# A Framework for Deep Learning Emulation of Numerical Models With a Case Study in Satellite Remote Sensing

Kate Duffy<sup>®</sup>, Thomas J. Vandal, Weile Wang, Ramakrishna R. Nemani<sup>®</sup>, and Auroop R. Ganguly<sup>®</sup>

Abstract-Numerical models based on physics represent the state of the art in Earth system modeling and comprise our best tools for generating insights and predictions. Despite rapid growth in computational power, the perceived need for higher model resolutions overwhelms the latest generation computers, reducing the ability of modelers to generate simulations for understanding parameter sensitivities and characterizing variability and uncertainty. Thus, surrogate models are often developed to capture the essential attributes of the full-blown numerical models. Recent successes of machine learning methods, especially deep learning (DL), across many disciplines offer the possibility that complex nonlinear connectionist representations may be able to capture the underlying complex structures and nonlinear processes in Earth systems. A difficult test for DL-based emulation, which refers to function approximation of numerical models, is to understand whether they can be comparable to traditional forms of surrogate models in terms of computational efficiency while simultaneously reproducing model results in a credible manner. A DL emulation that passes this test may be expected to perform even better than simple models with respect to capturing complex processes and spatiotemporal dependencies. Here, we examine, with a case study in satellite-based remote sensing, the hypothesis that DL approaches can credibly represent the simulations from a surrogate model with comparable computational efficiency. Our results are encouraging in that the DL emulation reproduces the results with acceptable accuracy and often even faster performance. We discuss the broader implications of our results in light of the pace of improvements in high-performance implementations of DL and the growing desire for higher resolution simulations in the Earth sciences.

*Index Terms*—Bayesian deep learning (DL), emulation, surrogate modeling, uncertainty quantification.

Manuscript received January 26, 2021; revised October 8, 2021; accepted April 15, 2022. This work was supported by the NASA Earth eXchange (NEX). (*Corresponding author: Kate Duffy.*)

Kate Duffy is with the Sustainability and Data Sciences Laboratory, Department of Civil and Environmental Engineering, Northeastern University, Boston, MA 02115 USA, also with the NASA Ames Research Center, Mountain View, CA 94035 USA, and also with the Bay Area Environmental Research Institute, Moffett Field, CA 94035 USA (e-mail: duffy.k@northeastern.edu).

Thomas J. Vandal is with the NASA Ames Research Center, Mountain View, CA 94035 USA, and also with the Bay Area Environmental Research Institute, Moffett Field, CA 94035 USA.

Weile Wang is with the NASA Ames Research Center, Mountain View, CA 94035 USA, and also with the Department of Applied Environmental Science, California State University, Seaside, CA 93955 USA.

Ramakrishna R. Nemani is with the NASA Ames Research Center, Mountain View, CA 94035 USA.

Auroop R. Ganguly is with the Sustainability and Data Sciences Laboratory, Department of Civil and Environmental Engineering, Northeastern University, Boston, MA 02115 USA.

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TNNLS.2022.3169958.

Digital Object Identifier 10.1109/TNNLS.2022.3169958

#### I. INTRODUCTION

ATELLITE remote sensing and climate modeling, which are our best tools for monitoring the climate and projecting the future climate, rely on powerful computers to process massive data streams and execute complex simulations. Over the last several decades, strides in high-performance computing have enabled the development of climate models that solve equations at increasingly fine spatial resolutions. Past iterations of the Coupled Model Intercomparison Project (CMIP) have seen the average horizontal resolution decrease from 250 km (CMIP3), to 150 km (CMIP5) and 25-50 km in the high-resolution experiments of the latest generation (CMIP6) [1]. Despite these strides, even the latest generation of climate models does not provide the type of robust regional climate information required by decision-makers, and the extant literature suggests that high-resolution simulations of fine-scale processes are critical to reproducing extremes and change [2]. Scientific and societal demands for higher resolution simulations easily consume available resources, leaving less capacity to invest in model improvement runs, such as parameters' optimization and characterization of model sensitivity and variability. Lack of resources for science runs such as these may contribute to the relatively modest increases in historical skill and model consensus over time [3]-[5], despite major investments in climate modeling and computing.

In climate science, the only source of data growing faster than archived model simulations is remote sensing observations of the environment [6]. Remotely sensed atmospheric and surface states document historical and current climates, are used to constrain the components of Earth system models [7], and are assimilated into climate reanalysis models [8]. Recent decades have seen an ever-growing fleet of Earth-orbiting satellites and continuous-resolution improvements. Depending on the spatial and temporal resolutions of gridded satellite datasets, as well as the number of attribute fields generated in postprocessing, a low Earth orbit satellite can generate tens to hundreds of gigabytes each year, while geostationary sensors, such as GOES-16, generate nearly one terabyte per day [9], [10].

The cost of running complex models has motivated many studies of surrogate modeling in the Earth sciences. Surrogate models capture the essential attributes of the full high-fidelity models, including their statistical variability, by projecting high-dimensional systems into low-order functions. These

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

proxy models save orders of magnitude in computing time so that scientists can quickly study complex systems where repeated simulations would be prohibitively slow. For example, statistical modeling has been used to emulate the output of climate models under any arbitrary forcing scenarios, given a set of full models runs [11], [12]. Simplified approximations have also been evaluated as potential replacements for time-consuming portions of model physics [13]. In remote sensing, a range of emulation techniques has been proposed for radiative transfer models (RTMs), which represents the scattering and absorption of solar radiation [14]. Even simpler reductions, such as lookup tables, are used operationally to reduce computation time [15], while more complex emulators, such as those based on deep learning (DL), have begun to show encouraging results [16].

DL has transformed the state of the art in computer vision [17], natural language processing [18], and even complex games, such as chess and go [19]. In the Earth sciences, DL has begun to show promise for traditionally challenging problems, such as El Niño/Southern Oscillation (ENSO) forecasting [20] and downscaling of geophysical data [21], [22]. It has been suggested that the ability to extract features from spatial and temporal contexts is at the root of DL's promising results [23], but, in the absence of solid theoretical explanations for DL's success, performance improvements have been greeted as "unreasonable effectiveness" [24]. In scientific domains, where process understanding and consistency with the laws of physics are key, "black box" machine learning models have not been universally embraced. However, given that Earth systems are governed by complex, multiscale, and nonlinear dynamics and contain spatial and temporal structures, several previous studies have begun to examine whether complex models may better capture the complexity of Earth science data than simpler, process-based models [25], [26].

