

Spectral Synthesis for Geostationary Satellite-to-Satellite Translation

Thomas J. Vandal^{id}, Daniel McDuff^{id}, *Member, IEEE*, Weile Wang, Kate Duffy^{id},
Andrew Michaelis, and Ramakrishna R. Nemani^{id}

Abstract—Earth-observing satellites carrying multispectral sensors are widely used to monitor the physical and biological states of the atmosphere, land, and oceans. These satellites have different vantage points above the Earth and different spectral imaging bands resulting in inconsistent imagery from one to another. This presents challenges in building downstream applications. What if we could generate synthetic bands for existing satellites from the union of all domains? We tackle the problem of generating synthetic spectral imagery for multispectral sensors as an unsupervised image-to-image translation problem modeled with a variational autoencoder (VAE) and generative adversarial network (GAN) architecture. Our approach introduces a novel shared spectral reconstruction loss to constrain the high-dimensional feature space of multispectral images. Simulated experiments performed by dropping one or more spectral bands show that cross-domain reconstruction outperforms measurements obtained from a second vantage point. Our proposed approach enables the synchronization of multispectral data and provides a basis for more homogeneous remote sensing datasets.

Index Terms—Geophysical image processing, neural networks (NNs), remote sensing, unsupervised learning.

I. INTRODUCTION

CLIMATE change and related environmental issues—including the loss of biodiversity and extreme weather—are listed by the World Economic Forum as the most important risks to our planet [1]. Monitoring the Earth is critical to mitigating these risks, understanding the effects, and making future predictions [2]. Multispectral and hyperspectral satellite-based remote sensing enables global observation of the Earth, allowing scientists to study large-scale system dynamics and inform general circulation models [3]. In weather forecasts, satellite data initialize the atmospheric state for future predictions. On longer time

scales, these data are used to measure the effects of climate change, such as land-cover variations, temperature trends, solar radiation levels, and the rate of snow/ice melt. In the coming decades, increased investments from the public and private sectors in satellite-based observations will continue to improve global monitoring, as highlighted in NASA’s decadal survey [4].

Satellites are designed based on specifications for a given set of applications with fiscal, technological, and physical constraints that limit their temporal, spatial, and spectral resolutions. Geostationary (GEO) satellites rotate with the Earth to stay over a constant position above the equator at a high elevation of 35786 km. This position enables GEO satellites with onboard multispectral imagers to take continuous and high-temporal snapshots over large spatial regions and is ideal for monitoring diurnal and fast moving events. Spectral bands measure the brightness and radiance intensities of the electromagnetic spectrum at a specified center wavelength and bandwidth. Bands are selected to satisfy defined variables of interest constrained by technological cost and accuracy. Applications of GEO sensors include atmospheric winds measurement [5], tropical cyclone tracking [6], wildfire monitoring [7], and short-term forecasting [8]. Multiple GEO satellites are needed to generate global high-temporal resolution datasets to better monitor these events around the world. However, variations in resolutions, sensor uncertainties, and temporal life spans lead to a set of separate datasets that are not consistent, making this process very challenging [3]. Developing consistent and homogeneous global datasets would relieve many of these challenges.

The current generation of GEO satellites (shown in Fig. 1) is now exception. The GOES-16/17 satellites operated by NASA/NOAA (cost: \$11 billion) have a set of 16 imaging bands covering the visible, near-, and thermal-infrared spectral range [9]. The Himawari-8 satellite operated by the Japanese Space Agency (cost: \$800 million) similarly has 16 bands but swaps an NIR (1.38 μm) band for a green channel (0.51 μm), enabling the construction of true-color images [10]. The 1.38- μm band is ideal for measuring Cirrus clouds, composed of ice particles in the upper troposphere, a major contributor to regulating the Earth’s climate that is not yet well understood [11], [12]. Without capturing this band, directly observing Cirrus clouds over Japan, East Asia, and Western Pacific region from Himawari-8 is not possible. Synthetic observations via virtual spectral sensors could be a low-cost solution to improving coverage availability and consistency with current satellites.

Manuscript received January 29, 2021; revised April 16, 2021 and May 27, 2021; accepted May 29, 2021. Date of publication June 23, 2021; date of current version January 12, 2022. This work was supported in part by the NASA Earth eXchange (NEX) and in part by the NASA Research Opportunities in Space and Earth Science (ROSES) under Grant 19-ESROGSS-19-0123. (Corresponding author: Thomas J. Vandal.)

Thomas J. Vandal is with the NASA Ames Research Center, Moffett Field, CA 94035 USA, and also with the Bay Area Environmental Research Institute, Moffett Field, CA 94035 USA (e-mail: vandall@baeri.org).

Daniel McDuff is with Microsoft AI, Redmond, WA 98052 USA.

Weile Wang is with the NASA Ames Research Center, Moffett Field, CA 94035 USA, and also with the School of Natural Sciences, California State University at Monterey Bay, Marina, CA 93933 USA.

Kate Duffy is with the NASA Ames Research Center, Moffett Field, CA 94035 USA, and also with the Civil and Environmental Engineering Department, Northeastern University, Boston, MA 02115 USA.

Andrew Michaelis and Ramakrishna R. Nemani are with the NASA Ames Research Center, Moffett Field, CA 94035 USA.

Digital Object Identifier 10.1109/TGRS.2021.3088686

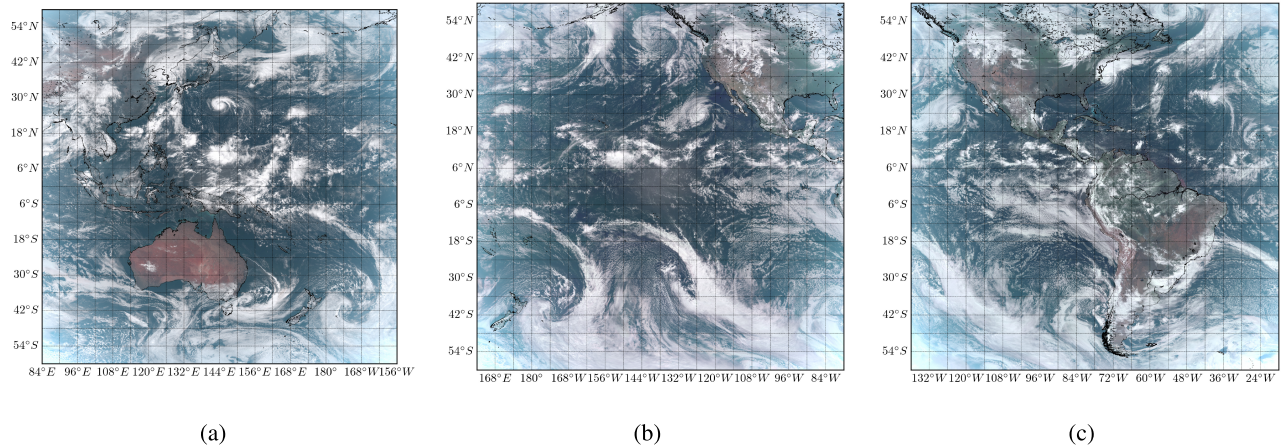


Fig. 1. Full-disk GEO true-color images from (a) Himawari-8 AHI, (b) GOES-17 ABI, and (c) GOES-16 ABI. Green is synthetically generated using the proposed method for ABI (b) and (c).

We present an approach to generate synthetic spectral channels from a multidomain unpaired satellite dataset. We treat satellites with either dissimilar spectral coverage or varying vantage points as separate spectral sets. In this way, the problem closely resembles that of colorization [13] and image-to-image translation tasks [14]–[16] in the case where paired images are not available but with the added complexity of a large number of spectral bands. We use a combination of variational autoencoder (VAE) and generative adversarial network (GAN) [17] architectures adapted to our problem to model a shared latent space, as in unsupervised image-to-image translation [14]. Generating synthetic bands is an underconstrained problem that paired with an adversarial loss in high dimensions, which promotes overfitting. Our approach mitigates these challenges by leveraging a weak supervision signal based on partial overlap in spectral bands between domains. By including a reconstruction loss on overlapping spectral bands between domain pairs, we can substantially improve spectral band synthesis.

