at the multi-kilometer scale). Second, satellites are unable to measure soil moisture in areas of dense vegetation, snow cover, or extremely dry surfaces; this results in gaps in the data. Fig. 1 shows an example of the monthly averages of daily soil moisture data for December 2000. The figure shows spatial information gaps across the globe due to (a) dense vegetation over the Amazon or snow and (b) ice cover across Greenland, for example. To use the spatial representation of soil moisture for applications in environmental sciences (e.g., ecological niche modeling, carbon monitoring systems, and other Earth system models) and precision farming (e.g., optimizing irrigation practices and other land management decisions), we need to increase the spatial resolution of information and predict values in areas with missing data.

In this paper, we address the two shortcomings associated with satellite data (i.e., coarse-grained resolution and spatial information gaps) by providing a modular SOil MOisture SPatial Inference Engine (SOMOSPIE). SOMOSPIE consists of modular components including input of available data at its native spatial resolution, selection of a geographic region of interest, prediction of missing values across the entire region of interest (i.e., gap-filling), analysis of generated predictions,




Fig. 1: Monthly soil moisture ( $\text{m}^3/\text{m}^3$ ) averages for December 2000 with gaps where data cannot be collected accurately because of dense vegetation, snow cover, and extremely dry surfaces. Averaged from daily data from the ESA-CCI soil moisture database [9].

and visualization of both predictions and analyses. To predict soil moisture, our engine leverages hydrologically meaningful terrain parameters (e.g., slope and topographic wetness index) calculated using an open source platforms for standard terrain analysis (i.e., SAGA-GIS or System for Automated GeoScientific Analysis-Geographical Information System) and various machine learning methods. The engine combines the publicly available datasets of satellite-derived soil moisture measurements from the European Space Agency (ESA) and generates fine-grained, gap-free soil moisture predictions using three implementations of machine learning algorithms: a kernel-based approach (kernel-weighted  $k$ -nearest neighbors or KKNN), the Hybrid Piecewise Polynomial approach (HYPPO), and a tree-based approach (Random Forests or RF). Data processing functionality of our engine includes selection of a region of interest, which we demonstrate using ecoregions as defined for North America by the Commission for Environmental Cooperation [10]. We exhibit the full functionality of our engine on the Middle Atlantic Coastal Plains in the eastern United States, a region with a diverse soil moisture profile.

The main contributions of this paper are:

- 1) A spatial inferences engine (SOMOSPIE) and all the data and components needed to generate viable soil moisture predictions;
- 2) An empirical study of the engine's functionality including an assessment of data processing and fine-grained predictions over a United States ecological region with a highly diverse soil moisture profile (i.e., the Middle Atlantic Coastal Plains).

The rest of the paper is structured as follows. Section II describes the datasets used for this project. Section III consists of a breakdown of the components of SOMOSPIE throughout its three stages: data processing, prediction generation, and

prediction analysis and visualization. Section IV contains results, including soil moisture predictions that leverage various modular elements of SOMOSPIE and analyses thereof. Section V describes the present research within the broader context of related work. Section VI closes the paper with our conclusion.

## II. OUR DATASETS

Our work builds upon publicly available data collections associated with remotely-sensed soil moisture information, topographic characteristics derived from quantitative land surface analysis, and eco-regionalization of North America. These diverse datasets are cornerstones of SOMOSPIE: (a) moisture records, (b) a digital surface model (DSM), and (c) boundaries for ecoregions. Table I reports the data resolutions and sources.

Satellite-derived soil moisture data are downloaded from the ESA-CCI soil moisture initiative [7], [11]–[13]. The ESA-CCI soil moisture data are collected in a raster format and have an original spatial resolution of  $0.25 \times 0.25$  lat-lon degrees (about  $27 \times 27$  km). Fig. 2 portrays a satellite collecting raster data. Each pixel in the raster file corresponds to a square of land and contains the satellite-derived soil moisture value for that land surface. The value is a ratio (between 0 and 1), the number of  $\text{m}^3$  of water per  $\text{m}^3$  of surface soil, where 0 indicates dry soil and 1 indicates water-saturated soil. This dataset is representative for the first 0 to 5 cm of soil surface [14]. The original temporal scale of the ESA-CCI is daily, but for this study we move from daily to monthly time steps by averaging all daily values in a given pixel across an entire month.

Our topography dataset, consisting of multiple terrain parameters, is based on DSMs, which are available at several resolutions and are useful to represent multiple terrain characteristics [15]. Topography is an important factor affecting

TABLE I: List of datasets used in this study.

