
Domain Adaptation for Sustainable Soil Management using Causal and Contrastive Constraint Minimization

Somya Sharma^a

Swati Sharma^{b*}

Rafael Padilha^a

Emre Kiciman^b

Ranveer Chandra^b

*swatisharma@microsoft.com

^a University of Minnesota Twin Cities, USA

^b Microsoft Research, Redmond, USA

Abstract

Monitoring organic matter is pivotal for maintaining soil health and can help inform sustainable soil management practices. While sensor-based soil information offers higher-fidelity and reliable insights into organic matter changes, sampling and measuring sensor data is cost-prohibitive. We propose a multi-modal, scalable framework that can estimate organic matter from remote sensing data, a more readily available data source while leveraging sparse soil information for improving generalization. Using the sensor data, we preserve underlying causal relations among sensor attributes and organic matter. Simultaneously we leverage inherent structure in the data and train the model to discriminate among domains using contrastive learning. This causal and contrastive constraint minimization ensures improved generalization and adaptation to other domains. We also shed light on the interpretability of the framework by identifying attributes that are important for improving generalization. Identifying these key soil attributes that affect organic matter will aid in efforts to standardize data collection efforts.

1 Introduction

Background Measuring and monitoring soil organic matter (OM) is pivotal in fighting against climate change [6]. Apart from contributing to improving soil health, organic matter has several co-benefits - such as fighting against soil erosion and preserving the water table [2, 5, 9]. These characteristics make it imperative to preserve organic matter, especially in regions that are becoming vulnerable to droughts and landslides as the climate changes. While monitoring OM is an important problem, sampling and measuring OM (and other soil attributes) is expensive (e.g., the average cost of soil testing in the US is \$1444) [23]. Furthermore, to ensure accurate and reliable measurement of soil properties, it is necessary to conduct sampling at intervals of every tens of meters [3].

Related Work Alternatively, insights about OM can be inferred from other soil characteristics and remote sensing information using data-driven ML methods [8, 11, 22, 24, 13, 15, 4, 17, 19, 7, 20]. While sensor data input offers high-fidelity insights and captures the variability in the dominant soil processes in the region, collecting data from sensors is expensive. On the other hand, remote sensing data is inexpensive and widely available for different parts of the world. Notably, while remote sensing data are underutilized in low-income regions, the benefits derived from use of remote sensing will be the most in these regions where the effect of climate change is prominent[12]. However,

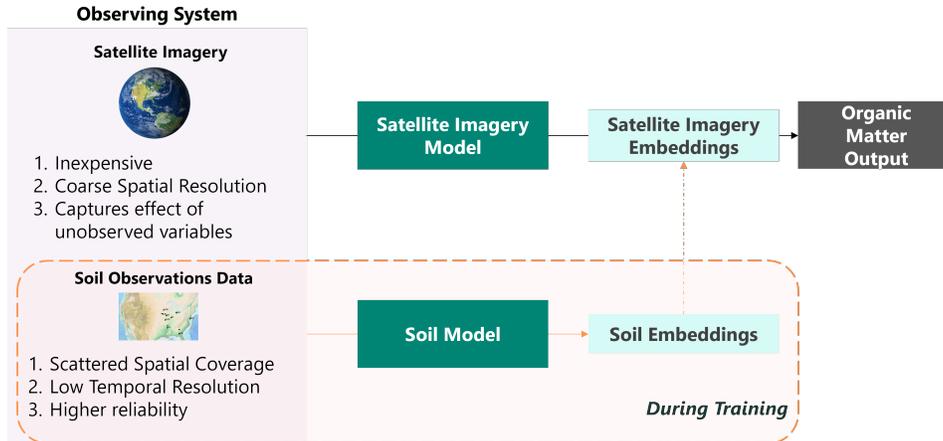


Figure 1: Domain Adaptation framework for Organic Matter Modeling using Causal and Contrastive Constraint Minimization. Conditioning the satellite image embedding using soil attribute embeddings provides additional context or guidance on how the underlying soil properties impact OM.

insights from remote sensing data may be biased due to noise and coarse-scale resolution [1, 10, 26]. In the context of OM mapping, domain adaptation has been useful in overcoming some of the challenges due to remote sensing data and in adapting models to differing land types and countries [16, 21]. While these studies use local soil information as input variables, which may be costly to collect, the studies conclude that domain adaptation helps in improving the generalization to other domains or regions.