To summarize our contributions, we: 1) introduce a shared spectral reconstruction loss to a VAE-GAN architecture for synthetic band generation; 2) test our methodology on real-world scenarios; and 3) present and release a test dataset of 2000 tiles of paired observations from the GeoNEX LIG GEO imagery for future research. In Sections II–IV, we will introduce related work in remote sensing and image-to-image translation, describe the architecture, and review experiments. Finally Section V, we will discuss the implications on this work and conclude with future directions.

II. BACKGROUND

A. Remote Sensing

Current generation GEO satellites observe 16 spectral bands over large regions every 10–15 min at a 0.5–2-km resolution. At a suboptimal 2 km, this produces full-disk images of size $5424 \times 5424 \times 16$, which causes storage constraints while being computationally expensive to process. Physical and statistical models are used to convert these images into more easily

interpreted variables, such as precipitation, cloud cover, and surface temperature [18]. Multiple GEO satellites, currently in orbit, extend the spatial ranges to actively monitoring larger regions. However, differences in spectral bands and sensor uncertainties/biases present challenges to commonly used sensor-specific models, and especially, existing downstream models do not generalize well to missing spectral information.

Spectral band adjustment is often applied to cross calibration of sensors using relative spectral responses with a hyperspectral sensor, such as Hyperion [19], [20]. This approach uses Hyperion to calculate spectral band adjustment factors (SBAFs) from relative spectral responses and has been applied to a number of datasets. For instance, SBAF was applied to Harmonized Landsat and Sentinel-2 for accurate cross comparison with the MODIS dataset [21]. Similarly, the work [20] used SBAF to evaluate long-term AVHRR surface reflectance datasets with a quadratic normalized difference vegetation index (NDVI). However, as outlined in these studies, applying SBAF requires an intermediate simulated or hyperspectral sensor to perform this translation. In contrast, our approach learns this intermediate sensor as a “shared” latent space capturing dependent information content between sensors. The following sections will develop our approach in the context of neural networks (NNs) and deep generative modeling.

Neural models have long been applied to process remote sensing data and generate downstream products. Hsu *et al.* [22] presented some of the first work that showed NNs could generate accurate and high-resolution precipitation products from satellite observations. In recent years, convolutional neural networks (CNNs) have been found to further improve this task [23]. Similarly, CNNs have successfully been applied to poverty mapping [24], super-resolution [25], sub-pixel classification [26], model emulation [27], and land-cover classification [28], all from low-level satellite products. In terms of spectral synthesis, a few studies have explored reconstruction of hyperspectral bands from RGB bands with supervised approaches [29], [30]. While many of these problems are within the class of image-to-image translation, they

generally assume that labels are widely available and focus on individual sensors. To the best of our knowledge, no studies have developed approaches to synthesize spectral information by learning across satellites in the unsupervised setting.

B. Image-to-Image Translation

Many problems can be defined as an image-to-image translation task, including super-resolution, style transfer, and colorization. Approaches to image-to-image translation have been developed for both supervised and unsupervised settings to map images from one domain to another. In the supervised setting, image pairs are available to learn a direct mapping from one to the other. GANs have been shown to be highly successful at this task [31], [32]. Numerous unsupervised learning methods have been developed for the common case of large unpaired datasets [14], [15], [33], [34]. CycleGAN, for instance, proposed an approach to directly map from one domain to another and back by incorporating a cycle-consistency loss with a GAN [15]. UNIT [14] proposed a probabilistic approach that uses an intermediate latent space between domains with a VAE [35] and GANs [17]. In contrast to prior work on image-to-image translation, our scenario specifically requires spectral translation and across multiple domains. Rather than translating between relatively low-dimensional RGB images and segmentation maps, as is found in traditional multimodal image-to-image translation [36]–[38], satellite imagery contains tens to hundreds of spectral bands. Domain adaptation is another area of active research, which also considers the case of effectiveness in unseen environments with cycle consistency and domain invariant [39], [40]. Lee *et al.* [41] used a shared content loss to translate between RGB image styles. Sanchez *et al.* [42] presented an application of image-to-image translation for four-band Sentinel-4 images between different times of day. Using VAEs and GANs, their approach provided 94.5% accuracy in classifying ten land-cover types in the EuroSat dataset. Our approach is based on the proven fundamental techniques of learning a shared latent space using cycle-consistency and adversarial losses extended in the spectral dimension. We also use the prior understanding of spatial consistency between domains to implement a partial skip connection.

C. Variational Autoencoders

Autoencoders (AEs) are widely used in deep learning to encode high-dimensional data to a lower dimension representation [43]. AEs consist of two stages, encoding and decoding networks, and are learned in an unsupervised manner. The encoder network $E(x)$ takes an input x to a low-dimensional representation and inputs a decoder, written as $G(E(x))$. AEs can then be trained with a mean square error loss to reconstruct the input, $\|x - G(E(x))\|^2$. While AEs are well suited for compressing a high-dimensional feature space, there is little control on the distribution of the latent space and can lead to severe overfitting.

In contrast to AEs, VAEs are generative models, which aims to learn an intermediate “latent” variable z as a compressed representation of the input. This is done by first encoding the

data to z followed by a sampling operation and decoding to the original domain [35]. More formally, VAEs model the latent space as a probability distribution such that the likelihood is written as $p(x|z) = p(z|x) * p(z)$. The prior distribution $p(z)$ is generally assumed to be Gaussian where $p(z) = N(0, I)$. The posterior distribution $p(z|x)$ is then approximated as $q(z|x) = \mathcal{N}(E(x), c * I)$, where c is a constant. Conditioning on the latent space z to reconstruct x is then written as $p(x|z) = G(z)$. The loss of a VAE is then written as $\|x - G(z)\|^2 - \text{KL}(q(z|x)||p(z))$ where KL is the Kullback–Leibler divergence measuring the relative entropy between two probability distribution. KL is formally defined as

$$\text{KL}(q(z|x)||p(z)) = \sum_x p(z) \frac{p(z)}{q(z|x)}. \quad (1)$$

With this approach, we are able to generate previously unseen examples from a constrained probabilistic latent space. Previous examples of VAEs in remote sensing include classification [44], hyperspectral unmixing [45], and feature extraction [46].

D. Generative Adversarial Networks

GANs have been found to be an effective approach to generate outputs that could realistically exist in the training set [17]. These models have been widely applied in computer vision for generating high-resolution images [47], medical imaging [48], and many others. In remote sensing, GANs have been shown to be effective at super-resolution [49], pan sharpening [50], and hyperspectral classification [51]. The idea of GANs is to learn a classifier to discriminate between real and generated examples, named an adversarial network. The generator then aims to fool the discriminator $D(x)$ by learning to produce realistic looking outputs from a random variable z . To learn these functions, the generators $G(z)$ and $D(x)$ compete with each other as a minimax optimization problem. The corresponding loss for GANs can be written as: $L(\Theta) = E_x(\log(D(x))) + E_z(\log(1 - D(G(z))))$ where z is a random variable. This loss can then be optimized through gradient descent.

III. APPROACH

VAEs and GANs are effective for image-to-image translation where pairs of images are not available [14]. This is the case for satellites with no space–time overlap. However, as in [14], a shared latent variable z can be used to approximate the joint distribution from marginal. An adversarial loss applied to cross reconstructions satisfies the shared latent space assumption but is underconstrained for high-dimensional, multispectral images. We shall observe that this leads to large errors in our task. To address this, we introduce a shared spectral reconstruction loss and skip connection to effectively generate synthetic spectral bands (see Fig. 2), and the result is a 50%–80% reduction in mean absolute error (MAE).

