

Resolution Revolution: A Physics-Guided Deep Learning Framework for Spatiotemporal Temperature Reconstruction

Shengjie Liu¹ Lu Zhang¹ Siqin Wang¹

¹University of Southern California

{skrisliu}@gmail.com, {lzhang63, siqinwan}@usc.edu

Abstract

Central to Earth observation is the trade-off between spatial and temporal resolution. For temperature, this is especially critical because real-world applications require high spatiotemporal resolution data. Current technology allows for hourly temperature observations at 2 km, but only every 16 days at 100 m, a gap further exacerbated by cloud cover. Earth system models offer continuous hourly temperature data, but at a much coarser spatial resolution (9-31 km). Here, we present a physics-guided deep learning framework for temperature data reconstruction that integrates these two data sources. The proposed framework uses a convolutional neural network that incorporates the annual temperature cycle and includes a linear term to amplify the coarse Earth system model output into fine-scale temperature values observed from satellites. We evaluated this framework using data from two satellites, GOES-16 (2 km, hourly) and Landsat (100 m, every 16 days), and demonstrated effective temperature reconstruction with hold-out and in situ data across four datasets. This physics-guided deep learning framework opens new possibilities for generating high-resolution temperature data across spatial and temporal scales, under all weather conditions and globally. Code and data are available at skrisliu.com/r2lst.

1. Introduction

Surface temperature is a critical physical property of the Earth's system and an important climate indicator [8]. Over land, land surface temperature (LST) can be heterogeneous due to complex surface characteristics, with urbanization further exacerbating these differences [17]. Earth system models are developed through synthesizing satellite and meteorological station observations, using physics-based energy balance modeling to simulate surface dynamics [7]. Currently, the best global simulations achieve a spatial resolution of 9 km over land areas and 31 km globally, much coarser than commonly used satellite data such as Landsat

(100 m) and GOES-16 (2 km) [11].

High-resolution temperature data are important for many real-world applications [19, 21]. However, the current lack of high spatiotemporal resolution temperature data has imposed limitations on many applications [30, 32]. For example, central to causal inference in climate and health research is the use of time-series data to track changes in critical health statistics under varying daily temperatures. Yet, due to cloud cover and long revisit times, temperature data obtained from satellites are often incomplete, sometimes offering only 1–2 observations every 16 days [22]. Cloud cover further exacerbates these limitations, with an average of 66% of the Earth's surface obscured by clouds [14]. The lack of seamless daily temperature data has constrained the use of satellite-derived temperature products in real-world applications. As a result, many existing studies rely on temperature data from sparsely distributed meteorological stations, which offer limited spatial coverage and are unable to capture temperature gradients within cities [37].

Integrating coarse-resolution, continuous temperature data from Earth system models with high-resolution, cloud-contaminated satellite observations offers a promising pathway toward generating seamless daily temperature datasets [22, 33, 36]. Despite the recent rise in high-resolution temperature data, the exponential growth in computational resources, and the rapid advancement of Earth system models, research in this area remains relatively limited, with existing studies adopting a variety of approaches from different perspectives. For example, the reanalysis and thermal merging (RTM) method integrates 30 km reanalysis temperature data from a land data assimilation system and the 1 km MODIS observational data through multiple steps, using the annual temperature cycle (ATC) to capture the average trend, employing random forests to model daily fluctuations, and using a search-window algorithm to account for spatial residuals, achieving 2.03–3.98 K accuracy [36]. Similarly, a three-step approach was proposed to integrate 30 km reanalysis data and 4 km geostationary satellite data, with reconstruction accuracy at 3.57–3.94 K [5]. A more recent study proposed a two-step approach combining an

enhanced ATC with ERA5 reanalysis temperature data and Gaussian processes to account for daily variation, achieving daily 30 m temperature data reconstruction at 1.48–2.11 K accuracy [22]. These methods can be categorized as two-stage or multi-stage approaches, which separately capture the annual trend (often through the ATC) and daily fluctuations. However, they are typically hand-crafted and require a significant amount of manual processing and thresholding. End-to-end training has not yet been achieved in temperature reconstruction tasks.