Dataset	Spatial resolution	Temporal resolution	Variable / Description	Source
ESA-CCI	0.25 degrees	Daily, 1978–2016	soil moisture ( $\text{m}^3/\text{m}^3$ )	European Space Agency
DSM	$\approx 30$ meters	Static ('Current')	Land surface characteristics	The Japan Aerospace Exploration Agency
CEC	n/a	Static ('Current')	Ecoregion boundaries	Commission for Environmental Cooperation

water distribution in soils since it directly affects overland flow and solar radiation rates [16]–[18]. DSMs are the main inputs of geomorphometry, which is the science of quantitative land-surface analysis [19]. The influence of topography on soil moisture prompts its present inclusion in soil moisture downscaling.




Fig. 2: Illustration of how satellites collect raster data across the surface of the Earth [12], [20].

To define the spatial limits of our soil moisture prediction, we use the 2011 update of the Commission for Environmental Cooperation (CEC) ecoregion dataset, developed jointly by Mexico, the United States, and Canada and based on the analysis of ecosystem elements such as geology, physiography, vegetation, climate, soils, land use, wildlife, and hydrology [10]. This approach divides North America into polygon-based ecoregions at three levels, which range from Level I (Fig. 3) to Level III and describe the similarity in the type, quality, and quantity of environmental parameters within the region. Level I regions are larger and more general and Level III regions are smaller and more specific.

### III. SOMOSPIE OVERVIEW

We present a modular SOil MOisture SPatial Inference Engine (SOMOSPIE) for prediction of missing soil moisture

information. SOMOSPIE includes three main stages, illustrated in Fig. 4: (1) data processing to select a region of interest, incorporate predictive factors such as topographic parameters, and reduce data redundancy for these new factors; (2) soil moisture prediction with three different machine learning methods (i.e., KKNN, RF, and HYPO); and (3) analysis and visualization of the prediction outputs.

#### A. Data and Data Preprocessing

With SOMOSPIE, data are separated into two independent groups (i.e., observed data and evaluation data) to be fed into one of our modeling approaches defined in Section III-B. Observed data are represented as vectors, one for each pixel in the satellite data. Each vector consists of the latitude and longitude of the centroid of the pixel in the satellite data, an average soil moisture ratio for that pixel, and (optionally) the values of 15 topographic parameters from the digital surface model (DSM) evaluated at that centroid. Additionally, a user can specify a percentage of the observed data to be randomly set aside as validation data. We discuss validation further in Section III-C. The remainder of the observed data not used for validation becomes our training data for generating models. Evaluation data are represented as vectors in the same format as observed data, except that the evaluation vectors do not have a soil moisture ratio, as this value is unknown for the corresponding pixel where a prediction is to be made. In this study, our desired resolution for soil moisture prediction is the  $1 \times 1$  km resolution of the DSM we are using.

Both climate and topography influence the spatial patterns of soil moisture [21]. Our approach selects a specific ecoregion, with similar climatic characteristics, and predicts the soil moisture profile at a fine-grain resolution for a region of interest within that ecoregion. Our engine uses the Commission for Environmental Cooperation (CEC) ecoregions as masks to select ecoregion. Consequently, topography becomes




Fig. 3: CEC Level I ecoregions across the conterminous United States.




Fig. 4: Overview of our modular SOil MOisture SPatial Inference Engine (SOMOSPIE) with its three stages: data preparation (left), prediction with machine learning methods (middle), and analysis and visualization of predictions (right).

the driving factor in our predictions of soil moisture spatial patterns. Specifically, our models use topography as the covariate space to downscale and gap-fill satellite-derived soil moisture, leveraging attributes derived from the DSM, such as the terrain slope or the aspect (i.e., the first and second derivatives of elevation data). These terrain parameters are surrogates of two main processes controlled by topography: the overland flow of water and the potential incoming solar radiation. Our topographic attributes are calculated using the SAGA GIS basic terrain parameters module<sup>1</sup> [17], [22].

Because terrain attributes can have significant correlation, SOMOSPIE allows users to apply Principal Component Analysis (PCA) to reduce the number of covariates. This transformation is relevant because the reduction in DSM topographic parameters can reduce the time needed for prediction. To perform this reduction on the data for the region of interest, the engine uses the PCA implementation from the Python package *sklearn*<sup>2</sup>, and selects components whose corresponding eigenvalues are at least one (a common convention [23]). Dimensional reduction of DSM data, when applied, is performed identically on both the training and evaluation data.

#### B. Prediction Models

SOMOSPIE presently supports three key machine learning modules for predictions of missing spatial soil moisture information: kernel-weighted  $k$ -nearest neighbor (KKNN), Random Forests (RF), and Hybrid Piecewise Polynomial (HYPPO). The selected methods have distinct character, with different assumptions about the model being generated. Yet they all share built-in automated parameter tuning. KKNN tunes for the kernel and number of neighbors, RF tunes for number of

variables per tree level, and HYPPO tunes for local polynomial degrees, all of which are explained below in the descriptions of the respective models. Our implementations of all three methods use 10-fold cross validation to accomplish the tuning. This is a standard technique [24] which involves dividing data into 10 roughly equal parts. For every possible parameter value, 10 different models are generated; each model builds from nine tenths of the data and is tested on the other tenth. The parameter value that minimizes cumulative error across all models is selected for generating a model with all the data. For our implementations of KKNN and RF, this tuning via cross-validation is performed with the R package *caret*. Having described the commonality of the methods, we now discuss the specific structure of each method in more detail.