Here, we present a case study comparing DL-based emulation to a traditional form of the surrogate model in remote sensing. A DL model is trained to approximate a lookup table capturing the RTM basis of the Multi-Angle Implementation of Atmospheric Correction (MAIAC) [15] algorithm using a dataset composed of lookup table inputs and outputs (see Fig. 1). We examine the hypothesis that DL approaches can credibly represent the simulations from a surrogate model and can achieve comparable or greater computational efficiency.

#### II. STUDY AREA AND DATASETS

Datasets used in this study are from the Advanced Himawari Imager (AHI) sensor carried by the Japanese geostationary satellite Himawari-8. In the GeoNEX processing pipeline, Himawari Standard Data (HSD) scans are georeferenced and converted to gridded data. The resulting gridded datasets follow a geographic coordinate system with a  $120^{\circ} \times 120^{\circ}$ extent (E85°–E205°, N60°–S60°). The domain is divided into 6° × 6° tiles defined by fixed latitude and longitude. Himawari-8's full disk, which encompasses the entire view as seen from the satellite, covers the continent of Australia and

TABLE I Himwari-8 AHI Solar Reflective Bands for Land Surface Observation

Himwari-8 AHI Band	Blue	Green	Red	NIR	SWIR1	SWIR2
Center Wavelength (µm)	0.46	0.51	0.64	0.86	1.6	2.3
Spatial resolution (km)	1.0	1.0	1.0	0.5	2.0	2.0

eastern Asia. Full disk scans are repeated every 15 minutes throughout the day.

The AHI has 16 observing bands encompassing visible, near-infrared (NIR), short wave infrared (SWIR), and thermal infrared with spatial resolution ranging from 0.5 to 2 km. Bands 1–6 are solar reflective bands, spectrally similar to NASA's Moderate Resolution Imaging Spectroradiometer (MODIS). All bands are resampled to common 0.01° resolution, which corresponds to 1 km at the equator (see Table I).

### A. Himawari-8 AHI TOA Reflectance

The top-of-atmosphere (TOA) reflectance is prepared by the GeoNEX processing pipeline from HSD scans, according to the processing procedure outlined in the Himawari 8/9 HSD User's Guide, Version 1.2 [27]. HSD consists of raw digital counts that are transformed to the bidirectional reflectance factor (bands 1–6) and the brightness temperature (bands 7–16). TOA reflectance data and additional documentation are available from www.nasa.gov/geonex.

# B. Himawari-8 AHI Surface Reflectance

Originally developed for the land-monitoring flagship MODIS, the MAIAC algorithm has been adapted to retrieve surface reflectance (SR) and atmospheric composition for the geostationary satellite Himawari-8. Running an atmospheric RTM for every measurement from a sensor is generally too slow and is not done in practice. Instead, MAIAC relies on the generation of lookup tables (LUTs) from which values are retrieved by linear interpolation [15]. As retrieval error is associated with step size, LUTs are precomputed at a grid density chosen with consideration to both accuracy and memory requirements. Given the numerous possible variations in atmospheric conditions, keeping the LUT within reasonable size guides many choices made in its construction [15]. In practice, the LUT approach gives nearly identical results as the inline RTM, given the necessary ancillary information, such as Sun-satellite geometries and atmospheric optical properties.

Geostationary SR is produced using the adapted MAIAC algorithm for all daylight observations [28]. MAIAC uses a time series of up to 16 days and a mixture of pixel and image-level processing for atmospheric correction with internal cloud detection, aerosol retrieval, and quality assurance flagging [15], [29], [30]. Multiangle determination of SR refers to viewing angle, which is fixed, and illumination angle, which varies continuously throughout the day. AHI MAIAC SR is released as a preliminary product and is available upon request at www.nasa.gov/geonex.



Fig. 1. Surrogate models save orders of magnitude in computing time so that scientists can quickly capture the essential attributes of complex simulations. In the context of remote sensing, we examine the hypothesis that a DL emulator performs credibly and achieves similar computational performance compared to a more traditional form of model order reduction.

# C. MODIS MCD12Q1 Land Cover Type

Land cover types are identified using MODIS global land cover classification, which is produced annually from combined Terra and Aqua observations [31]. MCD12Q1 incorporates five distinct classification schemes. We use the International Geosphere Biosphere Programme (IGBP) global vegetation classification scheme. This scheme delineates 17 distinct classes, including 11 natural, three developed/ mixed, and three nonvegetated. MCD12Q1 is resampled from 500-m resolution to the 0.01° grid of the prepared AHI datasets.

#### **III. METHODS**

# A. Emulator Model

Mapping from TOA to SR relies on computationally expensive radiative transfer equations that simulate nonlinear physics and incorporate ancillary information about atmospheric conditions. MAIAC is a state-of-the-art method for accomplishing atmospheric correction and is run exclusively on an LUT, which achieves low ( $\sim 0.2\%$ ) error with respect to inline calculations. Due to the unavailability of full model simulations in our domain of interest, in our approach to emulation, several deep networks are learned to approximate the input to output transformation contained in the MAIAC LUT.

1) Bayesian Deep Learning: Typical deep neural networks are learned as deterministic functions that fail to capture uncertainty in model parameters and data. The need to quantify these uncertainties has motivated the development of approaches combining Bayesian probability theory and neural networks. Bayesian neural networks learn probability distributions over the network parameters, where training data are used to transform the prior probability distributions into posterior distributions. Approximations of the Bayesian inference have been postulated to extract information about both epistemic uncertainty (due to randomness) and aleatory uncertainty (due to incomplete systemic knowledge) from DL models [32].

For a network with parameters **W**, inputs **X**, and outputs **Y**, the Bayesian theory estimates posterior distribution over the

weights as  $p(\mathbf{W}|\mathbf{X}, \mathbf{Y}) = p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{W})/p(\mathbf{Y}|\mathbf{X})$ . However, performing inference on the full posterior distribution is intractable for complex networks, where analytical solutions to the marginal probability  $p(\mathbf{Y}|\mathbf{X})$  are unavailable [33].