Deep learning for temperature reconstruction. Deep learning models present an opportunity for a unified, end-to-end framework to reconstruct seamless surface temperature data [18, 26, 31]. However, adapting deep learning approaches for satellite temperature reconstruction is challenging. A typical machine learning task requires pairs of training samples $\{x, y\}$, and the algorithm learns a mapping function f such that

$$y \sim f(x), \quad (1)$$

with x as the input features and y the target. Typically in Earth observation, x is multispectral imagery, and y is the corresponding label. For temperature data reconstruction, two factors limit its direct adaptation. First, temperature reconstruction involves time-series data that varies over time at a given location, yet the features x are often only available monthly or even annually [19]. Second, optical remote sensing features x can only be collected when clouds are not present, but the primary interest often lies in reconstructing temperature under cloudy conditions. Another potential framework is super-resolution, where x is coarse-resolution data and y the corresponding high-resolution data [9, 23]. However, integrating temporal dynamics remains difficult, and training still requires cloud-free pairs.

In this paper, we propose a different strategy: instead of directly predicting temperature, we predict the parameters of the ATC to enable full temperature reconstruction under all-weather conditions. We use the continuous coarse-resolution Earth system reanalysis temperature values as the signal and amplify them to capture daily fluctuations, and use convolutional layers (specifically U-Net) on features representing Earth surface properties to capture spatiotemporal dependencies and minimize bias. The proposed framework allows for unified, end-to-end training.

2. Related Work

2.1. Image Inpainting: Spatial Interpolation

In Earth observation, temperature data reconstruction is typically categorized into three distinct tasks. The first focuses on recovering missing values from a single time

snapshot, where portions of a temperature image are obscured due to cloud cover or other events. This task is analogous to the *image inpainting* vision task [34]. This is the most-studied framework for temperature data reconstruction, with various methods having been developed and applied, including graph-based modeling, Gaussian processes, Fourier convolution techniques, and adversarial training [3, 10, 13, 20, 27]. These approaches rely on the presence of at least some scattered observations at the same time within the spatial domain to guide the reconstruction. Unlike image inpainting, where visual plausibility is often sufficient and the plausible reconstruction solutions are acceptable, temperature reconstruction has hard ground truth, and it is essential to quantitatively evaluate the reconstruction results.

2.2. Two-Step Spatiotemporal Reconstruction

Spatiotemporal reconstruction methods aim at overcoming the limitation from the same-time valid observations by interpolating through the temporal direction. Specific to Earth observation, repeated acquisitions of imagery over the same geographic area are common, making time-series information particularly valuable for temperature reconstruction. A typical spatiotemporal reconstruction approach applies a temporal reconstruction function, such as the ATC, to estimate temperature trends over time first, and then residuals between the reconstructed and observed values are computed and modeled using spatial reconstruction techniques, such as Gaussian processes, graph-based models, or window-based pixel similarity searches [22, 35, 38]. These methods generally rely on hand-crafted features and involve separate steps for temporal and spatial modeling. As a result, the overall workflow is complicated and often requires manual intervention. Moreover, because temporal and spatial components are modeled independently, the final solution may be suboptimal due to the lack of joint optimization and end-to-end training.

2.3. Change Detection and Super-Resolution

An alternative framework, more directly aligned with modern deep learning paradigms, treats temperature reconstruction as a change detection problem. Given a fully observed thermal image and a subsequent image partially obscured by clouds, temperature changes can be computed over the observed pixels. With additional constraints, often derived from land cover information, these changes can be extrapolated to the obscured regions, enabling full-scene reconstruction [1, 24]. This approach can be further extended to super-resolution tasks [16]. Specifically, given a pair of aligned low-res and high-res thermal images, and a new low-res image, the goal is to generate the corresponding new high-res image. However, a common criticism of this framework is its reliance on clear-sky conditions [12, 33].

Under cloudy conditions, differences in surface energy balance result in temperature distributions that diverge significantly from those under clear skies. Consequently, change detection and super-resolution models trained on clear-sky data often fail to generalize, limiting their effectiveness in all-weather temperature reconstruction scenarios.