In the spectral domain, we consider the case of K satellites, $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K\}$, such that $\mathcal{S}_k \in \mathbb{R}^{H \times W \times B_k}$ is a set of B_k spectral bands with height H and width W , shown as a Venn

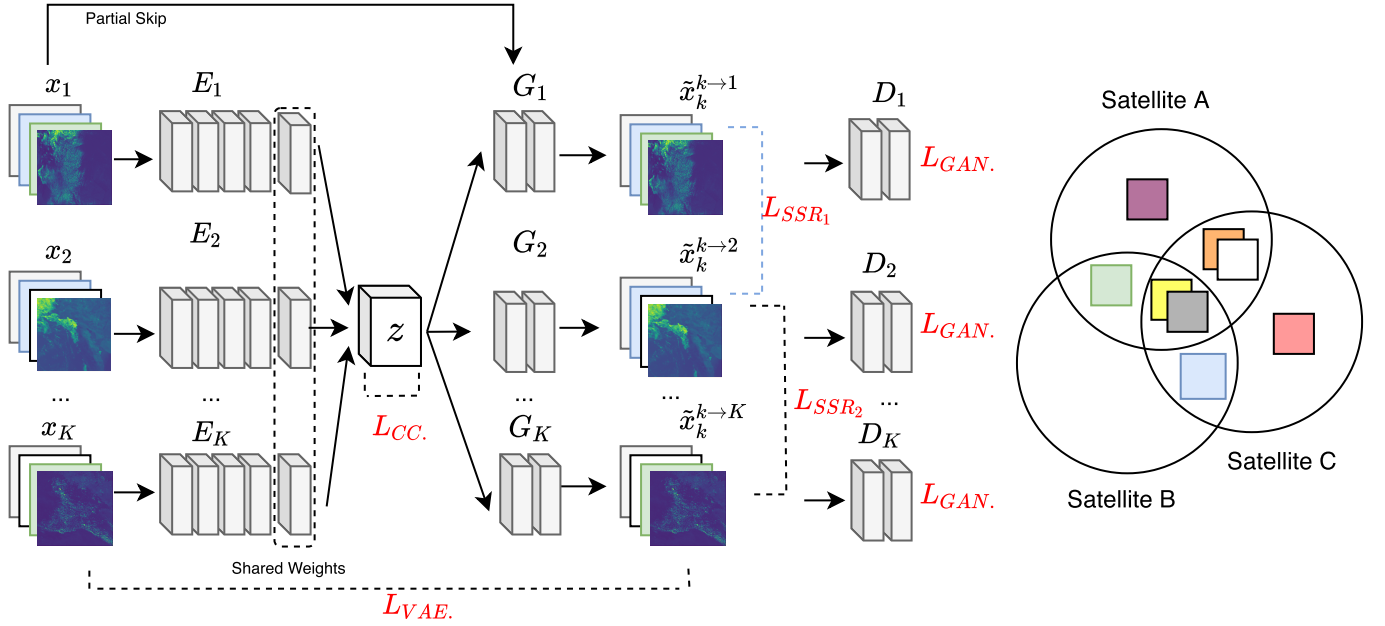


Fig. 2. (Left) Network architecture for $K = 3$ satellites. Encoders (E_k), decoders (G_k), and discriminators (D_k) are networks with residual blocks. Losses terms are highlighted in red. (Right) Venn diagram shows how spectral bands can overlap between pairs and multiple satellites.

diagram in Fig. 2. The union of all sets, $\cup_{i=1}^K \mathcal{S}_k$, represents the complete set of spectral channels in the data. We denote the intersection of two spectral sets as overlapping bands. Our goal is to generate synthetic bands well where $\mathcal{S}_i \cap \mathcal{S}_j^c \neq \emptyset$ for $\forall(i, j)$, with c the complement. A shared latent variable z is modeled with a Gaussian prior to learn a general representation for mapping between sets such that the assumptions of shared spectral reconstruction, weight sharing, cycle-consistency, and cross-domain adversarial losses are satisfied.

A. VAE-GAN

For a given spectral set k , we define encoder-generator pairs $\{E_k, G_k\}$ such that $q(z_k|s_k) = \mathcal{N}(E_k(s_k), I)$ and $\hat{s}_k^{k \rightarrow k} = G_k(z_k \sim q_k(z_k|s_k))$ for $s_k \in \mathcal{S}_k$. For any set j , $\hat{s}_k^{k \rightarrow j}$ corresponds to reconstruction from set k to j . The set of encoders $\{E_1, E_2, \dots, E_K\}$ shares their last layer of weights to constrain the latent space to high-level representations. Using prior $p_\eta(z) \sim \mathcal{N}(0, I)$, the VAE likelihood is defined as

$$\mathcal{L}_{\text{VAE}_k}(E, G) = \lambda_1 \mathbf{KL}(q_k(z_k|s_k) \| p_\eta(z)) - \lambda_2 \mathbb{E}_{z_k \sim q_k(z_k|s_k)} [\log p_{G_k}(s_k|z_k)]. \quad (2)$$

Distributions p_{G_k} are modeled as Laplacian distributions and a Gaussian latent space with prior $z \sim \mathcal{N}(0, I)$. GANs are used to enforce realistic spatial/spectral distributions of reconstructed images from the latent space. Discriminator networks D_1 – D_k compare observations with cross reconstructions from the latent space

$$\mathcal{L}_{\text{GAN}_k} = \lambda_3 \mathbb{E}_{s_k \sim P_{\mathcal{S}_k}} [\log D_k(s_k)] + \lambda_3 \sum_{j \neq k} \mathbb{E}_{z_j \sim P_{\mathcal{S}_k}} [\log (1 - D_k(G_k(z_j)))]. \quad (3)$$

Network architectures for E_k , G_k , and D_k follow those used in [14]. Our encoder, E_k , is a CNN with one downsampling layer and 64 units in the latent space. Weights of the last

layer are shared between encoders to constrain the latent space. The decoder, G_k , is also a CNN with four residual blocks and one convolutional transpose layer for upsampling. The discriminator, D_k , consists of two hidden layers with leakyrelu activation and average pooling and is learned with a least-square GAN loss. For details, we refer the reader to the code in the Supplementary Material.¹

B. Cycle Consistency

VAE and GAN losses are underconstrained and do not satisfy the shared latent space constraint alone. As in [14], a cycle-consistent loss is used such that $s_k = F_{j \rightarrow k}(F_{k \rightarrow j}(s_k))$ for all satellite pairs (j, k) , where $F_{k \rightarrow j}(s_k) = G_j(E_k(s_k))$. The loss between s_k and cycled reconstruction $\hat{s}_k^{k \rightarrow j \rightarrow k}$ is written as

$$\begin{aligned} \mathcal{L}_{\text{CC}_{k \rightarrow j}}(E_k, E_j, G_k, G_j) &= \lambda_4 \mathbf{KL}(q_k(z_k|s_k) \| p_\eta(z)) \\ &+ \lambda_4 \mathbf{KL}(q_j(z_j|s_k^{k \rightarrow j}) \| p_\eta(z)) \\ &- \lambda_5 \mathbb{E}_{z_j \sim q_j(z_j|s_k^{k \rightarrow j})} [\log p(G_k(s_k|z_j))]. \end{aligned} \quad (4)$$

With multiple domains, each domain should cycle through every other domain. The cycle-consistency loss for each permutation results in a complete cyclical graph. This loss is written as

$$\mathcal{L}_{\text{CC}_k} = \sum_{k \neq j} \mathcal{L}_{\text{CC}_{k \rightarrow j}}(E_k, E_j, G_k, G_j). \quad (5)$$

C. Shared Spectral Reconstruction Loss

Adversarial losses can be easily fooled with increased dimensions. To help avoid this, we introduce an additional

¹Code: <https://github.com/tjvandal/unsupervised-spectral-synthesis>

Algorithm 1 Generate a Synthetic Band by Translating From One Satellite to Another

Result: Synthetic spectral bandImage s_k from satellite k ;Encode to latent space $z = E_k(s_k)$;Decode to other satellite $\tilde{s}_j = G_j(z)$;Select synthetic band from \tilde{s}_j ;

loss, \mathcal{L}_{SSR_k} . In this problem, if the intersection of spectral channels $S_{k,j} = S_k \cap S_j$ between domains is not empty, then the difference between $p(s_k^{k \rightarrow k} | z_k)$ and $p(s_k^{k \rightarrow j} | z_k)$ can be minimized with KL divergence

$$\mathcal{L}_{SSR_k} = \lambda_6 \sum_{j \neq k} \mathbf{KL}(p(\tilde{s}_k^{k \rightarrow k} | z_k) || p(\tilde{s}_k^{k \rightarrow j} | z_k)) \quad (6)$$

where $\tilde{s}_k \in S_{k,j}$. The SSR loss encourages decoders to reconstruct identical spectral wavelengths with similar distributions while still synthesizing dissimilar bands. In this scenario, partial constraints are placed between domains and allow sampling of unobserved spectra from the shared latent space. By decreasing λ_6 , the bias between bands will be relaxed, which may reduce the effect of more uncertain domains.