Summary of Contributions We propose a scalable and generalizable domain adaptation framework for mitigating bias from remote sensing data using sparse sensor data as auxiliary data. Even if sensor data is not available during inference, leveraging sensor data while training can help capture the variability in the underlying soil processes and improve the generalization of the framework. Such a framework can be built on data-rich entities and be transferred to unobserved entities without further fine-tuning. Conventional ML methods overlook the underlying causal structure in the data, limiting their out-of-distribution generalization. To overcome this, we propose the use of causal constraint minimization to ensure that the relation between the sensor data attributes and OM is preserved in the posterior OM distribution across different regions. Traditional ML methods also struggle with generalization to locations with unobserved drivers (no training data). We leverage the spatial heterogeneity in the sensor attributes to improve generalization to out-of-distribution (OOD) locations by influencing embeddings via contrastive learning. Our analysis is also able to identify key soil attributes that affect organic matter, potentially improving understanding of how to optimize soil management practices and standardize data collection approaches.

2 Methods

The **backbone model** is a CNN autoencoder (Fig. 7) that estimates OM from satellite imagery data (more details in Appendix). Although this model uses satellite imagery input, which is more readily available and captures information about changing soil properties, vegetation, and climate, leveraging soil attributes can help distill the important signals from the images and improve the encodings learned by the autoencoder framework. To achieve this, we propose incorporating two regularization schemes - causal constraint minimization and contrastive learning.

Causally adaptive constraint minimization (CACM) has been implemented to improve generalization under distributional shifts in computer vision benchmark datasets [18]. The method utilizes different independence constraints based on how attributes in training data relate to response variables. These constraints are incorporated in the loss function to regularize the training by adaptively enforcing the correct independence constraints. We extend the **causal constraint minimization** framework to a setting with continuous attributes and continuous response variable on real-world dataset. We modify the framework to incorporate regularization loss terms that

ensure that the encoding space follows the distributional properties that reflect the causal relations among sensor attributes and OM instead of the output from the autoencoder. This helps ensure that the output from the model is not being over-smoothed by the constraint minimization. Following the CACM framework, our model also incorporates three types of causal relations - confounded with OM, caused by OM, or independent of OM. For attributes that are independent of OM, we enforce $\sum_{i=1}^{|A_{ind}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,ind}), P(\phi(x)|a_{j,ind}))$. For attributes that are caused by OM, we enforce $\sum_{i=1}^{|A_{cause}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,cause}, y), P(\phi(x)|a_{j,cause}, y))$. For attributes that are confounded with OM via a confounding variable, we enforce, $\sum_{\|E\|} \sum_{i=1}^{|A_{conf}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,conf}), P(\phi(x)|a_{j,conf}))$. Here, $\phi(x)$ are the encodings obtained from the satellite imagery CNN encoder. To compute the conditional probabilities, for each attribute a , we categorize the samples by attribute value below and above the mean attribute value. Here, A_{ind} , A_{cause} and A_{conf} refer to the sets containing independent, causal and confounded attributes, respectively. Adaptively choosing the regularization term depending on the causal relation enables us to ensure that the independence constraints are also reflected in the conditional embedding distribution.

We incorporate **contrastive learning** using another ANN autoencoder to learn representations for the soil attributes. The embeddings of this ANN model are used to provide additional context on how the underlying soil properties can impact the OM distribution. By incorporating contrastive learning, we capture location-specific patterns while learning to discriminate between the variations associated with different farming environments. We define positive and negative pairs based on differing locations. Samples from the same location are expected to have similar characteristics, including soil type, management practices, and environmental conditions, while different locations are more likely to have dissimilar characteristics. We include triplet loss or contrastive loss during training as $\mathcal{L}_{Contrastive} = D(z_a, z_p) - D(z_a, z_n)$. Here, z_a , z_p , and z_n refer to the embeddings from the anchor domain, positive pair of the domain, and negative pair of the domain. Here, we use Euclidean distance as our distance metric D . In each iteration, for each location, we randomly sample positive and negative pairs to regularize the loss. Here, the positive pair refers to a sample from the same location collected in a different year, and the negative pair refers to a sample from a different location.

The Causal and Contrastive Constraint Minimization approach is shown in Fig. 2.