3. Methodology

3.1. Overview: Physics-Guided Deep Learning With Time Consideration

In the classical machine learning paradigm, given training samples $\{x, y\}$, the goal is to learn a model \mathcal{M} such that

$$y \sim \mathcal{M}(x). \quad (2)$$

Typically, x is a set of features representing the Earth surface, which in vision tasks often correspond to images. This paradigm treats each sample independently, ignoring any temporal ordering. However, in temperature reconstruction tasks, time is a critical factor that must be explicitly incorporated:

$$y(t) \sim \mathcal{M}(x, t). \quad (3)$$

A naïve approach is to include time t as an additional input feature alongside x . Yet, this does not exploit the intrinsic temporal structure of the data, and due to the flexibility of neural networks, it risks producing unrealistic predictions that deviate from physical reality. To address this, we introduce the *Annual Temperature Cycle* (ATC) as a physical constraint within the network architecture. The ATC component explicitly models the overall seasonal trend, capturing significant time-series characteristics. To account for daily fluctuations, we introduce one additional linear term applied to the coarse-resolution temperature data from ERA5. Finally, the convolutional layers focus on learning the remaining spatiotemporal variations only. Specifically, the proposed model comprises three additive components, all embedded within a CNN:

$$y(t, \mathbf{X}) \sim \mathcal{M}_{\text{ATC}}(t | \phi_{\text{ATC}}) + \mathcal{M}_{\rho}(t | \phi_{\rho}) + \mathcal{M}_{\text{conv}}(\mathbf{X} | \phi_{\text{conv}}), \quad (4)$$

where t is time, \mathbf{X} is features representing the Earth surface, and $\phi = \{\phi_{\text{ATC}}, \phi_{\rho}, \phi_{\text{conv}}\}$ denotes the learnable parameters of each component.

3.2. Problem Setup: Temperature Reconstruction

Let $\mathbf{T} \in \mathbb{R}^{H \times W \times C}$ denote the input temperature tensor, where H and W are the height and width of the tensor, C is the number of time steps in the time-series, and some entries may be missing (NaN). We define a binary mask $\mathbf{M} \in \{0, 1\}^{H \times W \times C}$ indicating observed and missing values:

$$\mathbf{M}(i, j, t) = \begin{cases} 1, & \text{if } \mathbf{T}(i, j, t) \text{ is observed,} \\ 0, & \text{if } \mathbf{T}(i, j, t) \text{ is missing (NaN).} \end{cases} \quad (5)$$

We build a model g_{ϕ} for temperature reconstruction, which is to map \mathbf{T} to its estimate $\hat{\mathbf{T}}$, composed of three additive modules:

$$\hat{\mathbf{T}} = \underbrace{\mathcal{M}_{\text{ATC}}(\mathbf{T} | \phi_{\text{ATC}})}_{\text{annual temperature cycle}} + \underbrace{\mathcal{M}_{\rho}(\mathbf{T}_c | \phi_{\rho})}_{\text{daily fluctuation from ERA5}} + \underbrace{\mathcal{M}_{\text{resid}}(\mathbf{X} | \phi_{\text{resid}})}_{\text{spatiotemporal correction}}. \quad (6)$$

where:

- $\mathcal{M}_{\text{ATC}} : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{H \times W \times C}$ is a pixel-wise overall temperature trend over time, and we use the annual temperature cycle for modeling, with parameters ϕ_{ATC} . Each pixel (i, j) has its own learned parameters,
- $\mathcal{M}_{\rho} : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{H \times W \times C}$ is a pixel-wise module that introduces daily temperature fluctuations. It applies a linear transformation to the resampled ERA5 temperature, denoted as $\mathbf{T}_c \in \mathbb{R}^{H \times W \times C}$, and is parameterized by ϕ_{ρ} .
- $\mathcal{M}_{\text{resid}} : \mathbb{R}^{H \times W \times C \times x} \rightarrow \mathbb{R}^{H \times W \times C}$ is a learned, high-frequency correction that accounts for unresolved spatial and temporal structure. This component replaces the spatial filtering process commonly used in two-stage approaches and is the key of the proposed framework to achieve end-to-end training in temperature reconstruction.