D. Total Loss

The likelihood is maximized by optimizing the GAN mini-max problem such that the generator aims to fool the discriminator, alternating updates between (E, G) and (G, D)

$$\mathcal{L} = \min_{E,G} \max_D \sum_{k=1}^K [\mathcal{L}_{VAE_k} + \mathcal{L}_{CC_k} + \mathcal{L}_{GAN_k} + \mathcal{L}_{SSR_k}]. \quad (7)$$

The hyperparameters used correspond to those in [14] and set as $\lambda_1 = 1, \lambda_2 = 0.01, \lambda_3 = 1, \lambda_4 = 1, \lambda_5 = 0.01$, and $\lambda_6 = 0.1$. The Adam optimization is used to train the networks for 200 000 steps with a batch size of 8 with parameters $\beta_1 = 0.5, \beta_2 = 0.999$, and learning rate $1e-5$. The reader can find the detailed information in the Supplementary Material. In the following, we show the steps for generating a new band.

E. Data

Three GEO satellite imagery datasets, GOES-16 (G16), GOES-17 (G17), and Himawari-8 (H8), are used in our experiments. Each satellite captures hemispheric (full-disk) snapshots from a constant vantage point over time but of different regions. Examples are shown in Fig. 1. Images contain 16 bands (channels) in the visible, near-infrared, and thermal spectrum at 0.5–2-km spatial resolution (see Table I). G16 and G17 have identical specifications viewing the east and west regions of North America and include two visible (blue, red), four near-infrared (including cirrus), and ten thermal infrared bands. H8 has 15 overlapping bands with G16/G17 viewing the Pacific Ocean and East Asia, and this ensures similar information content. H8 captures three visible (blue, green, and red), three near-infrared (missing cirrus), and the same ten thermal infrared bands as G16/G17. Visible and near-infrared bands are measured as reflectances in the unitless

TABLE I

SPECTRAL BANDS AND WAVELENGTHS (μm) FOR GOES-16/17 ABI AND HIMAWARI-8 AHL. INCONSISTENT BANDS IN HIGHLIGHTED BOXES

Band	GOES-16/17		Himawari-8/9	
	Central Wavelength	Width	Central Wavelength	Width
1	0.470	0.040	0.455	0.050
2	0.640	0.080	0.510	0.020
3	0.864	0.035	0.645	0.030
4	1.373	0.014	0.86	0.02
5	1.61	0.04	1.61	0.02
6	2.24	0.05	2.26	0.02
7	3.90	0.19	3.85	0.22
8	6.19	0.80	6.25	0.37
9	6.93	0.42	6.95	0.12
10	7.34	0.19	7.35	0.17
11	8.44	0.43	8.60	0.32
12	9.61	0.38	9.63	0.18
13	10.33	0.30	10.45	0.30
14	11.21	0.78	11.20	0.20
15	12.29	0.93	12.35	0.30
16	13.28	0.57	13.30	0.20

range 0–1.6. Thermal infrared bands capture temperature in units kelvin (K). Thus, the G16 and G17 bands all overlap, cirrus (1.37 μm) exists in G16 and G17 but is not in H8, and green (0.51 μm) exists in H8 but not in G16 or G17. These differences cause difficulties when applying models relying on green or cirrus bands across satellite sets.

G16 observes the North, Central, and South Americas, capturing a good distribution of land and ocean. G17 observes the Pacific Ocean as well as most of North and Central America. However, G17 has known problems with its thermal cooling system causing the near- and thermal-infrared channels to be unusable during periods of high heat and biased throughout [52]. This further highlights the gain in replacing low-quality bands of G17 with a virtual sensor. Periods of high heat are filtered out of our training and test sets with quality control checks to eliminate the temporal periods of known uncertainty. After quality control, considerable space–time overlap between G16 and G17 can be used for testing. H8 observes East Asia, Australia, and the Western Pacific, partially overlapping with G17. Discrepancies are expected between sensors caused by different solar and sensor viewing angles, but we are not aware of a more appropriate dataset for evaluation. The data generated by Wang *et al.* [53] are used, which normalized G16, G17, and H8 to a common georeferenced gridding system in order to facilitate inter-comparisons and processed with the bidirectional reflectance distribution function (BRDF). Bands have resolutions varying from 500 m to 2 km which we interpolate to a common suboptimal resolution of 2 km. Full-disk images are on a common grid with tiles of size $300 \times 300 \times 16$. Training data are generated from the multipetabyte datasets. We randomly sample images to build a well-distributed and large training dataset from years 2018 (G16, H8) and 2019 (G17), which totaled 359 GB of data. Each tile is split into $64 \times 64 \times 16$ nonoverlapping patches for training, generating millions of samples.

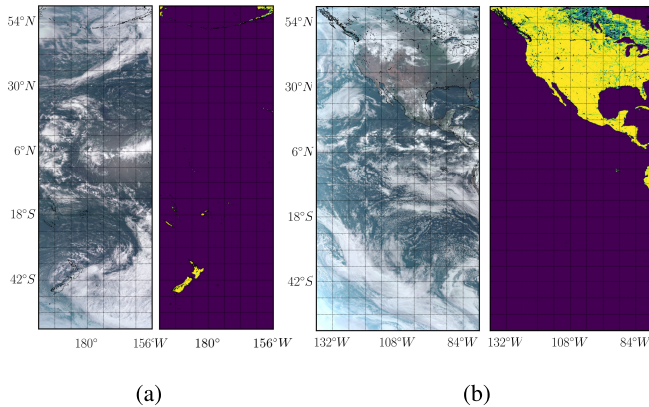


Fig. 3. Overlapping regions shown in true color between (a) H8/G17 and (b) G17/G16 with the corresponding land masks.

A test set including 2000 randomly selected tiles from 2020 from overlapping G16 and G17 observations (see Fig. 3). The random set of tiles assures a range of solar angles, system patterns, and land-cover types. This dataset will be made publicly available consisting of tiles from each satellite. Similarly, a pair of overlapping tiles of data from G17 and H8 on January 2, 2019, at 04:00UTC are selected to compare synthetically generated green and cirrus bands (spatial overlap of G17/H8 is mostly ocean).

IV. EXPERIMENTS AND DISCUSSION

In this section, we present a set of experiments to explore the properties of our approach by testing which bands can be robustly synthesized, how many bands can be generated, and how effectively the proposed loss performs. The metrics Bias, MAE, and Precision are used for evaluation. Relative metrics, RBias and RMAE, are computed by dividing by the mean intensity of the relevant band.

A. Individual Band Synthesis

Our experiments start with testing how well each spectral band can be synthesized. To do this, we remove individual bands from one satellite (G16) during training, synthesize these bands, and compare with the ground-truth observations. We use the full set of bands from the other two satellites during training. This approach is applied on G16 such that each model takes 15 bands of G16 and 16 of G17 and H8.

Three comparisons are used to help put the accuracy of synthesized bands for G16 into context. Ours refers to bands generated using our proposed approach with both SSR loss and skip connection. Ours w/o SSR refers to bands generated with our proposed approach without the shared spectral reconstruction loss. Ours w/o Skip refers to bands generated with our proposed approach without the skip connection between the input and generator. UNIT refers to the unsupervised image-to-image translation baseline as presented in [14] and is equivalent to ours without SSR and skip connection. Sensor refers to the performance if we simply use overlapping observations from another satellite (G17), and this acts as our lower bound in performance and is actually the status-quo

TABLE II
MEAN RBias, RMAE, AND Precision FOR VIS/NIR (REFLECTANCE IN THE RANGE 0–1.6) AND TIR (TEMPERATURE K) BANDS BY METHOD

Method	RBias		RMAE		Precision	
	VIS/NIR	TIR	VIS/NIR	TIR	VIS/NIR	TIR
Cross-Sensor	-0.085	0.0005	0.188	0.020	0.085	4.54
UNIT	0.003	-0.010	0.338	0.049	0.125	5.86
Ours w/o SSR	-0.004	-0.007	0.248	0.030	0.078	4.17
Ours w/o Skip	-0.015	-0.002	0.130	0.008	0.052	1.39
Ours	-0.014	-0.003	0.125	0.007	0.051	1.31
Reconstruct	-0.002	-0.002	0.098	0.005	0.036	0.94

(essentially substituting the missing band with images from the same band but from another satellite, which as we shall see is a suboptimal solution). Reconstruction refers to the images reconstructed from a full model trained on all satellites with no missing bands, and this acts as our upper bound in performance. Each of these signals is computed on our test set of 2000 overlapping tiles from G16/G17. Fig. 4 shows the RBias, RMAE, and Precision for each condition. Similarly, Table II shows the average metrics for VIS/NIR and TIR of each method.