The traditional  $k$ -nearest neighbor ( $k$ NN) regression technique builds many simple models from local data. Use of this technique assumes that the  $k$  points nearest in the prediction space to the data point one wishes to model are the most relevant and that points farther away have less influence on the point in question. The process begins with training data: points  $\langle x_1^1, \dots, x_1^d, z_1 \rangle, \dots, \langle x_n^1, \dots, x_n^d, z_n \rangle$ , where  $(x_i^1, \dots, x_i^d)$  are coordinates in the  $d$ -dimensional prediction space and  $z_i$  is the corresponding soil moisture ratio. To predict soil moisture ratios for a specific choice of  $(x^1, \dots, x^d)$ ,  $k$ NN selects the  $k$  nearest neighbors of  $(x^1, \dots, x^d)$  in the training data and uses the arithmetic mean of their associated soil moisture ratios.

A common generalization of  $k$ NN is to use a weighted mean of the  $k$  nearest soil moisture ratios, where values from points nearer to  $(x^1, \dots, x^d)$  are given higher weights. The variant of  $k$ NN in our engine is kernel-weighted  $k$ -nearest neighbors (KKNN) [25], implemented with the R package *kknn*. It uses a kernel function (i.e., rectangular, triangular, Epanechnikov, Gaussian, rank, or optimal) to compute neighbor weights

<sup>1</sup>[http://saga-gis.org/saga\\_tool\\_doc/2.1.3/ta\\_compound\\_0.html](http://saga-gis.org/saga_tool_doc/2.1.3/ta_compound_0.html)

<sup>2</sup>[sklearn.decomposition.PCA](http://scikit-learn.org/stable/modules/decomposition.html)

for the mean. Cross-validation is employed to determine the number of neighbors and which weighting kernel to use.

Random Forests (RF) consist of an ensemble of decision trees that are weighted via a statistical method called bagging. Each tree is grown with a random subset of predictors and of the training data. The tree's weight is determined by its "out-of-bag error", which is computed by testing the tree on the rest of the training data. To make a prediction at a new point, all decision trees in the ensemble are queried and their predictions are combined using weighted arithmetic mean. Such techniques do not assume a particular functional or geometric form of the model. Thus the techniques are suitable to deal with sparse datasets (e.g., areas with large gaps of soil moisture satellite estimates). SOMOSPIE employs the R package *quantregForest* and has two main parameters: (1) the number of trees to grow in the ensemble of regression trees; and (2) the number of covariates randomly selected at each level of tree growth. In the present study, we consider a total of 500 trees for the first parameter. The second parameter is bounded above by the total number of prediction parameters (i.e., 17 in our case with 2 spatial coordinates and 15 topographic predictors) and automatically selected using cross validation as described above.

The Hybrid Piecewise Polynomial (HYPPO) module builds upon and extends traditional  $k$ NN in a different way to mitigating some of its limitations [26]. Contrary to  $k$ NN, HYPPO allows local prediction models to be non-linear. In other words, the polynomial degrees in HYPPO become a flexible feature of the model. To build the prediction model with HYPPO, we start with  $n$  training points,  $\langle x_1, y_1, z_1 \rangle, \dots, \langle x_n, y_n, z_n \rangle$ . We want to predict the value of  $z$  (i.e., the soil moisture) for a new, specified coordinate  $(x, y)$ . Following the technique of  $k$ NN, we first find the  $k$  nearest neighbors of  $(x, y)$ , then using the data of the  $k$  nearest neighbors, HYPPO builds local models using a polynomial whose degree is selected using cross validation as described above.

The standard generation of a non-linear polynomial on a large set of variables requires a large number of data points. With 15 topographic parameters, the current implementation of HYPPO would require a larger number of neighbors than those always available in the regions of interest used in this paper. Thus, the present study demonstrates HYPPO using only latitude and longitude as parameters. Work in progress tackles the challenge of tuning HYPPO to efficiently work with a larger number of predictors.

### C. Analysis and Visualization

Our engine supports analysis and visualization to assess model output. As stated in Section III-A, a user can specify a percentage of the observed data to split off as validation data. All of the results in this paper are generated with 20% of the observed data points randomly set aside as validation data. After generating predictions, the engine compares the predicted soil moisture values to those of the validation points. In this study, we use the squared correlation coefficient ( $R^2$ -value) as explained variance between the validation data at the

coarse spatial resolution of  $27 \times 27$  km (i.e., satellite-derived soil moisture) and the prediction of soil moisture at  $1 \times 1$  km resolution. To this end, we first harmonize the predicted data with the validation data by computing the arithmetic mean of all the predicted values for  $1 \times 1$  km pixels that fall within a cell of the original grid ( $27 \times 27$  km). To account for the randomness of the subset selected for validation, every execution of the workflow for this study is repeated 10 times. For each execution, the resulting  $R^2$ -value is stored (rounded to two digits after the decimal); we give the arithmetic mean of the ten stored values. The standard deviation for the set of ten  $R^2$ -values fell between 0.088 and 0.199 in every case.

