
Image-Based Soil Organic Carbon Estimation from Multispectral Satellite Images with Fourier Neural Operator and Structural Similarity

Ken C. L. Wong*
IBM Research Almaden
San Jose, CA, USA

Levente Klein
IBM Research
Yorktown Heights, NY, USA

Ademir Ferreira da Silva
IBM Research
Rio De Janeiro, RJ, Brazil

Hongzhi Wang
IBM Research Almaden
San Jose, CA, USA

Jitendra Singh
IBM Research
Gurgaon, HR, India

Tanveer Syeda-Mahmood
IBM Research Almaden
San Jose, CA, USA

Abstract

Soil organic carbon (SOC) sequestration is the transfer and storage of atmospheric carbon dioxide in soils, which plays an important role in climate change mitigation. SOC concentration can be improved by proper land use, thus it is beneficial if SOC can be estimated at a regional or global scale. As multispectral satellite data can provide SOC-related information such as vegetation and soil properties at a global scale, estimation of SOC through satellite data has been explored as an alternative to manual soil sampling. Although existing works show promising results, most studies are based on pixel-based approaches with traditional machine learning methods, and convolutional neural networks (CNNs) are seldom used. To study the advantages of using CNNs on SOC remote sensing, in this paper, we propose the FNO-DenseNet based on the state-of-the-art Fourier neural operator (FNO). By combining the advantages of the FNO and DenseNet, the FNO-DenseNet outperformed the FNO in our experiments with hundreds of times fewer parameters. The FNO-DenseNet also outperformed a pixel-based random forest by 18% in the mean absolute percentage error. To the best of our knowledge, this is the first work of applying the FNO on SOC remote sensing.

1 Introduction

Carbon sequestration is the process of transferring atmospheric carbon dioxide (CO_2), which is the main cause of global warming, into a carbon pool and storing it securely [6]. Atmospheric CO_2 can be transferred and stored in soil in the form of soil organic carbon (SOC), and the effectiveness of SOC sequestration can be improved by methods such as forest management, conservation tillage of croplands, and agroforestry [9]. As per estimates, soils can sequester 1.85 petagrams of carbon per year from the atmosphere if soil management best practices can be implemented [20], thus SOC is a key component for limiting global warming to well below 2°C above pre-industrial levels. As SOC plays an important role in climate change mitigation [3], it is beneficial if the SOC level can be accurately and efficiently estimated to allow effective monitoring and management.

Estimating SOC at a global scale is a challenging task due to variability in soils, complex biochemical processes, and different management practices [13]. Traditionally, soil is sampled at different depths and lab analyzed to measure SOC, but manual sampling is impractical to monitor SOC change at

*Corresponding author (c1wong@us.ibm.com)




Figure 1: FNO-DenseNet. Multispectral images are stacked as multi-channel inputs. The Fourier layers are (1) implemented by the fast Fourier transform. $N = 8$ is used in the experiments, and all Conv and Fourier layers have 24 output channels. For the original FNO, we use $N = 8$ and 32 output channels without concatenations.

a global scale. In contrast, satellite data can provide a complementary, scalable, and cost-effective alternative. Remote sensing through satellites has the potential to enable measurement, reporting, and verification of SOC across the globe, with year-to-year tracking of carbon storage and carbon cycle disruptions when carbon sequestration practices are implemented [14].

As multispectral and hyperspectral satellite data can provide SOC related information such as vegetation, water, and soil properties, prediction and mapping of SOC using these data is an active research topic [19, 2, 11, 15, 16]. Although the results are promising, most studies use pixel-based regression with traditional machine learning methods such as support vector machines and random forests. Deep learning approaches, especially convolutional neural networks (CNNs), are seldom used [11, 16]. In this paper, to study the advantages of using CNNs on SOC remote sensing, we propose an image-based deep learning approach to estimate SOC from multispectral satellite imagery. By modifying the state-of-the-art Fourier neural operator (FNO) [8], we propose the FNO-DenseNet by combining the advantages of the FNO and DenseNet [4]. We also found that incorporating structural similarity (SSIM) [17] in the loss function can further improve SOC estimation. In our experiments, the FNO-DenseNet outperformed the FNO and had hundreds of times fewer parameters. The FNO-DenseNet also outperformed a modified V-Net [12, 18] and a pixel-based random forest. To the best of our knowledge, this is the first work of applying the FNO on SOC remote sensing.

