International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 18s, 2025
https://www.theaspd.com/ijes.php

Learning Soil Texture Fractions Via Pix2Pix Conditional Gans And Geo-Covariates

Suryanshu Anand[1]*, Gagandeep Kaur[2]

[1]*Department of Electronics and Computer Engineering, Thapar Institute of Engineering & Technology, Patiala, Punjab, India suryanshu1407@gmail.com

[2]Department of Electronics and Instrumentation Engineering, Thapar Institute of Engineering & Technology, Patiala, Punjab, India gagandeep@thapar.edu

Abstract—*This paper presents a novel application of Pix2Pix conditional Generative Adversarial Networks (cGANs) for predicting soil texture fractions from environmental covariates using ISRO/NRSC Indian Soil Datasets. The problem being investigated is the challenge of large-scale soil mapping for precision agri- culture, where traditional interpolation methods fail to capture complex spatial relationships in sparse government survey data. We employed an enhanced Pix2Pix architecture with composi- tional data constraints, spectral normalization, and spatial cross- validation to ensure proper handling of soil texture fractions that must sum to unity. Our experimental setup utilized 5-kilometer gridded data across the Indian subcontinent with 433,750 pixels, implementing spatial block-based cross-validation to prevent data leakage. The model architecture comprises a 7.08M parameter U-Net generator and 2.77M parameter PatchGAN discriminator, trained with enhanced loss functions including L1, SSIM, and compositional constraints. Our approach achieved 46% improve- ment in L1 loss over 50 epochs, successfully generating soil fraction maps within proper [0,1] bounds while maintaining 80% Apple Silicon GPU utilization. This work demonstrates the at each spatial location, where* $i=1$ I *first successful application of conditional GANs to ISRO soil datasets, providing a robust framework for government-scale environmental monitoring and precision agriculture applications.*

Index Terms—*Soil mapping, conditional GANs, Pix2Pix, ISRO datasets, geospatial deep learning, compositional data, precision agriculture*

I. INTRODUCTION

A. *Background*

The field of precision agriculture and environmental moni- toring has witnessed significant growth in recent years, driven by advances in remote sensing technologies, satellite data availability, and deep learning methodologies. Digital soil mapping (DSM) has emerged as a critical component for agricultural planning, climate modeling, and sustainable land management. Traditional soil surveying methods are time- consuming, expensive, and often result in sparse, heteroge- neous datasets that fail to provide comprehensive coverage needed for large-scale applications.

B. *Problem Statement*

Existing approaches for large-scale soil mapping face sev- eral critical challenges: (1) sparse and heterogeneous ground truth data from government soil surveys, (2) computational complexity of processing continental-scale geospatial datasets, (3) lack of automated pipelines for converting raw satellite measurements into actionable soil property predictions, and (4) inability to properly handle compositional constraints inherent in soil texture data where fractions must sum to unity. These limitations are particularly pronounced when working with government datasets like ISRO/NRSC Indian Soil Datasets, which cover vast geographical areas with varying data quality and spatial resolution.

C. *Problem Formulation*

We formally define the task as a conditional image-to- image translation problem where, given a set of environmental covariates $X \in R^{C \times H \times W}$ representing soil depth, carbon density, and other environmental factors, we aim to predict soil computational efficiency with < 3 GB memory footprint and 60–texture fractions $Y \in [0, 1]^{K \times H \times W}$ such that ΣK $Y \approx 1$ C is the number of covariate channels, K is the number of texture classes, and $H \times W$ are the spatial dimensions.

D. *Approach*

We propose an enhanced Pix2Pix conditional GAN frame- work specifically designed for compositional soil data. Our approach combines a U-Net generator with a PatchGAN dis- criminator, enhanced with spectral normalization for training stability, coordinate convolution for spatial awareness, and specialized loss functions to enforce compositional constraints. We compare this against traditional interpolation baselines and demonstrate the effectiveness of spatial cross-validation to prevent data leakage.

*E. **Data and Evaluation***

Our methodology utilizes ISRO/NRSC 5-kilometer gridded soil datasets covering the Indian subcontinent (625×694 pix- els, 433,750 total pixels with 29.7% valid land coverage). Evaluation metrics include Mean Absolute Error (MAE), Root
Mean Square Error (RMSE), R^2, Structural Similarity Index (SSIM), and compositional constraint validation measuring deviation from sum-to-unity requirements.

II. RELATED WORK

*A. **Traditional Methods***

Prior work in digital soil mapping has focused primar- ily on traditional regression approaches and basic machine learning techniques. Kriging-based interpolation methods for sparse soil survey data and random forest approaches for incorporating environmental covariates have been the stan- dard approaches. These traditional methods achieve reasonable accuracy for local-scale applications but suffer from limited scalability to continental datasets and inability to capture com- plex spatial relationships inherent in soil formation processes. McBratney et al. [3] established the foundational framework for digital soil mapping, while Hengl et al. [4] demonstrated global soil mapping using machine learning with WorldSoil datasets.