This has led to the development of efficient approximations of the Bayesian inference for DL [34]. One approach to the Bayesian DL defines an approximation of the true posterior distribution by Monte Carlo sampling. Dropout, the process of randomly removing nodes from deep neural networks, is applied as a regularization technique in many DL models to discourage overfitting [35]. When applied during training and testing, the stochastic dropout generates thinned networks, sampling of which is mathematically equivalent to sampling the approximate posterior distribution.

All models defined in this work are Bayesian models implemented with concrete dropout, a variant that dynamically adapts dropout probability to obtain well-calibrated uncertainty estimates [36]. For a model  $f^{\mathbf{W}}(X)$  with L layers, weights  $\mathbf{W}$ , and dropout probabilities p, the set of variational parameters is  $\theta = \{\mathbf{w}_l, p_l\}_{l=1}^L$ . Thus, the variational (factorized) approximation,  $q_{\theta}$ , of the true posterior distribution  $p(\mathbf{W})$  is formalized as  $q_{\theta}(\mathbf{W}) = \prod_{l=1}^L q_{\theta}(\mathbf{w}_l)$ . The optimization objective  $\mathscr{L}(\theta)$  for this variational interpretation of the model can be written as follows for data samples  $\{x_i, y_i\}_{i=1}^N$  [36]:

$$\mathscr{L}(\theta) = -\frac{1}{N} \sum_{i=1}^{N} \log p(y_i | f^{\mathbf{W}}(\mathbf{x}_i)) + \mathrm{KL}(q_{\theta}(\mathbf{W}) || p(\mathbf{W}))$$
$$= \mathscr{L}_{\mathcal{X}}(\theta) + \mathrm{KL}(q_{\theta}(\mathbf{W}) || p(\mathbf{W})). \quad (1)$$

The first term represents expresses the log-likelihood of the model, while the Kullback–Leibler (KL) divergence term acts as a regularizer by discouraging separation between the approximate posterior and the true posterior distribution. The loss function,  $\mathscr{L}_{\mathcal{X}}(\theta)$ , is written as follows for labels  $\{y_i\}_{i=1}^N$  and output values  $\{\hat{y}_i\}_{i=1}^N$ :

$$\mathscr{L}_{\mathcal{X}}(\theta) = -\frac{1}{N} \sum_{i=1}^{N} \frac{1}{2} \hat{\sigma_i}^{-2} ||y_i - \hat{y_i}||^2 + \frac{1}{2} \log \hat{\sigma_i}^2.$$
(2)

The noise in the input data is represented using the uncertainty regularization term  $\hat{\sigma}_i$  [32]. Through minimization of the loss function, the level of noise inherent in the data is learned implicitly. Furthermore, the model's epistemic uncertainty is characterized by the predictive variance [32]. During inference, stochastic forward passes generate *T* independent and identically distributed samples from *T* thinned networks. From these samples, we can empirically approximate the model's predictive distribution at each pixel. With *T* samples of  $[\hat{\mathbf{y}}, \hat{\sigma}]$  from the Bayesian network  $f^{W}(X)$ , the unbiased estimates of the first two moments of the predictive distribution are

$$E[y] = \frac{1}{T} \sum_{t=1}^{T} \hat{y}_t$$
(3)

$$\operatorname{Var}[y] = \frac{1}{T} \sum_{t=1}^{T} \hat{y}_t^2 - \left(\frac{1}{T} \sum_{t=1}^{T} \hat{y}_t\right)^2 + \sum_{t=1}^{T} \hat{\sigma}_t^2.$$
(4)

This estimate of uncertainty is obtained through supervision of the task rather than requiring "uncertainty labels." It theoretically comprises both aleatoric uncertainty, from measurement noise, and epistemic uncertainty, reducible through the collection of more data [37].

2) Discrete-Continuous Distribution: Prediction tasks generally fall into one of two categories: regression tasks predict a continuous quantity, while classification tasks are concerned with assigning a class label. MAIAC's atmospheric correction and cloud classification algorithms generate SR, a continuous variable ranging between 0 and 1, as well as binary cloud classification. We learn a discrete-continuous model to perform both regression and classification tasks in one probabilistic model [38]. To this end, the model is conditioned to predict the probability  $\hat{p}$  of a pixel being a clear sky. For the Bayesian network  $f^{W}(X)$  described in Section III-A1, the mean, variance, and probability are sampled as follows:

$$\hat{y}, \hat{\sigma}^2, \hat{\phi} ] = f^{\mathbf{W}}(X) \tag{5}$$

$$\hat{p} = \text{Sigmoid}(\hat{\phi}).$$
 (6)

This conditioning results in a two-part loss function with the first term capturing cross-entropy of predicted and cloud label and cloud prediction, and the second term capturing conditional regression loss at clear sky pixels. Here, y is the binary indicator of whether the classification is correct and D is the number of pixels with pixel index i. The loss function in (2) is modified as follows to incorporate classification loss:

L

$$\mathscr{L}_{\mathcal{X}}(\theta) = \overbrace{-\frac{1}{D}\sum_{i} \left[y_{i} \log(\hat{p}_{i})\right) + (1 - y_{i}) \log(1 - \hat{p}_{i})\right]}_{i} + \underbrace{\frac{1}{D}\sum_{i, y_{i} > 0} \frac{1}{2}\hat{\sigma}_{i}^{-2}||y_{i} - \hat{y}_{i}||^{2} + \frac{1}{2}\log\hat{\sigma}_{i}^{2}}_{\text{conditional regression loss}}$$
(7)

*3) Implementation and Training:* We implement three full Bayesian architectures conditioned to learn a discrete-continuous distribution as in [39] (see Fig. 2). The ML models are trained without the ancillary datasets (e.g., current atmosphere conditions) used as inputs to MAIAC.