2 Methodology

2.1 FNO-DenseNet for Image-Based SOC Estimation

The FNO is a deep learning model that was proposed to learn mappings between functions in partial differential equations (PDEs) [8]. As the formulations were developed based on the Green’s function in the continuous space, the model training does not require the actual PDEs, and the trained model is theoretically independent of data resolution and discretization [7]. Therefore, it is beneficial to use the FNO to learn the underlying relation between the multispectral images and SOC data. The neural operator is formulated as iterative updates:

$$v_{t+1}(x) := \sigma(Wv_t(x) + (\mathcal{K}v_t)(x)) \quad \text{with} \quad (\mathcal{K}v_t)(x) := \int_D \kappa(x-y)v_t(y)dy, \quad \forall x \in D \quad (1)$$

where $v_t(x) \in \mathbb{R}^{d_{v_t}}$ is a function of x . $W \in \mathbb{R}^{d_{v_{t+1}} \times d_{v_t}}$ is a learnable linear transformation. σ accounts for normalization and activation. In our work, $D \subset \mathbb{R}^2$ represents the 2D imaging space, and $v_t(x)$ are the outputs of hidden layers with d_{v_t} channels. \mathcal{K} is the kernel integral operator with $\kappa \in \mathbb{R}^{d_{v_{t+1}} \times d_{v_t}}$ a learnable kernel function. As $(\mathcal{K}v_t)(x)$ is a convolution operator, it can be efficiently computed by applying the convolution theorem with the Fourier transform \mathcal{F} :

$$(\mathcal{K}v_t)(x) = \mathcal{F}^{-1}(\mathcal{F}(\kappa) \cdot \mathcal{F}(v_t))(x) = \mathcal{F}^{-1}(R \cdot (\mathcal{F}v_t))(x), \quad \forall x \in D \quad (2)$$

For each k in the frequency domain, $R(k) \in \mathbb{C}^{d_{v_{t+1}} \times d_{v_t}}$ is learnable, and $(\mathcal{F}v_t)(k) \in \mathbb{C}^{d_{v_t}}$. Apart from efficiency, the Fourier transform also provides a global receptive field as all pixels are used to compute the value at each k during implementation.

To reduce the number of parameters and improve feature reuse and propagation, following DenseNet [4], we modify the original FNO by concatenating $v_t(x)$ (Fig. 1). Moreover, instead of using a different $R(k)$ for each k , we find that using the same (shared) $R(k)$ does not diminish the accuracy while the number of parameters can be significantly reduced (e.g., by more than 500 times).

Table 1: Testing results (mean \pm std) from five repeats on the same testing data. The unit of RMSE is g/kg, MAPE is shown in percentage, and SSIM is unitless. The best results of each column are highlighted. Each random forest model had 10 trees with a maximum depth of 10. The numbers of parameters of the modified V-Net, FNO, and FNO-DenseNet are 342K, 34M, and 64K, respectively. Note that the DSSIM loss cannot be applied to the pixel-based random forest.