An overview of the proposed framework is shown in Fig. 1. In the first branch, the incomplete temperature data are used to calculate the parameterization of the ATC model via optimization, and the ATC parameters are then decoded to reconstruct the annual trend. In the second branch, ERA5 reanalysis temperature data are used to parameterize a linear function, where the learned parameter ρ amplifies the ERA5 temperature values to capture daily fluctuations. The third branch is a U-Net-based CNN that uses surface reflectance as input to construct any remaining spatiotemporal residuals not captured by the first two components. The three components are then combined to generate the final reconstructed temperature output. This is a unified, end-to-end training framework.

3.3. Specific Model Components

Annual Temperature Cycle. The ATC component models the overall temporal temperature trend at each pixel (i, j) using a cosine function:

$$\mathcal{M}_{\text{ATC}}(i, j, t | \phi_{\text{ATC}}) = a_{i,j} + b_{i,j} \cos\left(\frac{2\pi t}{T} + \varphi_{i,j}\right), \quad (7)$$

where $a_{i,j}$, $b_{i,j}$, and $\varphi_{i,j}$ are learnable parameters corresponding to the annual mean temperature, amplitude, and phase shift, respectively, and T is the period (e.g., 365 days).

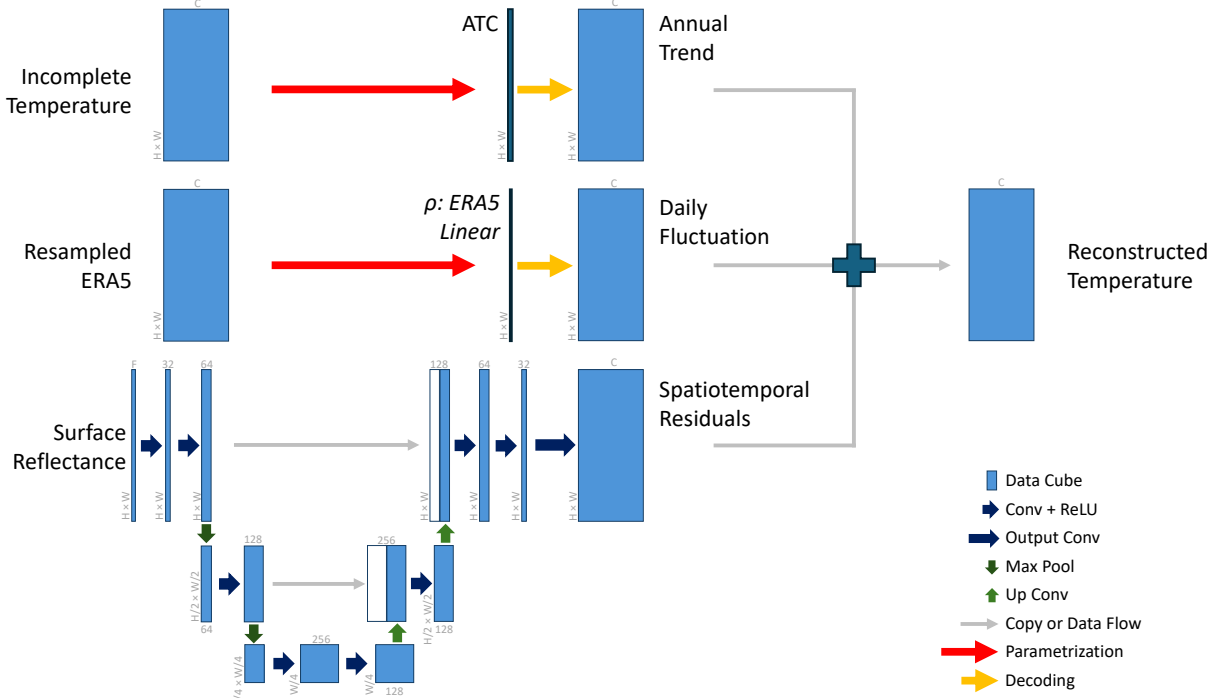


Figure 1. Overview of the proposed method. The annual trend and daily fluctuation are derived from physics-guided models (ATC and ERA5), while a U-Net-based CNN learns spatiotemporal residuals from surface reflectance. All components are combined in a unified, end-to-end trainable deep learning framework.