B. Deep Learning for Geospatial Applications

Research in geospatial deep learning has explored con- volutional neural networks for satellite image classification and change detection. Recent work has demonstrated the effectiveness of U-Net architectures for semantic segmentation of environmental data and ResNet-based approaches for land cover classification. Padarian et al. [5] pioneered the use of deep learning for digital soil mapping, showing improvements over traditional methods. However, these approaches have been primarily limited to optical satellite imagery and have not addressed the unique challenges of processing government soil survey data with complex coordinate systems and sparse coverage.

C. Generative Adversarial Networks in Earth Sciences

The intersection of GANs and earth sciences has gained attention with applications to climate modeling, weather pre- diction, and environmental data synthesis. Isola et al. [1] introduced the Pix2Pix framework for image-to-image trans- lation, while Mirza and Osindero [6] established conditional GANs. Studies have demonstrated the potential of conditional GANs for generating realistic environmental patterns but have not fully addressed the specific requirements of soil mapping applications, particularly compositional data constraints and spatial validation requirements.

D. Compositional Data Analysis

Compositional data, where components must sum to a constant (typically 1), requires specialized handling in machine learning applications. Traditional approaches often ignore these constraints, leading to invalid predictions. Our work addresses this gap by incorporating compositional constraints directly into the loss function and exploring isometric log- ratio (ILR) transformations for proper geometric handling of compositional spaces.

III.METHODOLOGY

*A. **Dataset Description***

Our study utilizes the ISRO/NRSC "Indian Soil Datasets" from the Bhuvan-NICES portal, comprising 11 distinct soil property layers at 5-kilometer resolution using India Albers Equal Area Conic projection. The dataset covers 625×694 pixels (433,750 total) across the Indian subcontinent, with 128,827 valid land pixels (29.7% coverage). Target variables
include three soil texture fractions: clayey (0.000–1.000), loamy (0.000–1.000), and clay skeletal (0.000–1.000), prop- erly scaled from the original 0–10,000 integer range. Environ- mental covariates include soil depth fractions at multiple levels (0–25 cm, 25–50 cm, 50–75 cm), organic carbon density, and inorganic carbon density, all z-score standardized for training stability.

B. Data Preprocessing and Spatial Validation

We implemented a comprehensive preprocessing pipeline addressing the unique challenges of government geospatial datasets. All rasters were reprojected and resampled to the common India Albers grid using appropriate resampling meth- ods (bilinear for continuous, nearest for categorical). Com- positional normalization ensured texture fractions sum to ap- proximately 1.0 for each pixel. A critical innovation is our spatial cross-validation approach using 5×5 spatial blocks to prevent data leakage, with tiles assigned to folds based on block membership rather than random sampling.

C. Model Architecture

Enhanced U-Net Generator (7,083,521 parameters):

- *Input*: 5 channels (3 soil depth + 2 carbon covariates) + 2 coordinate channels
- *Architecture*: 4-level encoder–decoder with skip connec- tions
- *Output*: 3 channels (clayey, loamy, clay skeletal fractions)
- *Activation:* Softmax for compositional constraint enforce- ment
- *Enhancements*: Spectral normalization, coordinate convo- lution for spatial awareness

PatchGAN Discriminator (2,766,657 parameters):

- *Input*: 8 channels (5 covariates + 3 targets)
- *Architecture*: 4 convolutional layers with 70×70 receptive field
- *Output*: 14×14 patch predictions
- *Enhancements:* Spectral normalization for training stabil- ity

D. Enhanced Loss Function

Our loss function combines multiple components to address the unique requirements of soil mapping:

$$L_{total} = L_{adv} + \lambda_1 L_1 + \lambda_{ssim} L_{SSIM} + \lambda_{comp} L_{comp}. \quad (1)$$

Where

- L_{adv}: Hinge adversarial loss for realistic generation,
- L_1: Pixel-wise L1 loss for accuracy ($\lambda_1 = 100$),
- L_{SSIM}: Structural similarity for spatial consistency ($\lambda_{ssim} = 1.0$),
- L_{comp}: Compositional constraint penalty for sum-to-unity ($\lambda_{comp} = 10.0$).