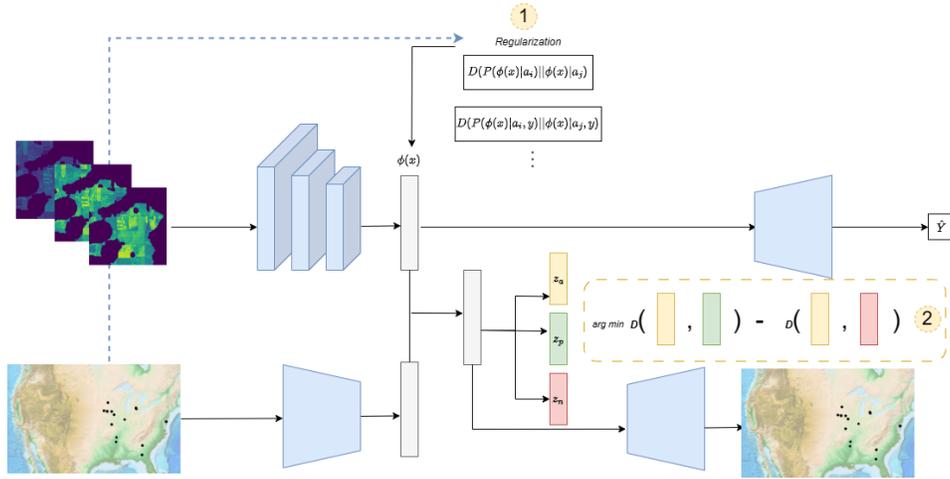


Figure 2: Domain Adaptation framework for Organic Matter Modeling using Causal and Contrastive Constraint Minimization. The bi-level optimization scheme first enforces causal independence constraints and then modifies the embeddings via contrastive learning.

3 Results

Model	Input Data	Auxiliary Training Data	MSE
Random Forest	Satellite Data	-	0.6351
Random Forest	Satellite, Sensor	-	0.2180
CNN	Satellite	-	0.3297
CNN	Satellite, Sensor	-	0.0814
CNN _{CACM}	Satellite	Sensor	0.0578
CNN _{CACM} + Contrastive	Satellite	Sensor	0.0513
CNN _{Contrastive}	Satellite	Sensor	0.0657

Fine-Tuning Env.	CACM	CACM+Contrastive	Contrastive
-	0.3072	0.2354	0.1637
random	0.1998	0.1702	0.1153
closest	0.0688	0.0509	0.0923
farthest	0.1134	0.0977	0.1150

(a) Out of Distribution Generalization (b) Domain Adaptation using CNN backbone model

Table 1: Results on OOD generalization and Domain Adaptation

Data The satellite imagery data is obtained from Sentinel-2. The sensor data is obtained from publicly available Genomes to Fields (G2F) dataset. The sensor data includes information about soil attributes (soil texture, micronutrients, fertilizer application), management practices and also the main variable of interest, OM. More details on included locations, datasets, and preprocessing are given in the Appendix.

Generalization To test the OOD generalization of the framework, we train and test on different locations. Table 1a presents results on OOD generalization results where the test set includes eastern US locations (Georgia and Delaware) while the train set includes locations from the western US. The table compares models that use only satellite data as input with gold standard models that are also able to use sensor data. While using sensor data as input improves generalization, sensor data may be expensive to collect and use during model inference. Therefore, we propose leveraging sensor data to improve embeddings in our model using causal constraint minimization. In such a scenario, sensor data is only used while training the model as an auxiliary dataset that influences encodings. We can further evaluate if we can leverage spatial heterogeneity among the different domains to improve generalization. We leverage contrastive learning to learn encodings that enable us to maximize the similarity between samples from the same location and maximize the dissimilarity between samples from different locations. The results suggest that apart from leveraging contrastive learning, successive implementation of causal and contrastive constraint minimization enables the model to improve OM estimation.

Domain Adaptation To further test how the out-of-distribution generalization can benefit from domain adaptation, we evaluate the model performance in K-fold cross-validation splits where K is the number of different states. The value of K is 6 in these experiments. Table 1b reports how model performance is impacted by fine-tuning the model in different locations relative to the test environment. For each of the K states, we subset one state as fine-tuning location. This way (K-2) locations are used for pertaining to the model, one location is used to fine-tune the model, and one location is used to test in each of the folds. We report the average test MSE for each test split in Table 2. The table shows that fine-tuning using constrain minimization is useful as opposed to not using any fine-tuning data. This may suggest that fine-tuning on another domain allows the model to escape any local minima that is achieved without fine-tuning. Among the three ways of choosing the fine-tuning location, when the closest location is used for fine-tuning, the framework is able to improve generalization. This may be because locations that are closely located may have similar physical and soil farm attributes, improving the transfer of knowledge between domains. As opposed to fine-tuning based on a randomly selected location, fine-tuning on the farthest location results in improved performance. It has been shown that training on heterogeneous locations improves OOD generalization [25]. Training on the farthest locations enables the model to learn to discriminate among different locations. Since causal constraint minimization enables the framework to preserve the underlying relation among the dominant soil processes, the improvement of farthest location-based fine-tuning over random location and no fine-tuning is more evident.