The discrete-continuous fully connected (DCFC) neural network has three layers with 512 filters per layer. The discrete-continuous convolutional neural network (DCCNN) is of similar width and depth to DCFC and has a convolutional filter size of three. Finally, the discrete-continuous very deep super-resolution network (DCVDSR), inspired by image super-resolution networks, is a convolutional neural network similar to DCCNN but incorporates a skip connection between the first and last hidden layers. The presence of the skip connection reformulates the task of the network as learning a residual function or the difference between the output mapping and input [39]. This reformulation improves the performance by preconditioning the function as an identity mapping. Rather than adopting the common bottleneck type architecture, we opt to exclude pooling layers in all three designs. This choice is consistent with the task's dependence on relatively local information and helps to maintain temporal consistency in

DUFFY et al.: FRAMEWORK FOR DL EMULATION OF NUMERICAL MODELS



Fig. 2. Proposed model architectures for DL emulation. We implemented three full Bayesian architectures conditioned to learn a discrete-continuous distribution for SR and cloud mask prediction: (a) DCFC neural network, (b) DCCNN, and (c) DCVDSR.

local features in the presence of global changes in the images. All models use ReLU activations.

Over 60 TB of data from a two-year period is divided into training (January 1, 2016, to December 31, 2016) and testing (January 1, 2017, to December 31, 2017) sets. The training set was used to experiment with model designs, while the test set was used exclusively to evaluate the model and prepare figures. Models are implemented in TensorFlow 2.0 and trained on  $50 \times 50$  pixel patches for approximately 8 hours using the stochastic gradient descent and the Adam optimization with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1e - 7$ , a batch size of 16, and a learning rate of 1e - 4 [40]. All three models use concrete dropout, as described in Section III-A1, for weight regularization and uncertainty quantification. For concrete dropout, hyperparameters tau and prior length scale are set to 1e - 5 and 1e - 14. Code is available at github.com/KateDuffy/maiac-emulator.

## B. Assessment of Emulated SR and Cloud Products

Validation of SR products against *in situ* measurements would result in uncertainties due to the presence of mixed pixels at the remote sensor's spatial resolution [41], [42]. Therefore, comparison with an existing, comprehensively validated product can be used to assess the performance of a new reflectance product [43]. Here, it is worthwhile to note that the

analysis and interpretation of the emulated data are performed in comparison to the lookup table prepared for MAIAC. These lookup tables are an approximation of a fundamentally approximate model and contain associated uncertainties [15]. The task of the emulator can be posed as learning a smooth interpolation of the MAIAC lookup table. As both the MAIAC SR and emulator SR are predicted from AHI TOA reflectance, all pixels are guaranteed coincident, coangled, and colocated, and can be directly compared. In addition, we evaluate the ability of the emulator to discriminate between clear sky and nonclear sky pixels. To evaluate the stability of emulator performance under varied conditions, results are presented for performance for the nine common most MODIS land cover classifications and four seasons.

## IV. RESULTS AND DISCUSSION

## A. Model Evaluation

We adapt three DL architectures for comparison with MAIAC's SR and cloud retrieval lookup tables (see Fig. 3). We compare the three models to identify the best-performing architecture based on the error in SR prediction, accuracy in cloud identification, and uncertainty quantification. Most in-depth analyses are focused on DCVDSR, which is the best-performing model in uncertainty calibration.



Fig. 3. Correction of atmospheric effects and cloud masking performed on true color images from the Japanese geostationary satellite Himawari-8. (a) TOA reflectance before processing (**Input**). (b) Retrieval of SR and cloud mask (in white) produced by the surrogate MAIAC model (**Target**: MAIAC lookup table). (c) Prediction of SR and cloud mask by the Bayesian DL emulator (**Prediction**: DL emulator). (d) Visualization of predictive undertainty from the emulator model (**Uncertainty**). **Rows from top:** southwestern Australia, eastern China, and the Bohai Sea.

		Blue	Green	Red	NIR	SWIR1	SWIR2
$\overline{SR}$	MAIAC	0.064	0.084	0.155	0.307	0.327	0.231
$\overline{SR}_{MAIAC} - \overline{SR}_{emulator}$	DCFC	-0.001	<b>-0.003</b>	-0.015	-0.012	-0.012	-0.031
	DCCNN	0.002	0.006	<b>0.008</b>	0.012	0.014	<b>0.006</b>
	DCVDSR	-0.004	-0.009	0.010	<b>0.005</b>	<b>0.002</b>	0.008
CV (%)	MAIAC	88	80	62	34	38	51
$\frac{CV_{MAIAC} - CV_{emulator}}{CV_{MAIAC}} (\%)$	DCFC	<b>9.1</b>	<b>4.2</b>	<b>2.3</b>	6.8	3.7	<b>1.0</b>
	DCCNN	31	24	10	<b>1.7</b>	<b>-0.84</b>	4.1
	DCVDSR	30	11	12	11	3.4	4.8
Correlation coefficient (r)	DCFC	0.960	<b>0.962</b>	0.842	0.893	<b>0.984</b>	<b>0.963</b>
	DCCNN	0.948	0.950	0.818	0.866	0.977	0.960
	DCVDSR	0.792	0.860	<b>0.930</b>	<b>0.918</b>	0.977	0.956
Conditional RMSE	DCFC	0.0221	0.0217	0.0235	<b>0.0196</b>	0.0186	0.0387
	DCCNN	0.0228	0.0239	0.0228	0.0228	0.0228	0.0235
	DCVDSR	0.0211	<b>0.0210</b>	<b>0.0220</b>	0.0211	<b>0.0153</b>	0 <b>.0233</b>

 TABLE II

 Evaluation of SR From the Three Candidate Emulator Models

1) Surface Reflectance: A comparison between basic statistics of SR datasets obtained from MAIAC and the emulators is presented in Table II. Here, the mean value of SR is the average intensity of clear sky pixels in each band. Coefficient of variation (CV) is a measure of relative dispersion of the data calculated as the ratio between the standard deviation and mean of a distribution. For SR, CV relates to the radiometric stability characteristic of the sensor, with lower CV indicating greater stability. A comparison of MAIAC and emulator CV suggests that the predictions of the fully connected model (DCFC) most closely match the dispersion of MAIAC SR. The emulator models generally capture the relative magnitudes of variation in each wavelength while underestimating the variation of the SR distribution. Underestimation of the observed variation is most pronounced for the blue band across models.