Loss	MAE			DSSIM			MAE + DSSIM		
Metric	RMSE	MAPE	SSIM	RMSE	MAPE	SSIM	RMSE	MAPE	SSIM
Random forest	2.50 \pm 2.37	45.40 \pm 46.30	0.07 \pm 0.08	—	—	—	—	—	—
Modified V-Net	2.12 \pm 2.26	30.12 \pm 23.00	0.07 \pm 0.06	2.21 \pm 2.25	33.44 \pm 22.23	0.20\pm0.13	2.00 \pm 1.99	29.79 \pm 21.11	0.17 \pm 0.12
FNO	1.97\pm2.03	28.08 \pm 21.13	0.10 \pm 0.09	2.31 \pm 2.34	35.57 \pm 24.22	0.20\pm0.13	1.96 \pm 1.96	28.32 \pm 18.29	0.18\pm0.13
FNO-DenseNet	1.98 \pm 2.09	27.24\pm19.25	0.11\pm0.10	2.16\pm2.13	32.33\pm19.48	0.20\pm0.13	1.89\pm1.75	27.75\pm18.12	0.18\pm0.12

2.2 Loss Function with Structural Similarity

In our early experiments, we used the mean absolute error (MAE) as the loss function as it outperformed the mean squared error. Nevertheless, we found that using the MAE alone produced blurry predictions that lacking the structural details of the ground truths. This is most likely due to the texture-like appearances of the SOC data and the pixel-based MAE cannot account for the local dependencies among pixels. Therefore, SSIM is used [17], whose value ($\in [-1, 1]$) between two nonnegative images \mathbf{x} and \mathbf{y} is:

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \frac{1}{M} \sum_{j=1}^M \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (3)$$

where C_1 and C_2 are small constants to avoid divide by zero. The means (μ_x, μ_y), variances (σ_x^2, σ_y^2), and covariance (σ_{xy}) are computed from a local window j of size 11×11 pixels using a Gaussian weighting function with standard deviation of 1.5 pixels. The SSIM compares the luminance, contrast, and structure between \mathbf{x} and \mathbf{y} . Combining with the MAE, the overall loss function is given as:

$$L = w \times \text{MAE} + 0.5 \times (1 - \text{SSIM}) = w \times \text{MAE} + \text{DSSIM} \quad (4)$$

where $\text{DSSIM} = 0.5 \times (1 - \text{SSIM}) \in [0, 1]$ is the structural dissimilarity. w is a scalar accounting for the difference in magnitude, which is 0.01 in our experiments.

2.3 Data and Training Strategy

Six spectral bands from the Moderate Resolution Imaging Spectroradiometer (MODIS) [1], i.e., blue (459–479 nm), green (545–565 nm), red (620–670 nm), near infrared (841–876 nm), and shortwave infrared (SWIR₁: 1230–1250 nm, SWIR₂: 2105–2155 nm), were used as the predictors. The SOC data in the top 5 cm from SoilGrids [13] were the predictands. Sampling from the regions of USA, Mexico, and Canada, a total of 3059 samples were generated. The dataset was partitioned into 50% for training, 20% for validation, and 30% for testing, and identical partitions were used in all experiments. All images were resampled to 128×128 pixels with a spatial resolution of 500 m. Each model was trained for 400 epochs with the batch size of 32. The Adamax optimizer [5] was used with the cosine annealing learning rate scheduler [10], with the maximum and minimum learning rates as 10^{-2} and 10^{-4} , respectively. Image augmentation with rotation ($\pm 30^\circ$), shifting ($\pm 20\%$), scaling ($\in [0.8, 1.2]$), and flipping were used, and each image had an 80% chance to transform.

3 Experiments

Our proposed FNO-DenseNet was compared with the original FNO [8] and a modified V-Net [12, 18]. Each architecture was trained with three different combinations of the MAE and DSSIM losses. Furthermore, to compare with the pixel-based approach, a random forest with 10 trees and a maximum depth of 10 was trained with the MAE loss by using each pixel as a sample². To account for the randomness in machine learning models, five models (repeats) were trained for each setting (i.e., 45 neural networks and five random forests in total) and the average results are reported.

²We tried up to 50 trees with different maximum depths and the results were similar.