Sensitivity Analysis In order to gain a better understanding of what variables are more impactful in improving generalization as auxiliary variables, we measure variable importance by variable removal. We leave one variable out at a time during training and evaluate how much test MSE increases due to removal. Figure 3 reports the standardized test MSE gain when each of the soil attribute variables is

Test Environment		Delaware	Georgia	Germany	Iowa	Illinois	Michigan	
Fine Tune Environment		Georgia	Illinois	Delaware	Illinois	Iowa	Illinois	Average
CNN _{CACM} + Contrastive	-	0.1843	0.1327	0.3265	0.6043	0.0716	0.0927	0.2354
CNN _{CACM} + Contrastive	FT	0.0235	0.0336	0.0131	0.1528	0.0418	0.0404	0.0509
CNN _{CACM}	-	0.0491	0.0406	0.3921	1.2433	0.0796	0.0382	0.3072
CNN _{CACM}	FT	0.0552	0.2127	0.0136	0.0715	0.0309	0.0289	0.0688
CNN _{Contrastive}	-	0.0195	0.0338	0.1049	0.7750	0.0380	0.0112	0.1637
CNN _{Contrastive}	FT	0.0195	0.0367	0.0162	0.1534	0.3171	0.0110	0.0923

Table 2: Domain Adaptation Model Performance on closest environment. FT: Fine-tuned

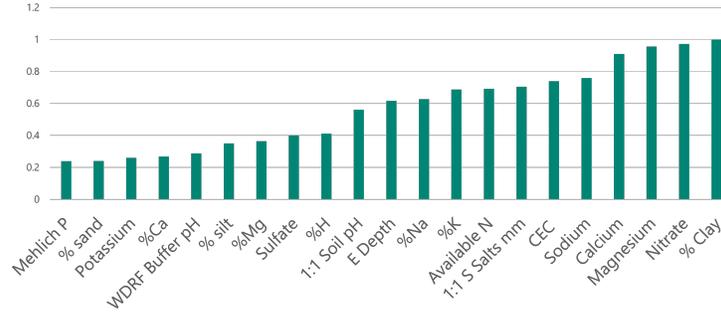


Figure 3: Standardized MSE Gain on Variable Removal

removed from the auxiliary training dataset. The higher the variable importance the more significant the influence on changes in OM. For instance, clay soils tend to have a higher capacity to retain moisture and nutrients, which in turn affects the decomposition of organic matter. Similarly, nitrate, a form of nitrogen readily available to plants, influences plant growth and increases OM through crop residues. This analysis helps shed light on which variables are important in improving out-of-sample generalization as auxiliary variables. The collection of data on these attributes may help farmers draw better insights about their farms. More results provided in the Appendix.

4 Discussion and Future Work

In this paper, we propose using a causal and contrastive constraint minimization mechanism to improve the estimation of organic matter (OM) from remote sensing data. This framework enables us to transfer knowledge from locations with soil information to locations where collecting this information may be infeasible, improving scalability to other regions in the world where data collection is cost-prohibitive. The sensitivity analysis may also help identify key soil characteristics influencing OM. Farmers can use the information on important soil characteristics to make targeted decisions on soil amendments. For example, if nitrogen levels are identified as crucial for organic matter, farmers might adjust their fertilizer application strategies. Moreover, if soil characteristics favoring organic matter are identified, farmers may consider reduced tillage practices to preserve soil structure and organic matter content. These efforts can help optimize soil management practices to allow more precision and efficient use of resources. There are several extensions to the study that can be explored. For example, although our study accounts for changes due to soil characteristics, it is also important to consider changes in OM due to weather dynamics. Two locations with similar soil texture can have widely varying weather and climate. Due to the limited availability of reliable soil attribute information, the use of more readily available weather data along with remote sensing data will be an important future extension of this work.

References

- [1] Sornkitja Boonprong et al. “The classification of noise-afflicted remotely sensed data using three machine-learning techniques: effect of different levels and types of noise on accuracy”. In: *ISPRS International Journal of Geo-Information* 7.7 (2018), p. 274.
- [2] Per Schjønning et al. “Chapter Two - The Role of Soil Organic Matter for Maintaining Crop Yields: Evidence for a Renewed Conceptual Basis”. In: ed. by Donald L. Sparks. Vol. 150. *Advances in Agronomy*. Academic Press, 2018, pp. 35–79. DOI: <https://doi.org/10.1016/B978-0-12-813333-3.00002-1>.