The engine also provides a suite of visualization tools for soil moisture predictions and evaluations. It uses R and Python scripts to perform standard geographic information system (GIS) tasks for both imagery and tabular forms of data. The current study demonstrates the creation of heatmaps for soil moisture predictions and scatter plots for comparisons of predictions to validation data. In both cases, rather than show all ten plots generated for every usage of the engine, we select an instance whose corresponding  $R^2$ -value is closest to the mean value for that usage. For the selected instance, we show the prediction heatmap as representative of the set of ten predictions.




Fig. 5: Selected Level III ecoregion for this study: the Middle Atlantic Coastal Plains (8.5.1).

## IV. PREDICTION RESULTS

For this study we select a small Commission for Environmental Cooperation (CEC) ecoregion: the Middle Atlantic Coastal Plains (Level III ecoregion 8.5.1, see Fig. 5). This region has a broad range of moisture ratios with which to test the capabilities of SOMOSPIE. The soil moisture ratios we use for the observed data of all of these demonstrations are from April 2016. For each latitude and longitude coordinate pair in the satellite data, we take the average of all soil moisture ratios available in that pixel that month.

SOMOSPIE can increase the granularity of soil moisture information from coarse-grained satellite data using machine learning techniques. Presently we downscale from the original satellite-derived soil moisture native resolution ( $27 \times 27$  km) to the  $1 \times 1$  km resolution of our topographic predictors. In Section IV-A, we demonstrate this with three machine learning methods: KKNN, RF, and HYPPO. In Section IV-B, we take a

step back in the SOMOSPIE workflow and examine the effects of data decisions on prediction. In particular, we generate predictions for the same ecoregion (8.5.1) using KKNN and RF, but with training data from a larger region and with PCA reduction applied to the 15 topographic dimensions of the training data.

#### A. Soil Moisture Predictions

We present soil moisture prediction on our local Level III ecoregion (i.e., Middle Atlantic Coastal Plains) as heatmaps with soil moisture ratios between 0.19 and 0.31. Fig. 6 shows predictions generated by our engine with three supported machine learning methods. Note the warmer colors in the southern portion (in South Carolina) representing lower ratios of moisture to land on the surface.

Qualitatively, the patterns of soil moisture trends show agreement across model predictions. We observe a more noisy prediction from KKNN, but unrealistic spatial artifacts (e.g., sharp jumps in North and South Carolina) from RF and, to a lesser extent, HYPPO. We also observe the KKNN predictions having fewer extreme values (blue for wetter and dark orange for drier). Despite the spatial artifacts, RF and HYPPO do not share this limitation in the range of soil moisture ratios predicted for such a region as the Middle Atlantic Coastal Plains. This is consistent with the claim in [26] that HYPPO was “designed to effectively and accurately model non-smooth [...] surfaces without the need for extensive sampling” since “most traditional techniques are designed to produce smooth models.”

Quantitatively, we investigate the relationship between the predicted soil moisture data and the initial coarse-grained observations (Fig. 7). We use explained variance ( $R^2$ -value) as an accuracy measure of the overall modeling performance, calculated from the relationship of the validation subset of the original satellite-based soil moisture data and the predicted

soil moisture estimates. The KKNN predictions show the lowest explained variance (a mean  $R^2$ -value of 0.296) between observed and predicted satellite soil moisture ratios; RF and HYPPO models have similar accuracy (with mean  $R^2$ -values of 0.575 and 0.557, respectively).

#### B. Impact of Data Processing Decisions

One of the features in our engine when selecting a region of interest is the use of a larger region for model generation beyond the boundary of an ecoregion of interest. The assumption is that between neighboring ecoregions there is not necessarily a sharp separation but rather some sort of transition with multiple ecological gradients that serves as a buffer. Buffer selection can be useful for quantifying spatial gradients of ecosystem functional diversity and soil moisture feedbacks at the borders of ecological regions. Thus, the engine facilitates predictions using a larger region in two ways: (1) users can specify a buffer distance to be automatically added around the region of interest; or (2) users can opt for prediction to use the lower level (i.e., larger, less ecologically specific) ecoregion containing the region of interest. To investigate the effect of regional restrictions, we run the same prediction methods (KKNN and RF) on three larger regions containing ecoregion 8.5.1, the Middle Atlantic Coastal Plains. Two of the three are ecoregion 8.5.1 with a fixed-width buffer, one with a 50 km buffer and one with a 100 km buffer (Fig. 8a). The third enlarged region is Level II ecoregion 8.5, the Mississippi Alluvial and Southwest USA Coastal Plains (Fig. 8b).