Daily fluctuation from ERA5 Data with Partial Satellite Observations. This component models the daily fluctuations using coarse-resolution Earth system model temperature data, denoted as $\mathbf{T}_c(i, j, t)$, where $\mathbf{T}_c \in \mathbb{R}^{H \times W \times T}$ represents the resampled data at the target resolution:

$$\mathcal{M}_\rho(i, j, t | \phi_\rho) = w_{i,j} \cdot \mathbf{T}_c(i, j, t), \quad (8)$$

where $w_{i,j}$ are pixel-wise learnable weights. The resampled temperature $\mathbf{T}_c(i, j, t)$ is obtained via

$$\mathbf{T}_c(i, j, t) = \text{Resample}(\tilde{\mathbf{T}}_c(t)) \quad \text{at pixel } (i, j), \quad (9)$$

where $\tilde{\mathbf{T}}_c(t)$ denotes the Earth system model temperature at its native resolution, and $\text{Resample}(\cdot)$ is an interpolation operator that maps the coarse-grid data to the fine grid indexed by (i, j, t) . In practice, the best available ERA5 temperature data has a spatial resolution of approximately 9 km over land, whereas the target resolutions in this study are 2 km and 100 m.

Spatiotemporal Correction Using U-Net-based CNN.

The third component is a U-Net-based CNN designed to capture spatiotemporal residuals by modeling the discrepancy between the previous intermediate output and valid observations. The underlying assumption is that these residuals are associated with Earth surface properties, represented

as $\mathbf{X} \in \mathbb{R}^{H \times W \times C_x}$. The residual field is reconstructed using the following convolutional mapping:

$$\mathcal{M}_{\text{resid}}(\mathbf{X} | \phi_{\text{conv}}) = f_{\text{CNN}}(\mathbf{X}; \phi_{\text{conv}}), \quad (10)$$

where f_{CNN} denotes a CNN with a U-Net architecture, allowing effective multi-scale feature extraction and integration of both local details and global context.

Specifically, we implement a simplified U-Net-style CNN to predict the remaining spatiotemporal residuals from spectral reflectance that are not captured by the previous two components. The U-Net architecture has a depth of three layers and takes the surface reflectance tensor \mathbf{X} as input. It transforms \mathbf{X} through the U-Net to reconstruct the temperature residual surfaces $\mathbf{T}_{\text{resid}}$, which are then added to the outputs from the ATC (annual trend) and the ERA5 linear term (daily fluctuation) to generate the final temperature reconstruction.

3.4. Loss Function

To handle missing data, the reconstruction loss is computed only over observed values, using the mask \mathbf{M} :

$$\mathcal{L}_{\text{rec}}(\phi) = \frac{1}{\sum_{i,j,t} \mathbf{M}(i, j, t)} \left\| \mathbf{M} \odot (\hat{\mathbf{T}} - \mathbf{T}) \right\|_1, \quad (11)$$

where \odot denotes element-wise multiplication, and $\|\cdot\|_1$ represents the element-wise L1 norm (i.e., the sum of absolute differences). We adopt the L1 loss instead of mean squared error (MSE) loss because it is less sensitive to outliers and typically yields better convergence during training.

4. Results and Analysis

We test the proposed framework and method on four datasets obtained from two different satellite sensors (Landsat, GOES-16), using hold-out and *in situ* validation. We compare the proposed method with two exiting approaches: 4-step ATC [38] and GEC-SEB [6].