The MAE in the sensor condition is substantial and largely caused by clouds/aerosols in the vertical direction (see gif in the Supplementary Material). On the other side, synthetically generating bands using our approach substantially reduces MAE by over 40% compared to both this baseline and UNIT (see Table II). Similarly, synthesized bands also improve upon the view from G17 even though during training they did not see examples of the corresponding band from G16. Mean biases for all approaches are all near zero but have cross-sensor, UNIT, and ours w/o SSR each have high variance, corresponding to similarly low precision. In contrast, our approach performs similar to reconstruction with lower variance around the mean bias and high precision. Ablation experiments removing the shared spectral reconstruction loss and the skip connection show their effectiveness. SSR is critical to learning a robust latent space and the skip connection improves both VIS/NIR and TIR predictions in terms of RBias, RMAE, and precision. We observe that without introducing the SSR loss, performance is even worse than the sensor baseline. From this, we learn that applying an existing image-to-image translation model [14] to our task, without adaptation, performs poorly. We find that band 7, the shortwave infrared band ($3.9 \mu\text{m}$), is particularly difficult to synthesize with RMAE, RBias, and precision significantly above that of the full reconstruction. This result suggests that the shortwave infrared band captures information that cannot be inferred from the others. Notice how the wavelength gap between bands 7 and 8 is relatively large ($2.3 \mu\text{m}$), and this may explain why the performance is poor. In the future, a similar and more extensive analysis could be used to inform future satellite design configurations.

We show qualitative examples of generating synthetic bands in Figs. 1 and 7. GOES-16 and GOES-17 images in Fig. 1 show the examples of true-color images generated from a synthetic green band. This process is applied to Himawari-8

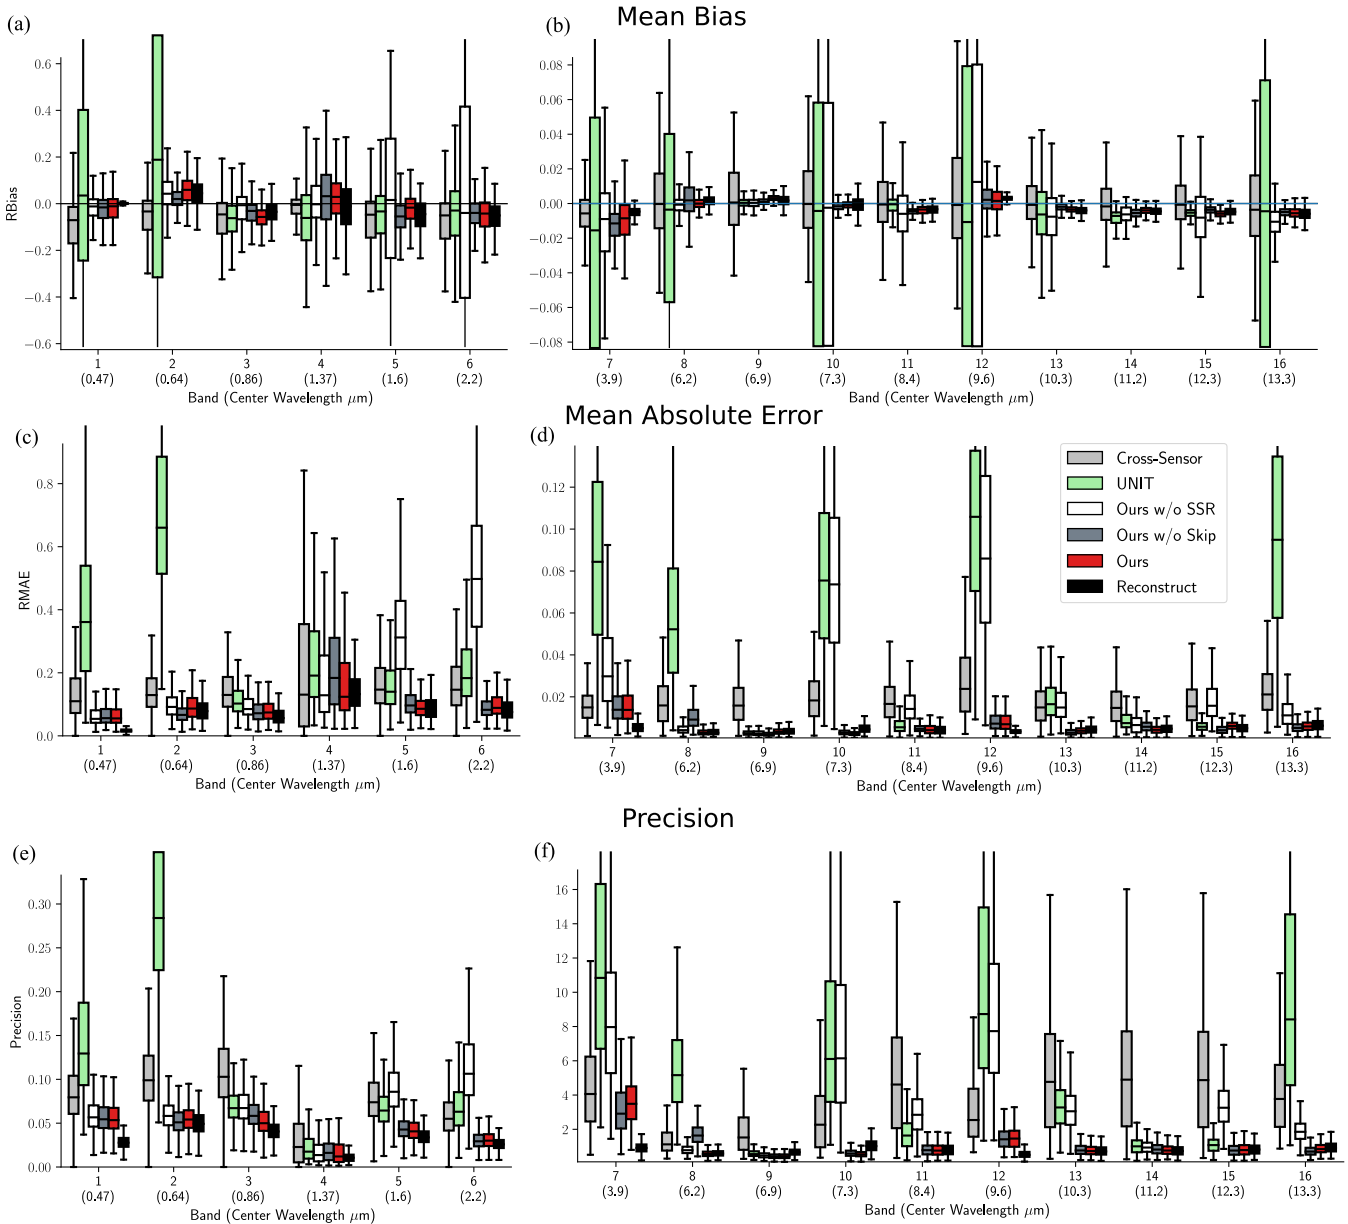


Fig. 4. MAE of a substitute sensor (GT observations from a separate satellite), synthetic sensor without SSR, synthetic sensor without skip, synthetic sensor, and reconstructed sensor (reconstructed images from a full model with access to all bands) for each band. The wavelength of each band is shown below the band number. (a), (c), and (e) include visible and near infra-red (no physical units), while (b), (d), and (f) include thermal infrared measured in kelvin.

to generate a cirrus band (shown in Fig. 7). While there may be challenges in synthetically generating all bands, most can be reconstructed with a high signal-to-noise ratio and this suggests that our approach could be used to make software updates to current satellite datasets.