The base results for comparison are from the predictions already generated for the region of interest, ecoregion 8.5.1, using KKNN (Fig. 9a) and RF (Fig. 10a). We generate models on the three larger regions, then use the models to obtain soil moisture ratios only for our region of interest. With KKNN, when implementing a 50 km or 100 km buffer around the Level III ecoregion (Fig. 9b and 9c), we observe slightly higher




Fig. 6: Example of prediction maps (1 × 1 km resolution) for ecoregion 8.5.1 (i.e., the Middle Atlantic Coastal Plains) generated by three different ML algorithms.




Fig. 7: Scatter plots comparing the modeled soil moisture from Fig. 6 and the validation data from the original satellite-based product ( $27 \times 27$  km resolution).




Fig. 8: Regions containing ecoregion 8.5.1, the Middle Atlantic Coastal Plains: Level III ecoregion 8.5.1 with a 100 km buffer (a, left); and Level II ecoregion 8.5 (b, right).

(darker green and blue-green) predictions in the northern third of the region (i.e., in Virginia, Maryland, and Delaware). Despite this shift, the spatial patterns are visually preserved across the KKNN predictions. With RF, when we add a 50 km or 100 km buffer (Fig. 10b and 10c), we see the spatial artifacts (e.g., sharp jumps in soil moisture ratios) become more pronounced. When using training data from the entire Level II ecoregion, KKNN appears to predict a narrower range of values in southern part of the region (Fig. 9d) and RF still exhibits sharp lines but with smaller value jumps across those lines (Fig. 10d).

For these models generated on various supersets of the region of interest, we now move from qualitative observations to quantitative analyses. We find that using training data further outside the region of interest sharply diminishes the explained variance between the KKNN predictions and the satellite observations. The original predictions had a mean  $R^2$ -value of 0.296, that dropped to 0.240 with a 50 km buffer, further to 0.204 with a 100 km buffer, and even further to 0.099 using the Level II ecoregion 8.5. For RF predictions, the mean  $R^2$ -value is generally unaffected, with 0.575 for the original

region, 0.577 with a 50 km buffer, 0.600 with a 100 km buffer, and 0.551 using the Level II ecoregion. This indicates that KKNN is more prone than RF to be negatively effected by extraneous data; yet our findings support the hypothesis that restriction to a region of common ecological character may be beneficial for some soil modeling efforts.

Another optional features in our engine is the use of Principal Component Analysis (PCA) to reduce the number of covariates from terrain attributes before running a machine learning algorithm. We perform ten rounds of predictions with KKNN and RF over the region of interest on 80% of the observed data with PCA reduction applied to the 15 topographic dimensions. In nine of the rounds, six of the principal components are used (having eigenvalues above the standard threshold of one); seven principal components are used in the other round. Fig. 11 demonstrates the consequence of using the PCA dimension reduction on soil moisture prediction for the Middle Atlantic Coastal Plains. Specifically, Fig. 11a and 11c show predictions when using all 15 DSM predictors while Fig. 11b and 11d show the predictions with the PCA-reduced training data. In both cases, the use of PCA reduction appears to cause sharper local contrasts.

To assess the accuracy of these predictions, we again compare our findings with the validation subset of the observed satellite soil moisture estimates. We find that the KKNN model using all the topographic prediction factors shows only a slightly higher explained variance with a mean  $R^2$ -value of 0.296 against the original product compared to the predictions based on the PCA transformation with a mean  $R^2$ -value of 0.289. On the other hand, we see a distinct rise in the mean explained variance for the RF predictions, which produces a mean  $R^2$ -value of 0.575 with all the topographic dimensions, yet a mean  $R^2$ -value of 0.693 after leveraging PCA. This indicates that a PCA dimension reduction can




Fig. 9: Prediction maps for the Middle Atlantic Coastal Plains (Level III ecoregion 8.5.1) using KKNN on training data from areas containing Level III ecoregion 8.5.1.




Fig. 10: Prediction maps for the Middle Atlantic Coastal Plains (Level III ecoregion 8.5.1) using RF on training data from areas containing Level III ecoregion 8.5.1.

be useful to smooth out the data for modeling methods that tend to overfit. Moreover, there is significant reduction to the number of prediction factors (from 15 to 6 or 7), for this input data. This indicates that PCA can be used to reduce the computational demand for generating predictions at a finer scale. Both the statistical and computational performance are important considerations in data-driven modeling frameworks in order to maximize accuracy of results while constraining computation time.

Overall we observe how SOMOSPIE is able to perform different methods for predicting the spatial pattern of soil moisture across ecoregions with varying soil moisture profiles. The use and assessment of various modeling approaches is helpful to represent the complex variability of soil moisture and its dynamics, as each method is able to capture distinct elements of the soil moisture variability across the region of interest. Moreover, the study on the impact of data processing demonstrates how initial data processing decisions such as region selection can impact the prediction output, even for a fixed selection of machine learning model. Such decisions cannot be arbitrary and must be motivated by established

scientific knowledge and by the accuracy of the predictions in comparison to existing data.