- 1016/bs.agron.2018.03.001. URL: <https://www.sciencedirect.com/science/article/pii/S0065211318300245>.
- [3] Zerina Kapetanovic et al. “FarmBeats: Improving Farm Productivity Using Data-Driven Agriculture”. In: *SIAM News* (July 2019). <https://sinews.siam.org/Details-Page/farmbeats-improving-farm-productivity-using-data-driven-agriculture>. URL: <https://www.microsoft.com/en-us/research/publication/farmbeats-improving-farm-productivity-using-data-driven-agriculture/>.
 - [4] Hamza Keskin et al. “Digital mapping of soil carbon fractions with machine learning”. In: *Geoderma* (2019). DOI: 10.1016/j.geoderma.2018.12.037.
 - [5] E. E. Oldfield, M. A. Bradford, and S. A. Wood. “Global meta-analysis of the relationship between soil organic matter and crop yields”. In: *SOIL* 5.1 (2019), pp. 15–32. DOI: 10.5194/soil-5-15-2019. URL: <https://soil.copernicus.org/articles/5/15/2019/>.
 - [6] Nils Droste et al. “Soil carbon insures arable crop production against increasing adverse weather due to climate change”. In: *Environmental Research Letters* 15.12 (2020), p. 124034.
 - [7] Mostafa Emadi et al. “Predicting and Mapping of Soil Organic Carbon Using Machine Learning Algorithms in Northern Iran”. In: *Remote Sensing* (2020). DOI: 10.3390/rs12142234.
 - [8] Gerard B.M. Heuvelink et al. “Machine learning in space and time for modelling soil organic carbon change”. In: *European Journal of Soil Science* (2020). DOI: 10.1111/ejss.12998.
 - [9] Rattan Lal. “Soil organic matter content and crop yield”. In: *Journal of Soil and Water Conservation* 75.2 (2020), 27A–32A. ISSN: 0022-4561. DOI: 10.2489/jswc.75.2.27A. eprint: <https://www.jswconline.org/content/75/2/27A.full.pdf>. URL: <https://www.jswconline.org/content/75/2/27A>.
 - [10] Guo Yanan et al. “Cloud detection for satellite imagery using deep learning”. In: *Journal of physics: Conference series*. Vol. 1617. 1. IOP Publishing. 2020, p. 012089.
 - [11] Thu Thuy Nguyen. “Predicting agricultural soil carbon using machine learning”. In: *Nature Reviews Earth & Environment* (2021). DOI: 10.1038/s43017-021-00243-y.
 - [12] Esther Rolf et al. “A generalizable and accessible approach to machine learning with global satellite imagery”. In: *Nature communications* 12.1 (2021), p. 4392.
 - [13] Ali Sakhaee et al. “Performance of three machine learning algorithms for predicting soil organic carbon in German agricultural soil”. In: *null* (2021). DOI: 10.5194/soil-2021-107.
 - [14] Mingmin Zhao, Peder A. Olsen, and Ranveer Chandra. “Seeing Through Clouds in Satellite Images”. In: *CoRR* abs/2106.08408 (2021). arXiv: 2106.08408. URL: <https://arxiv.org/abs/2106.08408>.
 - [15] Di An and Yangquan Chen. “A Soil Carbon Content Quantification Method Using A Miniature Millimeter Wave Radar Sensor and Machine Learning”. In: *IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications* (2022). DOI: 10.1109/mesa55290.2022.10004474.
 - [16] Petar Bursać, Miloš Kovačević, and Branislav Bajat. “Instance-based transfer learning for soil organic carbon estimation”. In: *Frontiers in Environmental Science* 10 (2022), p. 1003918.
 - [17] Hassan Fathizad et al. “Spatiotemporal Assessment of Soil Organic Carbon Change Using Machine-Learning in Arid Regions”. In: *Agronomy* (2022). DOI: 10.3390/agronomy12030628.
 - [18] Jivat Neet Kaur, Emre Kiciman, and Amit Sharma. “Modeling the data-generating process is necessary for out-of-distribution generalization”. In: *arXiv preprint arXiv:2206.07837* (2022).
 - [19] Xiangtian Meng et al. “An advanced soil organic carbon content prediction model via fused temporal-spatial-spectral (TSS) information based on machine learning and deep learning algorithms”. In: *Remote Sensing of Environment* (2022). DOI: 10.1016/j.rse.2022.113166.
 - [20] Ali Sakhaee et al. “Spatial prediction of organic carbon in German agricultural topsoil using machine learning algorithms”. In: *Soil* (2022). DOI: 10.5194/soil-8-587-2022.
 - [21] Zefang Shen et al. “Deep transfer learning of global spectra for local soil carbon monitoring”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 188 (2022), pp. 190–200.
 - [22] Zhi Tang et al. “Estimation of National Forest Aboveground Biomass from Multi-Source Remotely Sensed Dataset with Machine Learning Algorithms in China”. In: *Remote Sensing* (2022). DOI: 10.3390/rs14215487.