B. Land Versus Ocean Conditions

We analyze land versus ocean conditions using the models trained in the previous experiment synthesizing individual bands. Specifically, for each sample, we compute RBias, RMAE, and Precision for land and ocean pixels separately using the MODIS land/water mask shown in Fig. 3. The results are shown in Fig. 5 boxplots such that the columns represent metrics and rows correspond to VIS/NIR and TIR bands.

Overall, we find larger errors and reduced precision over land conditions compared to ocean areas. However, the bias, error, and precision of our approach outperform baselines cross sensor and UNIT consistently. Precision for VIS/NIR suggests that land areas are more challenging to recover compared to ocean but continuously outperform the baselines. Future work extended to surface reflectance may consider a similar analysis over different land-cover types. Results on TIR bands show high precision with very low bias.

C. Synthesizing Multiple Bands

Generating synthetic channels from satellites with a limited number of spectral bands could be of significant value for long-term analysis. For example, older generation satellites

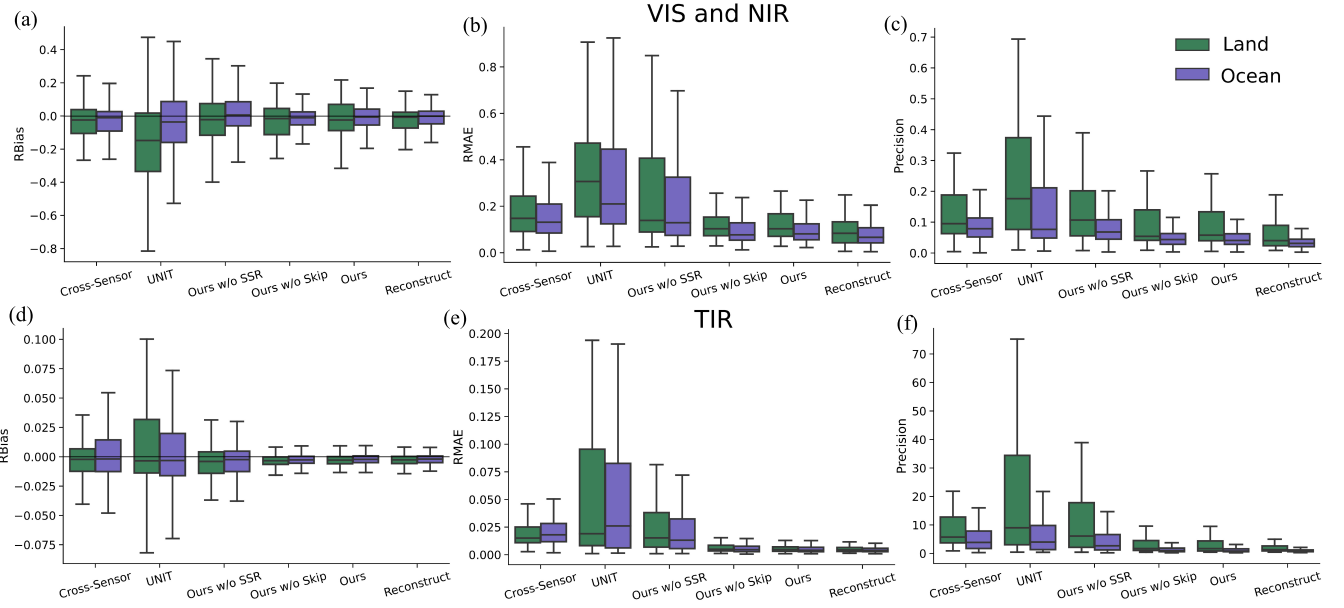


Fig. 5. Performance of models over land versus ocean. Columns correspond to (a) and (d) RBias, (b) and (e) RMAE, and (c) and (f) Precision, while (a)–(c) rows correspond to visible and near-infrared (reflectances) and (d)–(f) thermal infrared (temperature K).

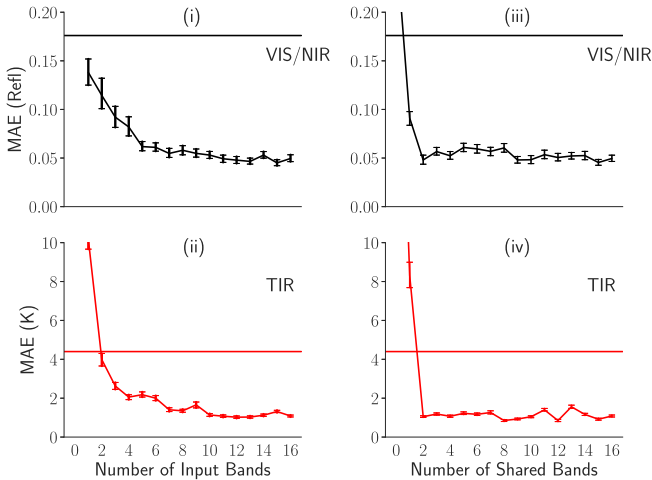


Fig. 6. Horizontal lines correspond to the sensor signal and error bars represent 95% confidence intervals of tile MAE. (i) and (ii) Results of an experiment synthetically generating more and more bands. (iii) and (iv) MAE as a function of number of bands shared between spectral sets. (i) and (iii) correspond to visible/near infrared and (ii) and (iv) correspond to thermal infrared.

often have fewer channels and could provide greater utility in downstream tasks if it was possible to generate images in additional frequency bands. Therefore, we set up an experiment to test how many additional bands can be synthesized reliably and how many initial bands are required. A set of synthesis models were trained on G16, removing bands one by one until just one band was left and while keeping all 16 G17 and H8 bands. For simplicity, and to reduce computation, we dropped bands in a fixed order: 9, 4, 13, 2, 15, 12, 6, 3, 10, 8, 14, 5, 11, 7, and 16. In the most extreme case, we use visible band 1 and attempt to synthesize the remaining 15. As above, results are computed on the test set of 59 overlapping G16 and G17 tiles. The results presented in Fig. 6 (left) show how

the number of available input bands affects the MAE for VIS/NIR (bands 1–6) and TIR (bands 7–16). As expected, MAE falls more or less monotonically as more bands are given as inputs. When just two bands, 1 (blue) and 16 (TIR), are used as inputs, the synthetic TIR reconstruction of G16 still has a lower error than the observed sensor difference between G16 and G17. These results show that few bands are needed to synthesize images that improve upon the status quo. In the TIR range, we find that MAE plateaus after 3–4 bands are used as inputs. These results suggest that the information content in a subset of bands may be sufficient for many applications. However, we should be prepared that some bands may contain specific information useful for monitoring rare events. Overall, these results show that a good proportion of bands can be synthesized remarkably well.