## V. RELATED WORK

Our work builds on recent technological advance in satellite-derived soil moisture (European Space Agency Climate Change Initiative [11]–[13]). Other satellite-derived datasets of soil moisture can be used in our workflow such as AMSR-E (Advanced Microwave Scanning Radiometer - Earth Observing System Sensor on the NASA Aqua Satellite [27]), ASCAT (Advanced SCATterometer aboard the EUMETSAT MetOp satellite [28]–[30]), and AQUARIUS (Satellite instrument from NASA SMAP mission [31]). Despite technological advances, satellite datasets still have coarse spatial resolution and present temporal gaps making support tools such as SOMOSPIE useful to provide insights for research, environmental management, and precision farming based on remote sensing data.

Our work complements recent efforts [32], [33] and provides the building blocks for interdisciplinary work and software development for soil moisture products. This project builds on the increasing recognition of the importance of




Fig. 11: Effect of modeling over ecoregion 8.5.1 with KKNN (top) or RF (bottom) using all 15 DSM predictors (left) or with auto-selected PCA components (right).

spatial and temporal dependency of environmental data [34]–[36] and its application to precision farming [37]. We take on the under-utilization of computer science techniques and computational resources to downscale satellite-derived soil moisture data, in order to describe trends in soil moisture across CONUS. This project focuses on how topography and environmental variability influence soil moisture [38] across CONUS, and we postulate that our cyberinfrastructure tool has worldwide applicability to predict and downscale soil moisture from available coarse resolution satellite information.

## VI. CONCLUSIONS

We develop the SOMOSPIE spatial inference engine for soil moisture data. This suite of cyberinfrastructure tools tackles the two main limitations of satellite-based soil moisture information: coarse granularity and spatial gaps. We demonstrate the potential of our engine by testing and comparing modeling decisions to predict the spatial variability of soil moisture across the Middle Atlantic Coastal Plains region of the United States. The modeling functionalities of SOMOSPIE include options for variable selection, data preprocessing, and method selection. Along with satellite-based soil moisture information, we integrate hydrologically meaningful prediction factors for soil moisture based on topography. Data preprocessing ca-

pabilities include training domain selection and data dimension reduction. For modeling method selection, our inference engine includes standard machine learning methods based on kernels (i.e., KKNN) and regression trees (i.e., Random Forests), and also integrates new modeling functionality with novel methods, such as HYPPO, not previously used for downscaling spatial data.

To assess modeling decisions, SOMOSPIE includes tools to validate and visualize predictions. For our study region, KKNN performs poorly with the 15 topographic covariates (with a mean  $R^2$ -value of 0.296 when predictions are compared to the original satellite observations) as compared to RF with the 15 topographic covariates and HYPPO without (with means  $R^2$  values of 0.575 and 0.557, respectively). We additionally demonstrate preprocessing decisions suitable for maximizing the effectiveness of data-driven soil moisture inference. One preprocessing option is PCA dimension reduction, which we use on data fed to KKNN and RF. Validation of the output indicates that PCA is a viable tool for lightening computational load without significantly affecting the result. Another option is to expand the area of training data beyond our ecologically specific region of interest. This negatively affects the prediction capabilities of KKNN, indicating both the sensitivity of the particular method to the decision and the importance of carefully choosing one's training area within the prediction domain. Overall, we demonstrate the functionality of our tool to provide insights into where and why different methods yield distinct predictions.

Motivation for the SOMOSPIE system derives from the pressing need to improve spatial representation of soil moisture across the world for several applications in environmental sciences. Due to climate change (specifically, increasing temperatures), arid environments are expected to increase and ecosystem services (e.g., water and carbon cycling) across these areas may be at risk. Therefore, accurate soil moisture estimates are necessary to identify priority areas for soil resource conservation efforts and improve management decisions and Earth system models. Future work will consider validation against field observations (e.g., from the American Soil Moisture Database), comparison with other soil moisture information sources (e.g., NASA-SMAP), and more in-depth data-driven model tuning, with the ultimate goal to provide accurate soil moisture estimates using a globally applicable modeling framework.

## ACKNOWLEDGMENT

We note equal contributions of the first two authors, D. R. and M. G. We thank Travis Johnston for a Python implementation of HYPPO and Anita Schwartz for helping prepare the 15 topographic parameters from the DSM. We also thank Paula Olaya and Elizabeth Racca for contributing to preliminary prototyping, brainstorming, and experimenting during a summer 2017 research program at the University of Delaware. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) [39] Jetstream at Indiana University through allocation TG-EAR180011. M. G.

acknowledges support from a Conacyt fellowship; M. T. and R. V. acknowledge support from NSF (#1724843, #1854312); and R. V. acknowledges partial support from NSF (#1652594).