- [23] Katy Willis. “How much does it cost to test soil and what does it include?” en. In: *Angi* (Nov. 2022). URL: <https://www.angi.com/articles/how-much-does-testing-soil-cost.htm>.
- [24] Xianglin Zhang et al. “Digital Mapping of Soil Organic Carbon with Machine Learning in Dryland of Northeast and North Plain China”. In: *Remote Sensing* (2022). DOI: 10.3390/rs14102504.
- [25] Jared D Willard et al. “Time Series Predictions in Unmonitored Sites: A Survey of Machine Learning Techniques in Water Resources”. In: *arXiv preprint arXiv:2308.09766* (2023).
- [26] Mingmin Zhao, Peder Olsen, and Ranveer Chandra. “Seeing through clouds in satellite images”. In: *IEEE Transactions on Geoscience and Remote Sensing* (2023).

Appendix

Study Area

The sites included in the study are obtained from the publicly available Genomes to Fields (G2F) dataset. The locations of the sites are given in the Table 3. For these sites, the data set provides accurate and reliable soil attribute information. Modeling using these sites enables us to make accurate predictions. The G2F data contains information from 59 environments (where an environment refers to a location in a given year) over 6 locations. The pretrained encoder is pretrained using satellite imagery from all 59 environments on multiple days in the year. For the results on domain adaptation and generalization, the satellite imagery on the day of soil sampling is used as input to minimize the bias in input data since the organic matter values are available on the days when the soil was sampled in a given location in a year. Experiments on early predictions suggest that using satellite imagery data from January 1st of each year instead of the day of sampling also provide reasonable generalization. This enables to use the model even when we do not have information on when the soil was sampled in the year.

deh	gah	iah	ilh	mih	geh
Delaware	Georgia	Iowa	Illinois	Michigan	Germany

Table 3: Site Names



Figure 4: Sites with Soil Attribute Information

Soil Attributes Definition

While the complete information on data collection, variable definition for all variables included in this study are given in Genomes to Fields (G2F) dataset, we define the variables that are predicted to be important as auxiliary training variables.

Differences in Environments

We also show the differences in the environments using pairplot in the Figure 5. Note in particular the difference in the OM distribution for Delaware and Germany in Fig. 6. In Table 1b, the results show that the proposed approach learns better even when fine tuned with a location farthest from the test location with distribution differences.

Satellite Imagery processing

We pre-process the raw Sentinel 2 imagery to remove clouds using the SpaceEye algorithm [14]. The following bands are retained after pre-processing: B02, B03, B04, B05, B06, B07, B08, B8A, B11

Variable Name	Variable Definition
% Clay	Percentage of clay composition in soil sample
Nitrate	Available nitrate in parts per million (ppm)
Magnesium	Available Magnesium in ppm
Calcium	Available Calcium in ppm
Sodium	Available sodium in ppm
CEC	Cation Exchange Capacity (me/100g)
1:1 S Salts mm	Soluble salts concentration in soil
Available N	Amount of nitrogen in pounds per acre
% K	Percentage of Potassium
% Na	Percentage of Sodium
E Depth	Soil sample collection depth
1:1 Soil pH	Soil pH level in a mixture, by weight, one-part soil to one-part distilled H2O
% H	Percentage of Hydrogen
Sulfate	Available sulfate in ppm
% Mg	Percentage of Magnesium
% silt	Percentage of silt composition in soil sample
WDRF Buffer pH	Woodruff method for measuring total soil acidity
% Ca	Percentage of Calcium
Potassium	Available Potassium in ppm

Table 4: Variable Definitions

and B12. For both the satellite imagery and soil attributes, we use min-max scaling to standardize the data.

Model Overview

Backbone Model

The backbone model is a CNN autoencoder showed in the Fig. 7.

Causal and Contrastive Constraint Minimization

Causal Constraint Minimization

Following the work constraints enforced in CACM [18], the regularization constraints for attributes that are independent, caused by OM and confounded with OM, respectively, are as follows,

$$\sum_{i=1}^{|A_{ind}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,ind}), P(\phi(x)|a_{j,ind})) \quad (1)$$

$$\sum_{i=1}^{|A_{cause}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,cause}, y), P(\phi(x)|a_{j,cause}, y)) \quad (2)$$

$$\sum_{|E|} \sum_{i=1}^{|A_{conf}|} \sum_{j>i} MMD(P(\phi(x)|a_{i,conf}), P(\phi(x)|a_{j,conf})) \quad (3)$$

More details on how to leverage CACM can be found in the work Kaur et al. [18]. In our work, we use the causal graph given in Figure 8 to identify which constraints to enforce.

Results on Pre-training the Encoder

Figure 9 shows the change in test MSE when the encoder is pretrained using more satellite data from other time periods. Pretraining on a larger satellite dataset helps the encoder learn the underlying physical patterns that are present in the remote sensing data.