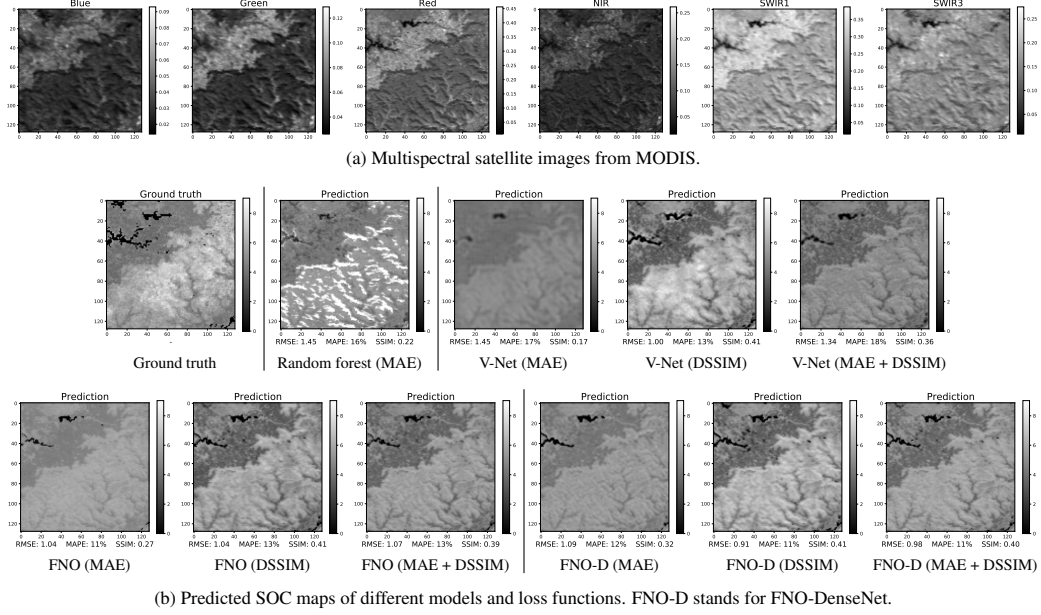


Figure 2: Visualization of a testing example.

3.1 Results

The root mean squared error (RMSE), mean absolute percentage error (MAPE), and SSIM between the predictions and the ground truths on the testing data averaged from the five repeats are reported in Table 1. Comparing among the image-based neural networks, the FNO-DenseNet had the overall best performance and the least number of parameters (64K). Although the original FNO had the second best performance, its number of parameters (34M) was more than 500 times larger than that of the FNO-DenseNet. Comparing among the different loss functions, using MAE alone without DSSIM resulted in better MAPE but worse SSIM, and vice versa. When using MAE + DSSIM, we had the best overall results. On the other hand, the random forest models had the worst performance. When using only the MAE loss, the random forest and the V-Net had worse SSIM than the FNO and the FNO-DenseNet. This may be due to the global receptive field advantage of the Fourier transform.

Fig. 2 shows the visual comparison of a testing example. While the multispectral images had a variety of intensity distributions (Fig. 2(a)), the predictions of the neural networks were similar to the ground truth regardless of the loss functions (Fig. 2(b)). When using only the MAE, regardless of the network architecture, the predictions captured less textural details than those using DSSIM, which is reflected by their relatively low SSIM values. Consistent with Table 1, when using only the MAE, the FNO and FNO-DenseNet had better SSIM values than the random forest and V-Net. The SOC map predicted by the random forest was the most different from the ground truth, which was unable to capture the structural details and some SOC values were overestimated. This is probably caused by the pixel-based nature that does not account for the local dependencies among pixels.

4 Conclusion

In this paper, we propose a satellite image-based approach for SOC remote sensing with multispectral satellite data. The experimental results show that by combining the advantages of the FNO and DenseNet, the FNO-DenseNet outperformed the FNO, V-Net, and random forest with the least number of parameters among the neural networks. Furthermore, by using the structural dissimilarity in the loss function, the learned models can provide predictions with structural details like the ground truths. Given these advantages, the proposed framework can enable measurement, reporting, and verification of SOC across the globe to allow year-to-year tracking of carbon storage and carbon cycle disruptions when carbon sequestration practices are implemented.