Histograms of differences between MAIAC and emulator SR are plotted in Fig. 4. For an ideal model, differences between observed and modeled values should be small and unbiased. Distributions are generally symmetric and centered around near-zero means, indicating a minimal bias toward overestimation or underestimation by the emulator. Similar results are obtained for the three candidate emulator models.

The correlation coefficient and the conditional root mean square error (RMSE) are also presented in Table II. The



Fig. 4. Histograms of difference between MAIAC lookup table SR and emulator (DCVDSR) SR for six bands indicate minimal bias toward overestimation or underestimation of SR by the emulator.



Surface Reflectance (MAIAC)

Fig. 5. Density scatter plots compare SR from the MAIAC lookup table and the DL emulator (DCVDSR) for the six AHI solar reflective bands. Red indicates a higher density of points and blue a lower density. The gray line is a 1:1 line.

conditional RMSE refers to RMSE evaluated at the pixels identified by both MAIAC and the emulator as clear sky. Correlation coefficients indicate the strongest linear relationship between MAIAC SR and DCFC emulator SR. This result is unsurprising as the convolutional filters in DCCNN and DCVDSR will produce smoother results than the pixelwise model, DCFC. These correlation coefficients reflect similar strength of the linear relationship as those between MAIAC SR and MODIS SR [28]. MODIS Terra and Aqua products provided the principle comparison for validation of AHI MAIAC SR. The evaluation of the performance across models suggests that mappings in some bands may be easier to learn (SWIR1 and SWIR2) and others more difficult to learn (blue and green). Conditional RMSE is generally the lowest for the DCVDSR model. As the square root of the variance of residuals, RMSE indicates the absolute fit of the model to the data and can be thought of as more germane to predictive ability than correlation. In the context of the accuracy requirements of higher level data products (Albedo: 0.02–0.05; NDVI: 0.03) [44], the emulators generally err within the reasonable range.

The pixel-by-pixel comparison is presented using density plots in Fig. 5. The plots suggest strong coherence between MAIAC and emulator SR. A 1:1 line is displayed for visual



Fig. 6. Analysis of cloud prediction by the DL emulators. (a) Cloud classification accuracy with varying decisions thresholds under the discrete-continuous distribution. Models achieve 86%–87% accuracy compared to MAIAC lookup table cloud masks. Dotted lines represent the optimal threshold for each model. (b) Area under the ROC curve for evaluation of the model's discrimination ability with an AUC value of 1 denoting perfect discrimination and the gray 1:1 line representing no skill.

comparison, while the slope and intercept of the data best fit line are displayed on the plots. Outliers are generally located above the 1:1 line, suggesting that the emulator model does not fully capture the upper tail of the SR distribution.

2) Cloud Identification: The evaluation of the emulator cloud prediction is performed by pixelwise comparison between the MAIAC cloud mask and the emulator cloud mask. Binary classification accuracy is defined as the fraction of true predictions, where TP = true positives, TN = true negatives, FP = false positives, and FN = false negatives

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}.$$
 (8)

As described in Section III-A2, the model is conditioned to predict  $\hat{p}$  as the probability of a pixel being a clear sky. By selecting a decision threshold value of **p**, cloud classification proceeds by casting pixels with  $\hat{p} < \mathbf{p}$  as nonclear sky and  $\hat{p} > \mathbf{p}$  as clear sky. Continuously varying the decision threshold **p** and calculating the resulting classification accuracy indicate the optimal mask probability, **p**, for each trained model. Fig. 6 presents a plot of classification accuracy with varying decision thresholds.

Fig. 6 also presents the receiver operating characteristic (ROC) curve, used to assess the discrimination ability of binary classifiers. The true positive rate (TPR) is plotted against the false positive rate (FPR) at various thresholds. The area under the curve (AUC) provides a threshold-invariant measure of how well the model can discriminate between two classes, with a maximum value of one for perfect classification.

The performance by accuracy for cloud classification is similar across the evaluated emulator models (see Table III).

TABLE III CLOUD CLASSIFICATION ACCURACY, SENSITIVITY, AND SPECIFICITY

	Accuracy (%)	Sensitivity	Specificity
DCFC	0.866	0.702	0.921
DCCNN	0.877	0.767	0.905
DCVDSR	0.864	0.771	0.872

Sensitivity refers to the TPR or proportion of clear sky pixels that are correctly classified. Specificity refers to the true negative rate or the proportion of nonclear pixels that are correctly classified. A high specificity classifier will erroneously screen high aerosol optical depth pixels, while a less conservative, higher sensitivity classifier carries more chance of cloud contamination. Such cloud contamination has a potentially strong negative effect on SR retrieval. The three emulator models are generally conservative, achieving greater classification accuracy for nonclear pixels than clear sky pixels.

It should be noted that the assessment of classification accuracy uses MAIAC cloud masks as ground truth. Cloud masks produced from MAIAC contain uncertainties and inaccuracies of their own, and it is possible that the ability of CNNs to incorporate spatial information produces an advantage in cloud classification. Visual assessment of cloud predictions from MAIAC and emulator often indicate greater spatial coherence of emulator cloud masks and lesser appearance of some undesirable model artifacts (see Fig. 3).

*3) Stability of Model Over Varied Conditions:* Homogeneous vegetation areas are identified using MODIS MCD12Q1 Land Cover Type I, and the performance of the MAIAC



Fig. 7. Analysis of emulator SR error and cloud classification accuracy across varying surface conditions. (a) Emulator (DCVDSR) performance is relatively stable across vegetated land types according to the IGBP global vegetation classification scheme. (b) Emulator (DCVDSR) performance in the northern hemisphere is slightly poorer in winter than in seasons. Dotted lines represent overall mean performance for all land covers and all seasons in the northern hemisphere, respectively.

emulator is evaluated for each land cover type separately. Performances including conditional RMSE of SR and cloud classification accuracy by the DCVDSR emulator are presented in Fig. 7. Results are presented for the nine most abundant classes in the test set. Both regression error and cloud classification accuracy are relatively stable for vegetated categories, including forests, shrubland, and savanna, while barren land results in poorer performance. The optical properties of highly reflective surfaces present a challenge to atmospheric correction, and such may also result in poor performance for MAIAC [45]. In addition to high reflectance, large solar angle, large zenith angle, and high aerosol optical depth also form challenging conditions for reflectance retrieval. More specific evaluations are needed to understand the performance of the emulator under these special conditions.