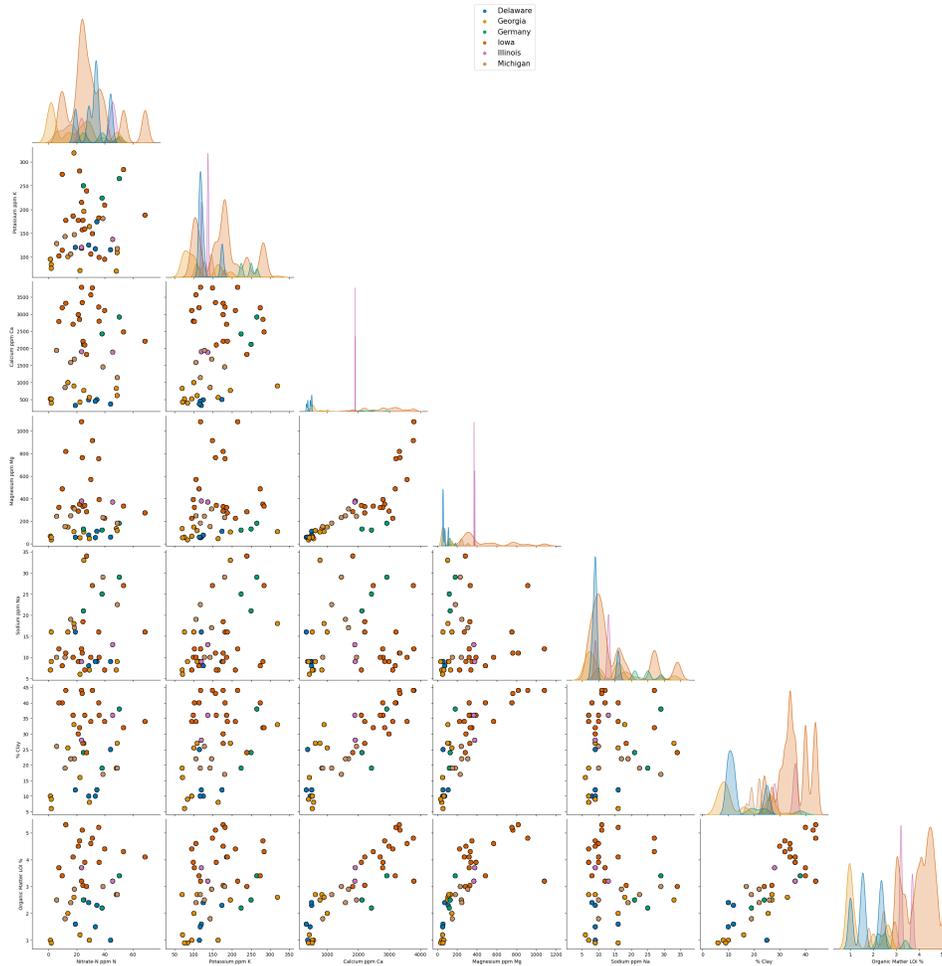


Figure 5: Pairplot showing relationships between different soil variables.

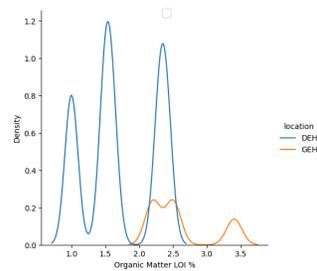


Figure 6: Distribution of OM for Delaware (DEH) and Germany (GEH) locations.

Results on CACM versus Encoding-based CACM

In this study, we report the results in Fig. 5 using encoding-based CACM, which regularizes the encoding space instead of the original variable space as was originally proposed. This reduces any over-smoothing in the output space. This further allows the decoder to focus on train parameters that only focus on OM estimation. We provide results on the domain adaptation experiments wherein we finetune the model on the farthest environment.

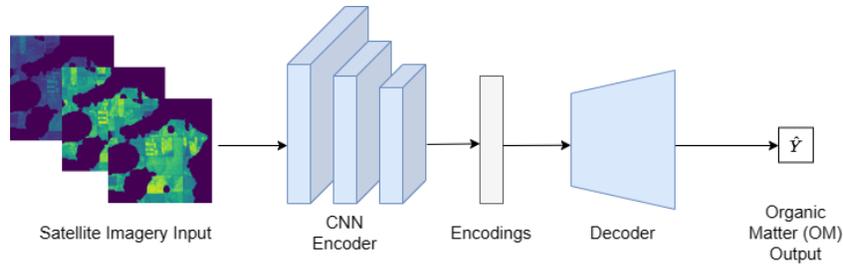


Figure 7: The backbone model is a CNN autoencoder. The input to the model is satellite imagery data and the response is organic matter value. The encoder includes 3 sets of convolution layers stacked with ReLU activation and max-pooling layers. The encoder maps the information to an array of encodings that are used as input to the decoder, which includes fully connected layers.