References

- [1] MODIS Overview. <https://lpdaac.usgs.gov/data/get-started-data/collection-overview/missions/modis-overview/>. Accessed on September 5, 2022.
- [2] Theodora Angelopoulou, Nikolaos Tziolas, Athanasios Balafoutis, George Zalidis, and Dionysis Bochtis. Remote sensing techniques for soil organic carbon estimation: A review. *Remote Sensing*, 11(6):676, 2019.
- [3] D. A. Bossio, S. C. Cook-Patton, P. W. Ellis, J. Fargione, J. Sanderman, P. Smith, S. Wood, R. J. Zomer, M. von Unger, I. M. Emmer, and B. W. Griscom. The role of soil carbon in natural climate solutions. *Nature Sustainability*, 3(5):391–398, 2020.
- [4] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017.
- [5] Diederik P. Kingma and Jimmy L. Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980 [cs.LG]*, 2014.
- [6] Rattan Lal. Soil carbon sequestration impacts on global climate change and food security. *Science*, 304(5677):1623–1627, 2004.
- [7] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv:2003.03485 [cs.LG]*, 2020.
- [8] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- [9] Klaus Lorenz and Rattan Lal. Soil organic carbon sequestration in agroforestry systems: a review. *Agronomy for Sustainable Development*, 34(2):443–454, 2014.
- [10] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.
- [11] Xiangtian Meng, Yilin Bao, Jiangui Liu, Huanjun Liu, Xinle Zhang, Yu Zhang, Peng Wang, Haitao Tang, and Fanchang Kong. Regional soil organic carbon prediction model based on a discrete wavelet analysis of hyperspectral satellite data. *International Journal of Applied Earth Observation and Geoinformation*, 89:102111, 2020.
- [12] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In *IEEE International Conference on 3D Vision*, pages 565–571, 2016.
- [13] Laura Poggio, Luis M. de Sousa, Niels H. Batjes, Gerard B. M. Heuvelink, Bas Kempen, Eloi Ribeiro, and David Rossiter. SoilGrids 2.0: producing soil information for the globe with quantified spatial uncertainty. *SOIL*, 7(1):217–240, 2021.
- [14] P. Smith, J. F. Soussana, D. Angers, L. Schipper, C. Chenu, D. P. Rasse, N. H. Batjes, F. Van Egmond, S. McNeill, M. Kuhnert, and C. Arias-Navarro. How to measure, report and verify soil carbon change to realize the potential of soil carbon sequestration for atmospheric greenhouse gas removal. *Global Change Biology*, 26(1):219–241, 2020.
- [15] Zander S. Venter, Heidi-Jayne Hawkins, Michael D. Cramer, and Anthony J. Mills. Mapping soil organic carbon stocks and trends with satellite-driven high resolution maps over South Africa. *Science of the Total Environment*, 771:145384, 2021.
- [16] Sheng Wang, Kaiyu Guan, Chenhui Zhang, DoKyoung Lee, Andrew J Margenot, Yufeng Ge, Jian Peng, Wang Zhou, Qu Zhou, and Yizhi Huang. Using soil library hyperspectral reflectance and machine learning to predict soil organic carbon: Assessing potential of airborne and spaceborne optical soil sensing. *Remote Sensing of Environment*, 271:112914, 2022.

- [17] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [18] Ken C. L. Wong, Mehdi Moradi, Hui Tang, and Tanveer Syeda-Mahmood. 3D segmentation with exponential logarithmic loss for highly unbalanced object sizes. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 11072 of *LNCS*, pages 612–619, 2018.
- [19] Daniel Žížala, Robert Minařík, and Tereza Zádorová. Soil organic carbon mapping using multispectral remote sensing data: Prediction ability of data with different spatial and spectral resolutions. *Remote Sensing*, 11(24):2947, 2019.
- [20] Robert J. Zomer, Deborah A. Bossio, Rolf Sommer, and Louis V. Verchot. Global sequestration potential of increased organic carbon in cropland soils. *Scientific Reports*, 7(1):1–8, 2017.