The seasonal analysis is used to evaluate the performance under annual fluctuations in vegetation phenology (see Fig. 7). Spring green-up, fall senescence, and transitions between wet and dry seasons result in SR variation of several absolute percent in vegetated areas [30]. SR error and cloud classification accuracy are generally stable throughout the year but evince slightly poorer performance and greater spread fall months (SR prediction) and winter months (cloud classification). Seasonal performance is evaluated separately for each hemisphere for consistency of seasons. Similar results were found in the Southern Hemisphere.

#### B. Uncertainty Quantification

Bayesian DL models capture predictive uncertainty in regression tasks by producing a probabilistic output. As described in Section III-A1, we use variational inference to produce an ensemble of predictions for each sample and then compute unbiased estimates of the first and second moments of the predictive distribution at each pixel. From the second moment, the standard deviation expresses the magnitude of predictive uncertainty at each location.

We assess the quality of the uncertainty measurements by evaluating the uncertainty calibration or whether the model captures the uncertainty in observed data. We compare the



Fig. 8. Assessment of predictive uncertainty across the three Bayesian DL emulator models. Uncertainty calibration, which evaluates the frequency of observed values (*y*-axis) within predicted probability ranges (*x*-axis), indicates the underestimation of uncertainty to varying extents. The gray 1:1 line indicates the behavior of a perfectly calibrated model.

model's predictive distribution to the observed values by evaluating the frequency of residuals lying in various probability thresholds within the predicted distribution [32] (see Fig. 8). A perfectly calibrated model, which captures the distribution of the observed data, would match the 1:1 line. All three models underestimate uncertainty to some extent, meaning that they are overconfident in their predictions. Of the three, DCVDSR has the most well-calibrated uncertainty.

While quantification of uncertainty in DL makes a step toward more informed use of artificial intelligence models, interpretability remains a key challenge in Earth science and other domains. Alternatives to dropout, such as Bayes by Backprop, have exhibited comparable performance in recent years [46]. Further theoretical development and experimentation with practical applications will improve the trustworthiness of AI.

#### C. Computational Performance

The spatiotemporal resolution and spatial extent of AHI scans result in the generation of 94 GB of new data per day. There is growing interest in using AI to exploit this data, as well as the growing volume from other remote sensors, climate reanalysis models, and Earth system models.

We quantify the nontrivial computing time necessary to retrieve SR using the MAIAC lookup table and the emulator models. To evaluate the DL emulator models, we assess inference from one forward pass (static network) and ten stochastic forward passes (Bayesian sampling network). A single inference with the static network is sufficient to produce SR and cloud products; Bayesian sampling produces the same with uncertainty quantification. Processing speeds are presented in Table IV. Emulator inference is evaluated on one GPU, while

TABLE IV PROCESSING SPEED OF MAIAC AND EMULATOR MODELS

	Examples per second			
Model	Static network	Bayesian sampling network		
MAIAC LUT	0.40			
DCFC	2.4	0.60		
DCCNN	1.8	0.33		
DCVDSR	1.5	0.25		

the MAIAC lookup tables are executed on one CPU. Among the compared emulator models, processing speed decreases with increasing complexity. Inference with Bayesian sampling is generally comparable in speed to the LUT, while inference on the static network represents between 3.75x (DCDVSR) and 6x (DCFC) speedup over the acceleration offered by the LUT.

GPU-based models also have advantages in scalability. Parallel workloads distributed across GPUs can achieve a speedup, which is nearly linear with the number of cores, while scaling across additional CPU cores lags in performance. These behaviors will tend to increase the advantages of GPU-based computing in higher throughput operations. In addition, advances in distributed computing and performance-aware implementations of DL, while outside the scope of this work, will further enhance efficiency. Meanwhile, improvements in the implementation of MAIAC and the design of lookup tables provide a moving target for both statistical and computational performance.

## V. CONCLUSION

The time and expense associated with running numerical models often create a bottleneck in research on complex systems. In this work, we present a case study evaluating DL-based emulation in comparison to a more traditional reduced-order approximation of the MAIAC model. We found initially promising results in the statistical and computational performance of the DL emulator. However, in counterbalance to DL's unfettered success in many commercial domains, we encountered challenges unique to working with geoscience data. Satellite observations exhibit significant heterogeneity in space and time with the joint distributions between variables fluctuating across locations and seasons. We mitigated this effect by using smaller patches to allow for larger batch sizes with more heterogeneity of samples. Whereas any form of model reduction represents a tradeoff between model fidelity and speed, the "black box" nature of DL models is a particular disadvantage in the geosciences, where physical consistency and interpretability are key. Appropriate adaptations of recent work in explainable and interpretable DL [47], as well as physics-guided neural networks [48], may offer greater reliability and new research directions in the context of emulation. For example, networks that trace salience back to the input information can be used to infer scientifically meaningful information about what the model has learned. Future directions should also include further evaluation of the model's

ability to generalize, especially under rare or challenging conditions.

As computation time is the main motivating factor in this work, it is important to consider the multiple factors contributing to the outlook on emulations' computational advantages. Just as high-performance implementations of DL and GPU hardware are active areas of research, MAIAC and its reduced order implementations are not static but are under ongoing development. Advances in numerical modeling and CPU-based computing will alter the landscape across many scientific disciplines and shift the benchmark for speedups. As such, we anticipate a framework for DL emulation to evolve continuously with new generations of hardware and software technology. Our current results suggest that, as advances continue in machine learning and highperformance computing, the credibility and computational performance of DL emulators may improve sufficiently to ease computing limitations on model development and big data processing across multiple scientific disciplines.