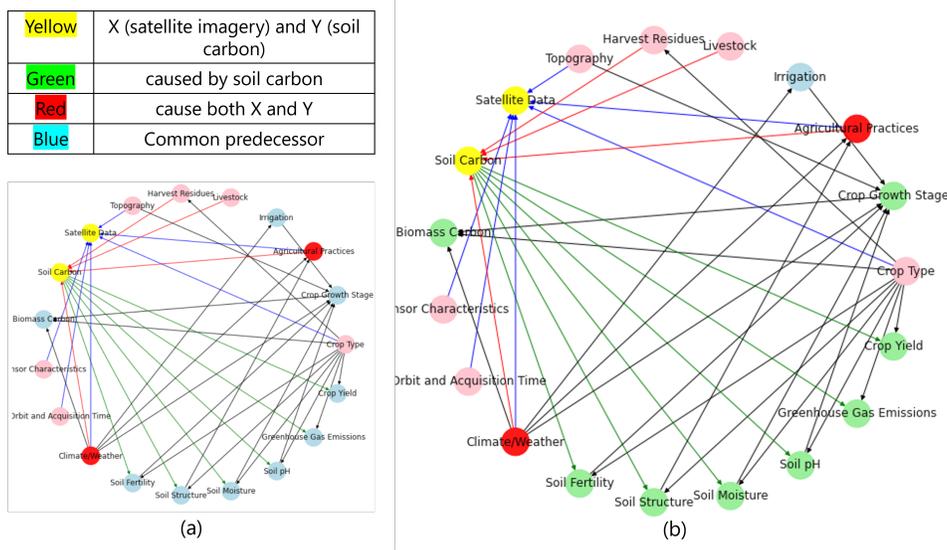


Figure 8: Causal graph among the soil attributes. Several attributes in the data are both caused by Y and confounded with Y . (a) The sub-figure shows a version of the causal graph where these nodes are treated as confounded with Y . (b) The sub-figure shows a version of the causal graph where these nodes are treated as caused by Y . The empirical generalization performance for the graph in sub-figure (a) is better.

Results on Sensitivity Analysis

Table 6 also provides a comparison of variable ranking obtained from this sensitivity analysis with ranking for similar analysis when the soil attributes are used as input data.

These rankings are in agreement for several of the variables, such as percentage of clay (% Clay) in soil, magnesium (ppm) in soil, and CEC. However, variables that have been known to significantly impact organic matter content, such as nitrate, have a disagreement in ranking when used as auxiliary variables and as input variables. This agreement may arise since the ranking derived based on auxiliary variables also accounts for input data satellite imagery that takes care of other unobserved variables whose impact we are unable to measure otherwise. The disagreement may also arise because of a difference in the role of the variables during training - while the impact of the input variable is directly mapped to response, the auxiliary variable in our framework is used indirectly to influence the training process. Interestingly, some variables, like the percentage of potassium in the soil, which are important in determining changes in organic matter, do not come out as important as the input variable. This may be due to the effect of unobserved confounders or noise in data since the correlation of this variable with OM is also relatively lower.

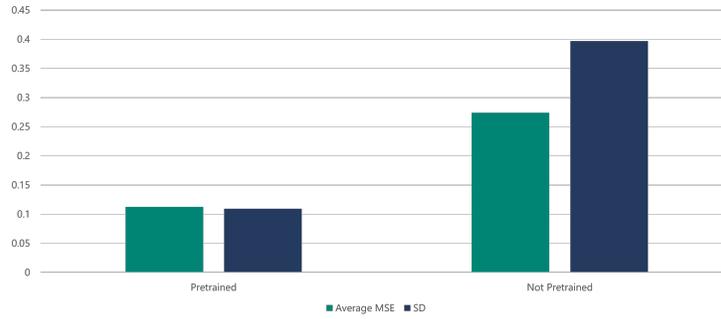


Figure 9: Pretrained Encoder

Model	MSE
CNN_{CACM}	0.1134
$CNN_{Original\ CACM}$	0.1308

Table 5: MSE comparison between original CACM and encoding based CACM.

Variable	Importance as auxiliary variable	Importance as input variable
% Clay	1	1
Nitrate	2	14
Magnesium	3	3
Calcium	4	10
CEC	6	4
% Potassium	9	27
Available Nitrogen	8	13

Table 6: Variable Ranking by Descending Importance