## REFERENCES

- [1] R. Koster, P. Dirmeyer, Z. Guo, G. Bonan, E. Chan, P. Cox, *et al.*, “Regions of strong coupling between soil moisture and precipitation,” *Science*, vol. 305(5687), pp. 1138–1140, 2004.
- [2] O. Denmead and R. Shaw, “Availability of soil water to plants as affected by soil moisture content and meteorological conditions,” *Agron. J.*, vol. 54, pp. 385–390, 1962.
- [3] C. Bauch, R. Sigdel, J. Pharaon, and M. Anand, “Early warning signals of regime shifts in coupled human–environment systems,” *Proc. National Academy of Sciences*, vol. 113(51), pp. 14560–14567, 2016.
- [4] A. Phillips, N. Newlands, S. Liang, and B. Ellert, “Integrated sensing of soil moisture at the field-scale: Measuring, modeling and sharing for improved agricultural decision support,” *Comput. Electron. Agric.*, vol. 107, pp. 73–88, Sept. 2014.
- [5] D. Griffin and K. Anchukaitis, “How unusual is the 2012–2014 California drought?,” *Geophys. Research Lett.*, vol. 41(24), pp. 9017–9023, 2014.
- [6] D. Entekhabi, E. Njoku, P. O’Neill, K. Kellogg, W. Crow, W. Edelstein, *et al.*, “The Soil Moisture Active Passive (SMAP) Mission,” *Proc. IEEE*, vol. 98(5), pp. 704–716, May 2010.
- [7] W. Dorigo, P. Oevelen, W. Wagner, M. Drusch, S. Mecklenburg, A. Robock, and T. Jackson, “A new international network for in situ soil moisture data,” *Eos, Transactions American Geophys. Union*, vol. 92(17), pp. 141–142, 2011.
- [8] National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL), “Soil Moisture Active Passive (SMAP).” <https://smap.jpl.nasa.gov/>.
- [9] European Space Agency (ESA) Climate Change Initiative (CCI), “Soil Moisture CCI.” <https://www.esa-soilmoisture-cci.org>.
- [10] Commission for Environmental Cooperation (CEC), “Ecological Regions of North America.” Geospatial Data Presentation Form: Vector digital data, Apr. 2009.
- [11] W. Wagner, W. Dorigo, R. de Jeu, D. Fernandez, J. Benveniste, E. Haas, and M. Ertl, “Fusion of active and passive microwave observations to create an essential climate variable data record on soil moisture,” *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. I-7, pp. 315–321, 2012.
- [12] Y. Liu, R. Parinussa, W. Dorigo, R. De Jeu, W. Wagner, A. van Dijk, and J. Evans, “Developing an improved soil moisture dataset by blending passive and active microwave satellite-based retrievals,” *Hydrology and Earth System Sciences*, vol. 15(2), pp. 425–436, 2011.
- [13] Y. Liu, W. Dorigo, R. Parinussa, R. de Jeu, W. Wagner, M. McCabe, and A. van Dijk, “Trend-preserving blending of passive and active microwave soil moisture retrievals,” *Remote Sensing of Environment*, vol. 123, pp. 280–297, 2012.
- [14] M. Hirschi, B. Mueller, W. Dorigo, and S. Seneviratne, “Using remotely sensed soil moisture for land-atmosphere coupling diagnostics: The role of surface vs. root-zone soil moisture variability,” *Remote Sensing of Environment*, vol. 154, pp. 246–252, 2014.
- [15] A. Nelson, H. Reuter, and P. Gessler, “Chapter 3: DEM Production Methods and Sources,” in *Geomorphometry* (T. Hengl and H. Reuter, eds.), vol. 33 of *Developments in Soil Science*, pp. 65–85, Elsevier, 2009.
- [16] I. Moore, R. Grayson, and A. Ladson, “Digital terrain modelling: A review of hydrological, geomorphological, and biological applications,” *Hydrological Processes*, vol. 5(1), pp. 1–30, 1991.
- [17] J. Wilson, “Digital terrain modeling,” *Geomorphology*, vol. 137(1), pp. 107–121, 2012. Geospatial Technologies and Geomorphological Mapping Proceedings of the 41st Annual Binghamton Geomorphology Symposium.
- [18] T. Tadono, H. Ishida, F. Oda, S. Naito, K. Minakawa, and H. Iwamoto, “Precise Global DEM Generation by ALOS PRISM,” *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. II-4, pp. 71–76, 2014.
- [19] R. Pike, I. Evans, and T. Hengl, “Chapter 1: Geomorphometry: A Brief Guide,” in *Geomorphometry* (T. Hengl and H. Reuter, eds.), vol. 33 of *Developments in Soil Science*, pp. 3–30, Elsevier, 2009.