D. Sharing Spectral Losses

The effectiveness of the shared spectral reconstruction loss is tested by gradually increasing the number of shared bands included in the loss one by one. Mathematically, this corresponds to the number of bands included in the set $\mathcal{S}_{k,j}$. In all runs, 16 bands of G16, G17, and H8 are used even if ignored by the SSR loss. Fig. 6(iii) and (iv) shows that the effect of adding shared bands during training leads to a dramatic decrease in MAE. Corresponding cross-sensor signals are shown as horizontal lines. In this setting, we find that using the SSR loss is critical to learning this model. Sharing two spectral bands in the loss function improves the signal and is almost all that is needed for accurate reconstruction. This further reinforces our insight above that a large amount of the information is captured in just a few spectral bands. In Table III, we further explore the SSR loss by testing a range of values for λ_6 from 0.01 to 10. Our results suggest that

TABLE III
MAE FOR VIS/NIR (REFLECTANCE) AND TIR
(TEMPERATURE K) BANDS BY λ_6

λ_6	0.01	0.10	1.0	10.0
VIS/NIR	0.068	0.048	0.042	0.049
TIR	5.03	1.48	1.34	1.28

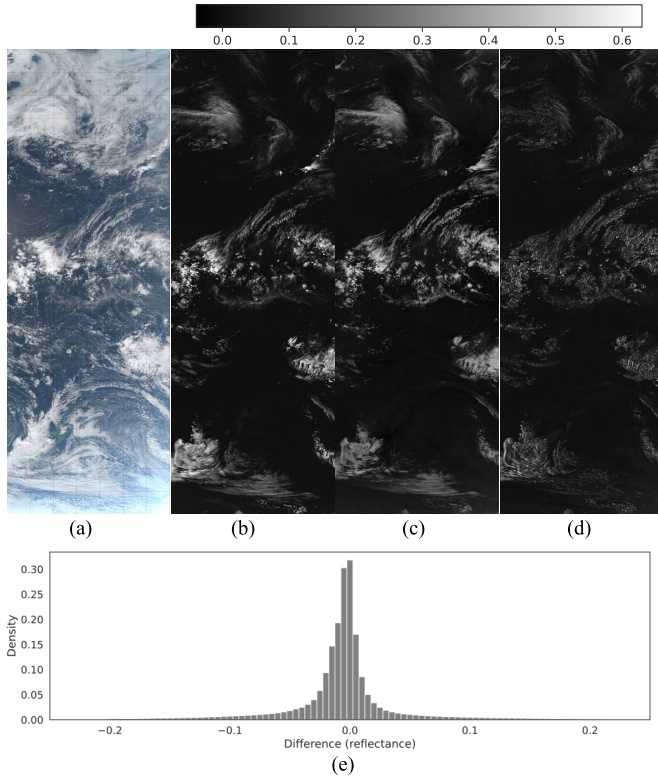


Fig. 7. Synthetically sensing cirrus from Himawari-8 using GOES-17. Images are taken from (a) test set (January 2, 2019, at 4:00UTC). (a) G17 H8 observed true color. (b) Cross-sensor observed G17 cirrus. (c) Synthetic H8 cirrus. (d) Absolute difference between cross sensor and synthetic. (e) Histogram of difference between cross sensor and synthetic.

increasing SSR weighting factor improves the performance on the test set.

E. Synthesizing Cirrus for Himawari-8

As discussed above, the cirrus band ($1.38 \mu\text{m}$) monitors ice particles in the upper troposphere which regulate the climate, and H8 is missing this band. These ice particles are often seen as thin clouds high in the atmosphere, which may be viewed in the visible range, along with other clouds. To generate a synthetic cirrus band, an H8 observation is translated to G17. In Fig. 7, we show four images where G17 and H8 have space-time overlap but different viewing angles. Fig. 7(b)–(e) shows the cross-sensor G17 cirrus band, corresponding H8 synthetic cirrus band, absolute difference between cross sensor and synthetic, and histogram of differences, respectively. This scene consists of clouds of multiple types and atmospheric heights on June 10, 2019, at 04:00UTC. Cirrus clouds are found high in the atmosphere and are seen as thin or wispy (see the lower right portion of the

images). Visually comparing images 7(b) and (c) shows the similarity between synthetic bands and observations. Lower level clouds, which can be seen throughout the true-color image, are ignored by both the observed and synthetic cirrus bands. However, from Fig. 4, we know that cross-sensor errors are large relative to real observations. Hence, the quantitative difference between synthetic and cross sensor is unsurprising. The differences are normally distributed, which indicates random noise. These results suggest that our learned latent space can represent unobserved bands, distinguishing different types of material.

F. Limitations

While the VAE-GAN architecture performs well overall, it does present some limitations. VAEs aim to explicitly model the data as a multivariate Gaussian and often produce blurry outputs. The GAN counteracts this effect by discriminating between real and generated images. However, there is a concern that this reduces data precision and fails to detect rare and anomalous events which may affect scientific applications. Extending our work to use normalizing flows, as in [16], may reduce this limitation.

V. CONCLUSION

We have presented an unsupervised learning approach for satellite-to-satellite translation that can be used to synthesize unobserved spectral bands. A novel shared spectral reconstruction loss is presented to further constrain learning and conserve spectral information and a partial skip connection maintains spatial consistency. Experiments with sensors on the GOES-16/17 and Himawari-8 satellites show that synthetic spectral bands can be generated through reconstruction from a shared latent space. For the first time, we are able to generate true-color images from GOES-16/17 and the cirrus band from Himawari-8, generating further value from these satellites. Future work may consider conditioning the shared latent space with known physical properties and extending to additional tasks.

ACKNOWLEDGMENT

The authors would like to thank the GOES-R data providers and the GeoNEX Team, Ames Research Center, Moffett Field, CA, USA for access to GEO L1G data. NASA’s Advanced Supercomputing Facility at the Ames Research Center was used for accelerated computing with Nvidia V100 GPUs.

REFERENCES

- [1] *Global Risks Report 2020*, W. E. Forum, Cologne, Switzerland, 2020.
- [2] M. A. Wulder and N. C. Coops, “Satellites: Make earth observations open access,” *Nature News*, vol. 513, no. 7516, p. 30, 2014.
- [3] 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.
- [4] S. S. Board and National Academies of Sciences, Engineering, and Medicine, *Thriving on our Changing Planet: A Decadal Strategy for Earth Observation From Space*. Washington, DC, USA: The National Academies Press, 2018. [Online]. Available: <https://www.nap.edu/catalog/24938/thriving-on-our-changing-planet-a-decadal-strategy-for-earth>