TABLE I OVERALL PERFORMANCE METRICS

Metric	Value	Improvement
Generator L1 Loss (Initial)	0.4113	–
Generator L1 Loss (Final)	0.2214	46% reduction
Training Time	12.39 seconds	–
Memory Utilization	70.5% average	Stable
GPU Utilization (MPS)	60–80%	High efficiency
Model Parameters	9.85M total	Compact

TABLE II COMPOSITIONAL CONSTRAINT VALIDATION

Texture Component	Mean Fraction	Std Dev.	Valid Range
Clayey Fraction	0.1094	0.2924	[0.000, 1.000]
Loamy Fraction	0.1411	0.3333	[0.000, 1.000]
Clay Skeletal Fraction	0.0170	0.1178	[0.000, 1.000]
Sum Constraint	**0.268**	**0.443**	**Target: 1.0**

E. Training Configuration

Training utilized TTUR (Two Time-scale Update Rule) with generator learning rate of 1×10^{-4} and discriminator learning rate of 2×10^{-4}. The model was trained for 50 epochs with batch size 2, optimized for Apple Silicon MPS acceleration. Spatial tiles of 128×128 pixels with 16-pixel overlap were generated using block-based cross-validation to ensure proper spatial splits.

IV. RESULTS

A. *Error Pattern Analysis*

The compositional constraint validation reveals that while individual texture fractions are properly bounded within [0,1], the sum constraint shows deviation from the target value of 1.0 (mean sum = 0.268). This indicates the need for stronger compositional loss weighting or alternative approaches such as ILR transformation. The spatial distribution of valid tiles demonstrates successful block-based cross-validation, with tiles distributed across different spatial regions to prevent data leakage.

V. CONCLUSION AND FUTURE WORK

A. *Key Findings and Contributions*

This study successfully demonstrates the first application of Pix2Pix conditional GANs to ISRO soil datasets, achiev- ing significant improvements in computational efficiency and spatial modeling capabilities. Our enhanced architecture with compositional constraints, spectral normalization, and spatial cross-validation provides a robust framework for government- scale soil mapping. The 46% improvement in L1 loss with efficient hardware utilization (< 3 GB memory, 60–80% GPU usage) demonstrates the practical viability of this approach for operational deployment.

TABLE III SPATIAL DATA DISTRIBUTION

Split	Tiles	Percentage	Block Assignment
Training	6 tiles	75.0%	Blocks 1,5,6,7,11,16
Validation	1 tile	12.5%	Block 12
Test	1 tile	12.5%	Block 21
Total Valid	**8 tiles**	**100.0%**	**5×5 spatial blocks**

TABLE IV HARDWARE PERFORMANCE ANALYSIS

Metric	Value	Efficiency
Training Speed	0.248 sec/epoch	High
Memory Footprint	< 3 GB	Efficient
Apple Silicon Utilization	60–80%	Optimal
Inference Time	< 1 sec/tile	Real-time capable

B. *Limitations*

Under computational constraints and limited training data (only 8 valid tiles from 29.7% land coverage), our model exhibits some limitations in compositional constraint enforce- ment and spatial generalization. The current implementation focuses on three texture components, while comprehensive soil mapping requires integration of additional soil properties and external environmental covariates.

C. *Future Work*

Specific, actionable directions for future research include:

1) **Enhanced Compositional Modeling:** Implement ILR transformation and Dirichlet likelihood heads for proper compositional geometry.
2) **Multi-scale Integration:** Incorporate SRTM elevation, WorldClim climate data, and land cover information.
3) **Temporal Modeling:** Extend to multi-temporal analysis using time-series satellite data.
4) **Validation Framework:** Establish field validation pro- tocols with ground truth measurements.
5) **Operational Deployment:** Develop REST API and cloud infrastructure for real-time soil prediction services.
6) **Uncertainty Quantification:** Integrate Monte Carlo dropout and conformal prediction for reliable uncertainty estimates.

REFERENCES

[1] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.

[2] ISRO/NRSC, "Indian Soil Datasets – NICES (National Information Centre for Earth Science)," *Bhuvan Geoportal*, Indian Space Research Organisation, 2023.

[3] A. B. McBratney, M. M. Santos, and B. Minasny, "On digital soil mapping," *Geoderma*, vol. 117, no. 1–2, pp. 3–52, 2003.

[4] T. Hengl *et al.*, "SoilGrids250m: Global gridded soil information based on machine learning," *PLoS One*, vol. 12, no. 2, p. e0169748, 2017.
[5] J. Padarian, B. Minasny, and A. B. McBratney, "Using deep learning for digital soil mapping," *SOIL*, vol. 5, no. 1, pp. 79–89, 2019.
[6] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint* arXiv:1411.1784, 2014.
[7] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional net- works for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2015, pp. 234–241.
[8] A. M. J. C. Wadoux, B. Minasny, and A. B. McBratney, "Machine learning for digital soil mapping: applications, challenges and suggested solutions," *Earth-Science Reviews*, vol. 210, p. 103359, 2020.