- [20] T. Kitson, P. Olaya, E. Racca, M. Wyatt, M. Guevara, R. Vargas, and M. Taufer, “Data analytics for modeling soil moisture patterns across United States ecoclimatic domains,” *2017 IEEE Conference on Big Data*, Dec. 2017. Unpublished conference poster.
- [21] I. Florinsky, “Chapter 8: Influence of Topography on Soil Properties,” in *Digital Terrain Analysis in Soil Science and Geology* (I. Florinsky, ed.), pp. 145–149, Boston: Academic Press, 2012.
- [22] J. Wilson and J. Gallant, “Terrain analysis: Principles and applications,” *Wiley, New York*, p. 479, 2000.
- [23] S. Wold, K. Esbensen, and P. Geladi, “Principal component analysis,” *Chemometrics and Intelligent Laboratory Systems*, vol. 2, p. 3752, Aug. 1987.
- [24] S. Borra and A. Di Ciaccio, “Measuring the prediction error: A comparison of cross-validation, bootstrap and covariance penalty methods,” *Computational Statistics & Data Analysis*, vol. 54(12), pp. 2976–2989, Dec. 2010.
- [25] K. Hechenbichler and K. Schliep, “Weighted k-Nearest-Neighbor Techniques and Ordinal Classification,” *Ludwig-Maximilians-Universität München*. Collaborative Research Center 386, Discussion Paper 399.
- [26] T. Johnston, C. Zanin, and M. Taufer, “HYPPO: A Hybrid, Piecewise Polynomial Modeling Technique for Non-Smooth Surfaces,” *Proc. IEEE 2016 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pp. 1–8, 2016.
- [27] M. Owe, R. de Jeu, and T. Holmes, “Multisensor historical climatology of satellite-derived global land surface moisture,” *J. Geophys. Research: Earth Surface*, vol. 113(F1), 2008.
- [28] W. Wagner, G. Lemoine, and H. Rott, “A Method for Estimating Soil Moisture from ERS Scatterometer and Soil Data,” *Remote Sensing of Environment*, vol. 70(2), pp. 191–207, 1999.
- [29] V. Naeimi, K. Scipal, Z. Bartalis, S. Hasenauer, and W. Wagner, “An Improved Soil Moisture Retrieval Algorithm for ERS and METOP Scatterometer Observations,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47(7), pp. 1999–2013, July 2009.
- [30] V. Naeimi, C. Paulik, A. Bartsch, W. Wagner, R. Kidd, S. Park, K. Elger, and J. Boike, “ASCAT Surface State Flag (SSF): Extracting Information on Surface Freeze/Thaw Conditions From Backscatter Data Using an Empirical Threshold-Analysis Algorithm,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50(7), pp. 2566–2582, July 2012.
- [31] R. Bindlish, T. Jackson, M. Cosh, T. Zhao, and P. O’Neill, “Global Soil Moisture From the Aquarius/SAC-D Satellite: Description and Initial Assessment,” *IEEE Geoscience and Remote Sensing Lett.*, vol. 12(5), pp. 923–927, May 2015.
- [32] V. Pallipuram, J. DiMarco, and M. Taufer, “Applying frequency analysis techniques to dag-based workflows to benchmark and predict resource behavior on non-dedicated clusters,” in *2014 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 29–37, Sept. 2014.
- [33] R. McKinney, V. Pallipuram, R. Vargas, and M. Taufer, “From HPC Performance to Climate Modeling: Transforming Methods for HPC Predictions into Models of Extreme Climate Conditions,” in *2015 IEEE 11th International Conference on e-Science*, pp. 108–117, Aug. 2015.
- [34] C. Ettema and D. Wardle, “Spatial soil ecology,” *Trends in Ecology & Evolution*, vol. 17(4), pp. 177–183, 2002.
- [35] V. Dakos, S. Kéfi, M. Rietkerk, E. van Nes, and M. Scheffer, “Slowing down in spatially patterned ecosystems at the brink of collapse,” *The American Naturalist*, vol. 177(6), pp. E153–E166, 2011. PMID: 21597246.
- [36] E. Leon, R. Vargas, S. Bullock, E. Lopez, A. Panosso, and N. La Scala, “Hot spots, hot moments, and spatio-temporal controls on soil CO<sub>2</sub> efflux in a water-limited ecosystem,” *Soil Biology and Biochemistry*, vol. 77, pp. 12–21, 2014.
- [37] A. McBratney, B. Whelan, T. Ancev, and J. Bouma, “Future directions of precision agriculture,” *Precis. Agric.*, vol. 6, pp. 7–23, 2005.
- [38] J. Schmidhuber and F. Tubiello, “Global food security under climate change,” *Proc. National Academy of Sciences*, vol. 104(50), pp. 19703–19708, 2007.
- [39] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, *et al.*, “XSEDE: Accelerating Scientific Discovery,” *Computing in Science & Engineering*, vol. 16(5), pp. 62–74, 2014.