- [5] C. S. Velden, C. M. Hayden, S. J. W. Nieman, W. P. Menzel, S. Wanzong, and J. S. Goerss, "Upper-tropospheric winds derived from geostationary satellite water vapor observations," *Bull. Amer. Meteorol. Soc.*, vol. 78, no. 2, pp. 173–196, 1997.
- [6] C. S. Velden, T. L. Olander, and R. M. Zehr, "Development of an objective scheme to estimate tropical cyclone intensity from digital geostationary satellite infrared imagery," *Weather Forecasting*, vol. 13, no. 1, pp. 172–186, Mar. 1998.
- [7] X. Zhang, S. Kondragunta, J. Ram, C. Schmidt, and H.-C. Huang, "Near-real-time global biomass burning emissions product from geostationary satellite constellation," *J. Geophys. Res., Atmos.*, vol. 117, no. D14, pp. 1–19, Jul. 2012.
- [8] W. P. Menzel *et al.*, "Application of GOES-8/9 soundings to weather forecasting and nowcasting," *Bull. Amer. Meteorol. Soc.*, vol. 79, no. 10, pp. 2059–2078, 1998.
- [9] T. J. Schmit, P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebar, "A closer look at the ABI on the GOES-R series," *Bull. Amer. Meteorol. Soc.*, vol. 98, no. 4, pp. 681–698, Apr. 2017.
- [10] K. Bessho *et al.*, "An introduction to himawari-8/9—Japan's new-generation geostationary meteorological satellites," *J. Meteorol. Soc. Japan. II*, vol. 94, no. 2, pp. 151–183, 2016.
- [11] K.-N. Liou, "Influence of cirrus clouds on weather and climate processes: A global perspective," *Monthly Weather Rev.*, vol. 114, no. 6, pp. 1167–1199, Jun. 1986.
- [12] A. J. Heymsfield *et al.*, "Cirrus clouds," *Meteorol. Monographs*, vol. 58, no. 1, pp. 1–2, 2017.
- [13] R. Zhang, P. Isola, and A. A. Efros, "Colorful image colorization," in *Computer Vision—ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham, Switzerland: Springer, 2016, pp. 649–666.
- [14] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 700–708.
- [15] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 2223–2232.
- [16] A. Grover, C. Chute, R. Shu, Z. Cao, and S. Ermon, "AlignFlow: Cycle consistent learning from multiple domains via normalizing flows," 2019, *arXiv:1905.12892*. [Online]. Available: <https://arxiv.org/abs/1905.12892>
- [17] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [18] T. J. Schmit, S. S. Lindstrom, J. J. Gerth, and M. M. Gunshor, "Applications of the 16 spectral bands on the advanced baseline imager (ABI)," *J. Oper. Meteorol.*, vol. 6, no. 4, pp. 33–46, 2018, doi: 10.15191/nwajom.2018.0604.
- [19] G. Chander *et al.*, "Applications of spectral band adjustment factors (SBAF) for cross-calibration," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 3, pp. 1267–1281, Mar. 2013.
- [20] A. Santamaria-Artigas, E. F. Vermote, B. Franch, J.-C. Roger, and S. Skakun, "Evaluation of the AVHRR surface reflectance long term data record between 1984 and 2011," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 98, Jun. 2021, Art. no. 102317.
- [21] M. Claverie *et al.*, "The harmonized landsat and sentinel-2 surface reflectance data set," *Remote Sens. Environ.*, vol. 219, pp. 145–161, Dec. 2018.
- [22] K.-L. Hsu, X. Gao, S. Sorooshian, and H. V. Gupta, "Precipitation estimation from remotely sensed information using artificial neural networks," *J. Appl. Meteorol.*, vol. 36, no. 9, pp. 1176–1190, Sep. 1997.
- [23] M. Sadeghi *et al.*, "PERSIANN-CNN: Precipitation estimation from remotely sensed information using artificial neural networks—convolutional neural networks," *J. Hydrometeorol.*, vol. 20, no. 12, pp. 2273–2289, Dec. 2019.
- [24] N. Jean, M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, vol. 353, no. 6301, pp. 790–794, Aug. 2016.
- [25] S. Lei, Z. Shi, and Z. Zou, "Super-resolution for remote sensing images via local-global combined network," *IEEE Geosci. Remote Sens. Lett.*, vol. 14, no. 8, pp. 1243–1247, Aug. 2017.
- [26] F. Ling *et al.*, "Measuring river wetted width from remotely sensed imagery at the subpixel scale with a deep convolutional neural network," *Water Resour. Res.*, vol. 55, no. 7, pp. 5631–5649, Jul. 2019.
- [27] K. Duffy, T. J. Vandal, S. Li, R. R. Nemani, and A. Ganguly, "Deep learning emulation of atmospheric correction for geostationary sensors," in *Proc. AGU Fall Meeting Abstracts*, vol. 2019, 2019, p. A41T-2671.
- [28] S. Basu, S. Ganguly, S. Mukhopadhyay, R. DiBiano, M. Karki, and R. Nemani, "DeepSat: A learning framework for satellite imagery," in *Proc. 23rd SIGSPATIAL Int. Conf. Adv. Geograph. Inf. Syst.*, Nov. 2015, pp. 1–10.
- [29] Z. Shi, C. Chen, Z. Xiong, D. Liu, and F. Wu, "HSCNN+: Advanced CNN-based hyperspectral recovery from RGB images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2018, pp. 939–947.
- [30] N. Akhtar and A. Mian, "Hyperspectral recovery from RGB images using Gaussian processes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 1, pp. 100–113, Jan. 2020.
- [31] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 1125–1134.
- [32] H. Tang, D. Xu, N. Sebe, Y. Wang, J. J. Corso, and Y. Yan, "Multi-channel attention selection GAN with cascaded semantic guidance for cross-view image translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 2417–2426.
- [33] Z. Yi, H. Zhang, P. Tan, and M. Gong, "DualGAN: Unsupervised dual learning for image-to-image translation," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2849–2857.
- [34] M.-Y. Liu *et al.*, "Few-shot unsupervised image-to-image translation," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2019, pp. 10551–10560.
- [35] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013, *arXiv:1312.6114*. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [36] J.-Y. Zhu *et al.*, "Toward multimodal image-to-image translation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 465–476.
- [37] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8789–8797.
- [38] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 172–189.
- [39] J. Hoffman *et al.*, "CyCADA: Cycle-consistent adversarial domain adaptation," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1989–1998.
- [40] W. Deng, L. Zheng, Q. Ye, G. Kang, Y. Yang, and J. Jiao, "Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 994–1003.
- [41] H.-Y. Lee, H.-Y. Tseng, J.-B. Huang, M. Singh, and M.-H. Yang, "Diverse image-to-image translation via disentangled representations," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 35–51.
- [42] E. H. Sanchez, M. Serrurier, and M. Ortner, "Learning disentangled representations via mutual information estimation," 2019, *arXiv:1912.03915*. [Online]. Available: <https://arxiv.org/abs/1912.03915>
- [43] G. E. Hinton, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, Jul. 2006.
- [44] X. Wang, K. Tan, Q. Du, Y. Chen, and P. Du, "CVA2E: A conditional variational autoencoder with an adversarial training process for hyperspectral imagery classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 8, pp. 5676–5692, Aug. 2020.
- [45] Y. Su, J. Li, A. Plaza, A. Marinoni, P. Gamba, and S. Chakraborty, "DAEN: Deep autoencoder networks for hyperspectral unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 7, pp. 4309–4321, Jul. 2019.
- [46] W. Yu, M. Zhang, and Y. Shen, "Spatial revising variational autoencoder-based feature extraction method for hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 2, pp. 1410–1423, Feb. 2021.
- [47] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," 2018, *arXiv:1809.11096*. [Online]. Available: <https://arxiv.org/abs/1809.11096>
- [48] K. Armanian *et al.*, "MedGAN: Medical image translation using GANs," *Comput. Med. Imag. Graph.*, vol. 79, Jan. 2020, Art. no. 101684.
- [49] K. Jiang, Z. Wang, P. Yi, G. Wang, T. Lu, and J. Jiang, "Edge-enhanced GAN for remote sensing image superresolution," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 8, pp. 5799–5812, Aug. 2019.
- [50] W. Xie, Y. Cui, Y. Li, J. Lei, Q. Du, and J. Li, "HPGAN: Hyperspectral pansharpening using 3-D generative adversarial networks," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 1, pp. 463–477, Jan. 2021.
- [51] L. Zhu, Y. Chen, P. Ghamisi, and J. A. Benediktsson, "Generative adversarial networks for hyperspectral image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 9, pp. 5046–5063, Sep. 2018.
- [52] G.-R. Series. *Goes-17 ABI Performance*. Accessed: May 30, 2020. [Online]. Available: <https://www.goes-r.gov/users/goes-17-abi-performance.html>
- [53] W. Wang *et al.*, "An introduction to the geostationary-NASA Earth exchange (GeoNEX) products: 1. Top-of-atmosphere reflectance and brightness temperature," *Remote Sens.*, vol. 12, no. 8, p. 1267, Apr. 2020.



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

He is the Research Scientist of the NASA Ames Research Center, Moffett Field, CA, USA, and the Bay Area Environmental Research Institute, Moffett Field. His research interests include machine learning and Earth science applications.



Kate Duffy received the B.Sc. degree in environmental engineering from Cornell University, Ithaca, NY, USA, in 2015, and the M.Sc. degree in civil engineering from Northeastern University, Boston, MA, USA, in 2019. She is pursuing the Ph.D. degree with Northeastern University, advised by Dr. Auroop Ganguly and a NASA Pathways Intern at the NASA Earth eXchange (NEX).

Her research is at the intersection of climate, remote sensing, and machine learning with applications to temperature variability and ecological resilience.



Daniel McDuff (Member, IEEE) received the bachelor's degree (Hons.) and the master's degree in engineering from the University of Cambridge, Cambridge, U.K., in 2008 and 2009, respectively, and the Ph.D. degree from the MIT Media Laboratory, while working with the Affective Computing Group, in 2014.

He is the Principal Researcher of Microsoft Research, Redmond, WA, USA. Previously, he was the Director of Research at MIT Media Laboratory Spin-Out Affectiva, Boston, MA, USA.



Andrew Michaelis received the B.Sc. degree in computer science from the University of Montana, Missoula, MT, USA, in 2003.

He is working as a Staff Software Engineer with the Data Visualization Group, NASA Advanced Supercomputing Facility, NASA Ames Research Center, Moffett Field, CA, USA. His core focus is supporting researchers with large scale remote sensing, modeling, and data analytics problems in the Earth Sciences domain.



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 the Research Scientist of the NASA Ames Research Center, Moffett Field, CA, USA, and 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 the Senior Earth Scientist of 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.