#### References

- F. O. Ajibola, B. Zhou, G. T. Gnitou, and A. Onyejuruwa, "Evaluation of the performance of CMIP6 HighResMIP on west African precipitation," *Atmosphere*, vol. 11, no. 10, p. 1053, Oct. 2020.
- [2] N. S. Diffenbaugh, J. S. Pal, R. J. Trapp, and F. Giorgi, "Fine-scale processes regulate the response of extreme events to global climate change," *Proc. Nat. Acad. Sci. USA*, vol. 102, no. 44, pp. 15774–15778, 2005.
- [3] D. Kumar, E. Kodra, and A. R. Ganguly, "Regional and seasonal intercomparison of CMIP3 and CMIP5 climate model ensembles for temperature and precipitation," *Climate Dyn.*, vol. 43, nos. 9–10, pp. 2491–2518, Nov. 2014.
- [4] R. Séférian *et al.*, "Tracking improvement in simulated marine biogeochemistry between cmip5 and cmip6," *Current Climate Change Rep.*, vol. 6, no. 3, pp. 1–25, 2020.
- [5] S. McKenna, A. Santoso, A. S. Gupta, A. S. Taschetto, and W. Cai, "Indian ocean dipole in CMIP5 and CMIP6: Characteristics, biases, and links to ENSO," *Sci. Rep.*, vol. 10, no. 1, pp. 1–13, Dec. 2020.
- [6] J. T. Overpeck, G. A. Meehl, S. Bony, and D. R. Easterling, "Climate data challenges in the 21st century," *Science*, vol. 331, no. 6018, pp. 700–702, 2011.
- [7] D. Ghent, J. Kaduk, J. Remedios, and H. Balzter, "Data assimilation into land surface models: The implications for climate feedbacks," *Int. J. Remote Sens.*, vol. 32, no. 3, pp. 617–632, Feb. 2011.
- [8] S. Saha et al., "The NCEP climate forecast system reanalysis," Bull. Amer. Meteorol. Soc., vol. 91, no. 8, pp. 1015–1058, 2010.
- [9] A. Y. Sun and B. R. Scanlon, "How can big data and machine learning benefit environment and water management: A survey of methods, applications, and future directions," *Environ. Res. Lett.*, vol. 14, no. 7, Jul. 2019, Art. no. 073001.
- [10] 2017: The Year of Goes-16. Accessed: Sep. 21, 2021.
   [Online]. Available: https://www.goes-r.gov/featureStories/2017T heYearOfGOES16.html
- [11] S. Castruccio, D. J. McInerney, M. L. Stein, F. L. Crouch, R. L. Jacob, and J. E Moyer, "Statistical emulation of climate model projections based on precomputed gcm runs," *J. Climate*, vol. 27, no. 5, pp. 1829–1844, 2014.
- [12] P. B. Holden and N. R. Edwards, "Dimensionally reduced emulation of an AOGCM for application to integrated assessment modelling," *Geophys. Res. Lett.*, vol. 37, no. 21, pp. 1–21, Nov. 2010.
- [13] V. M. Krasnopolsky, M. S. Fox-Rabinovitz, and D. V. Chalikov, "New approach to calculation of atmospheric model physics: Accurate and fast neural network emulation of longwave radiation in a climate model," *Monthly Weather Rev.*, vol. 133, no. 5, pp. 1370–1383, May 2005.
- [14] J. P. Rivera, J. Verrelst, J. Gómez-Dans, J. Muñoz-Marí, J. Moreno, and G. Camps-Valls, "An emulator toolbox to approximate radiative transfer models with statistical learning," *Remote Sens.*, vol. 7, no. 7, pp. 9347–9370, 2015.

- [15] A. Lyapustin, J. Martonchik, Y. Wang, I. Laszlo, and S. Korkin, "Multiangle implementation of atmospheric correction (MAIAC): 1. Radiative transfer basis and look-up tables," *J. Geophys. Res.*, vol. 116, no. D3, pp. 1–9, 2011.
- [16] M. F. Kasim et al., "Building high accuracy emulators for scientific simulations with deep neural architecture search," 2020, arXiv:2001.08055.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 2, pp. 84–90, Jun. 2017.
- [18] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *J. Mach. Learn. Res.*, vol. 12 pp. 2493–2537, Aug. 2011.
- [19] E. Gibney, "Google AI algorithm masters ancient game of go," *Nature News*, vol. 529, no. 7587, p. 445, 2016.
- [20] Y.-G. Ham, J.-H. Kim, and J.-J. Luo, "Deep learning for multiyear ENSO forecasts," *Nature*, vol. 573, no. 7775, pp. 568–572, Sep. 2019.
- [21] T. Vandal, E. Kodra, S. Ganguly, A. Michaelis, R. Nemani, and A. R. Ganguly, "DeepSD: Generating high resolution climate change projections through single image super-resolution," in *Proc. 23rd* ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, Aug. 2017, pp. 1663–1672.
- pp. 1663–1672.
  [22] Y. Liu, A. R. Ganguly, and J. Dy, "Climate downscaling using YNet: A deep convolutional network with skip connections and fusion," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2020, pp. 3145–3153.
- [23] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and M. Prabhat, "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, pp. 195–204, Feb. 2019.
- [24] T. J. Sejnowski, "The unreasonable effectiveness of deep learning in artificial intelligence," *Proc. Nat. Acad. Sci. USA*, vol. 117, no. 48, pp. 30033–30038, Dec. 2020.
- [25] F. Kratzert, D. Klotz, M. Herrnegger, A. K. Sampson, S. Hochreiter, and G. S. Nearing, "Toward improved predictions in ungauged basins: Exploiting the power of machine learning," *Water Resour. Res.*, vol. 55, no. 12, pp. 11344–11354, Dec. 2019.
- [26] R. Lagerquist, D. Turner, I. Ebert-Uphoff, J. Stewart, and V. Hagerty, "Using deep learning to emulate and accelerate a radiative-transfer model," *J. Atmos. Ocean. Technol.*, pp. 1673–1696, vol. 38, no. 10, Jul. 2021.
- [27] J. M. Agency, "Himawari-8/9 Himawari standard data user's guide version 1.2," Japan Meteorological Agency Otemachi, Chiyoda-ku, Tokyo, Japan, Tech. Rep. 100-8122, 2015.
- [28] S. Li et al., "First provisional land surface reflectance product from geostationary satellite Himawari-8 AHI," *Remote Sens.*, vol. 11, no. 24, p. 2990, Dec. 2019.
- [29] A. Lyapustin *et al.*, "Multiangle implementation of atmospheric correction (MAIAC): 2. Aerosol algorithm," *J. Geophys. Res.*, vol. 116, no. D3, pp. 1–15, 2011.
- [30] A. I. Lyapustin *et al.*, "Multi-angle implementation of atmospheric correction for MODIS (MAIAC): 3. atmospheric correction," *Remote Sens. Environ.*, vol. 127, pp. 385–393, Dec. 2012.
- [31] D. Sulla-Menashe and M. A. Friedl, "User guide to collection 6 MODIS land cover (MCD12Q1 and MCD12C1) product," USGS, Reston, VA, USA, Tech. Rep. 1, 2018.
- [32] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5574–5584.
- [33] A. Graves, "Practical variational inference for neural networks," in Proc. Adv. Neural Inf. Process. Syst., 2011, pp. 2348–2356.
- [34] H. Wang and D.-Y. Yeung, "A survey on Bayesian deep learning," 2016, arXiv:1604.01662.
- [35] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014.
- [36] Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in Proc. Adv. Neural Inf. Process. Syst., 2017, pp. 3581–3590.
- [37] Y. Gal, "Uncertainty in deep learning," Ph.D. dissertation, Dept. Eng., Univ. Cambridge, Cambridge, U.K., 2016.
- [38] T. Vandal, E. Kodra, J. Dy, S. Ganguly, R. Nemani, and A. R. Ganguly, "Quantifying uncertainty in discrete-continuous and skewed data with Bayesian deep learning," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 2377–2386.
- [39] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (CVPR), Jun. 2016, pp. 770–778.

- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014. arXiv:1412.6980.
- [41] S.-S. Li, L.-F. Chen, J.-H. Tao, D. Han, Z.-T. Wang, and B.-H. He, "Retrieval and validation of the surface reflectance using HJ-1-CCD data," Guang Pu Xue Yu Guang Pu Fen Xi= Guang Pu, vol. 31, no. 2, pp. 516–520, 2011.
- [42] H. Guo, X. Gu, F. Bao, and S. Shi, "Analysis of surface reflectance retrieval over four typical surfaces via gaofen-1 satellite WFV4 imagery," J. Indian Soc. Remote Sens., vol. 48, no. 5, pp. 1-12, 2020.
- [43] M. Feng, C. Huang, S. Channan, E. F. Vermote, J. G. Masek, and J. R. Townshend, "Quality assessment of landsat surface reflectance products using MODIS data," Comput. Geosci., vol. 38, no. 1, pp. 9-22, Jan. 2012.
- [44] E. Vermote et al., "Surface reflectance earth system data record/climate data record white paper," NASA, Greenbelt, MD, USA, Tech. Rep. 18, 2006. vol. 6.
- [45] A. Lyapustin, Y. Wang, I. Laszlo, and S. Korkin, "Improved cloud and snow screening in MAIAC aerosol retrievals using spectral and spatial analysis," Atmos. Meas. Techn., vol. 5, no. 4, pp. 843-850, Apr. 2012.
- [46] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in Proc. Int. Conf. Mach. Learn., 2015, pp. 1613–1622.
- [47] B. A. Toms, E. A. Barnes, and I. Ebert-Uphoff, "Physically interpretable neural networks for the geosciences: Applications to earth system variability," 2019, arXiv:1912.01752.
- [48] A. Daw, A. Karpatne, W. Watkins, J. Read, and V. Kumar, "Physicsguided neural networks (PGNN): An application in lake temperature modeling," 2017, arXiv:1710.11431.









Thomas J. Vandal received the B.S. degree in mathematics from the University of Maryland, College Park, MD, USA, in 2012, and the Ph.D. degree in interdisciplinary engineering from Northeastern University, Boston, MA, USA, in 2018.

He is currently a Research Scientist with the NASA Ames Research Center, Mountain View, CA, USA, and the Bay Area Environmental Research Institute, Moffett Field, CA, USA. His research interests include machine learning and Earth sciences.

Weile Wang received the B.S. and M.S. degrees in electrical engineering from Tsinghua University, Beijing, China, in 1998 and 2001, respectively, and the Ph.D. degree in physical geography from Boston University, Boston, MA, USA, in 2006.

He is currently the Research Scientist with the NASA Ames Research Center, Moffett Field, CA, USA, and the California State University at Monterey Bay, Marina, CA, USA. His research interests include remote sensing, global carbon cycle, and climate change.

Ramakrishna R. Nemani received the B.S. degree from Andhra Pradesh Agricultural University, Bapatla, India, in 1979, the M.S. degree from Punjab Agricultural University, Ludhiana, India, in 1982, and the Ph.D. degree from the University of Montana, Missoula, MT, USA, in 1987.

He is currently a Senior Earth Scientist with the Ecological Forecasting Laboratory, NASA Ames Research Center, Moffett Field, CA, USA. His research interests include ecological forecasting and collaborative computing in the Earth sciences.

Auroop R. Ganguly received the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002.

He has 24 years of professional experience in the U.S. spanning the private sector, government laboratories, and academia. He is currently a Professor with Northeastern University, Boston, MA, USA, and a Chief Scientist by joint affiliation with the U.S. DOE's Pacific Northwest National Laboratory, Richland, WA, USA.

Dr. Ganguly is also a fellow of the American Society of Civil Engineers (ASCE) and a Senior Member of the Association for Computing Machinery (ACM).



Kate Duffy received the B.Sc. degree in environmental engineering from Cornell University, Ithaca, NY, USA, in 2015, and the M.S. degree in civil engineering and the Ph.D. degree in interdisciplinary engineering from Northeastern University, Boston, MA, USA, in 2019 and 2021, respectively.

She is currently a Research Scientist with the NASA Ames Research Center, Mountain View, CA, USA, and the Bay Area Environmental Research Institute, Moffett Field, CA, USA. Her research interests include climate change, remote sensing, and applied machine learning.