4.1. Datasets and Experimental Setup

Landsat Data Two separate datasets from the Landsat satellites are constructed. The raw thermal pixel obtained from the Landsat satellites is with 100 m ground sampling distance, and all data pixels are aligned with the multispectral sensor at 30 m on delivery. The Landsat surface temperature data is derived using a single channel algorithm, with 60% of the observations within 2 K accuracy [15]. We obtained all data from 2023 over two separate locations where *in situ* LST sites are available. The first dataset is over the PSU site (40.7°N, 77.9°W) near State College, Pennsylvania. We obtain a total of 46 scenes within the 365-day period. After cloud-masking the data using the cloud masks provided along with the data, we subset a 256×256 area centering at the PSU site. The second dataset is constructed over the BON site (40.5°N, 88.4°W) near Bondville, Illinois. We obtain all scenes from 2023, resulting in a total of 91 available dates (within an overlapping area and thus doubled the frequency). We then again subset a 256×256 area centering at the BON site.

GOES-16 Geostationary Satellite Data Two separate datasets are constructed through the GOES-16 geostationary satellite. GOES-16, the first satellite in NOAA’s Geostationary Operational Environmental Satellites (GOES)-R series, was launched in November 2016 and enabled 2 km resolution temperature monitoring every 5 minutes for the first time [2]. The resulting LST data are produced hourly, with an accuracy of approximately 2.5 K when surface emissivity is known and proper atmospheric correction is applied, and around 5 K otherwise [28]. For the first dataset, we use all GOES-16 data from 2022 at 5:00 local time over a 700×550 pixel region in the east coast centered on New York City, spanning from Québec City to North Carolina and west to Detroit. The second dataset is another adjacent 700×550 pixel region in the Midwest centering around Chicago and includes the Great Lakes at 19:00 local time. We use five GOES-16 spectral bands centered at 0.47, 0.64, 0.86, 1.61, and 2.24 μm as input features.

Validation Data Apart from the standard hold-out validation, we obtain the *in situ* measurements from the Surface Radiation Budget (SURFRAD) network and calculate the *in situ* LST value with the emissivity estimated using the ASTER dataset, following standard procedures [4]. We used the annual mean spectral reflectance from the first seven spectral bands to represent the Earth’s surface. The total number of available *in situ* sites is limited (fewer than 10 in the U.S.), and in this study, each Landsat and GOES-16 dataset has one site within the coverage.

Experimental Setup We conducted all experiments using PyTorch. The model was trained with the Adam optimizer at a learning rate of 0.1 for 500 epochs. For the Landsat datasets, validation was performed exclusively using *in situ* measurements due to the limited temporal coverage. For the GOES-16 dataset, we employed both *in situ* validation and hold-out validation, with 20% of the valid observed data held out for testing.

4.2. Results on Landsat Data

The results on the Landsat datasets are summarized in Table 1 for the PSU site and in Table 2 for the BON site. This evaluation uses *in situ* measurements to assess all-weather LST reconstruction performance (a few days without valid *in situ* measurements were excluded, resulting in 362 days for PSU and 359 days for BON in 2023). At the PSU site, the proposed Physics-Guided CNN achieved the best overall performance (MAE = 2.75 K, RMSE = 3.44 K). Its bias was comparable to that of the 4-step ATC [38], an existing method, while achieving reconstructions across the full year (362 days), in contrast to only 48 days of the competitor. Another competing method published this year [6], based on lower-resolution inputs (1,000 m vs. 30 m), provided reconstruction for 229 days but yielded significantly inferior accuracy (RMSE = 4.93 K). Similarly, at the BON site, the proposed method outperformed all competitors across all evaluation metrics, achieving an MAE of 3.43 K and a bias of 0.59 K. Overall, the proposed method enables accurate, daily, all-weather LST reconstruction at very high spatial resolution (30 m).

Table 1. Comparison of Landsat temperature data reconstruction across different methods against *in situ* measurements (PSU).

Method	MAE (K)	RMSE (K)	Bias (K)	N	Resolution
4-step ATC [38]	–	3.79	0.12	48*	30 m
GEC-SEB [6]	–	4.93	-0.87	229*	1,000 m
ATC	3.97	5.19	0.66	362	30 m
ATC + ERA5	3.78	4.74	3.00	362	30 m
Naïve CNN	14.63	17.68	-14.36	362	30 m
Proposed	2.75	3.44	-0.16	362	30 m

*Existing methods are not daily for all-weather conditions and applied to selected days with observations only. For other days, existing methods reduce to ATC.

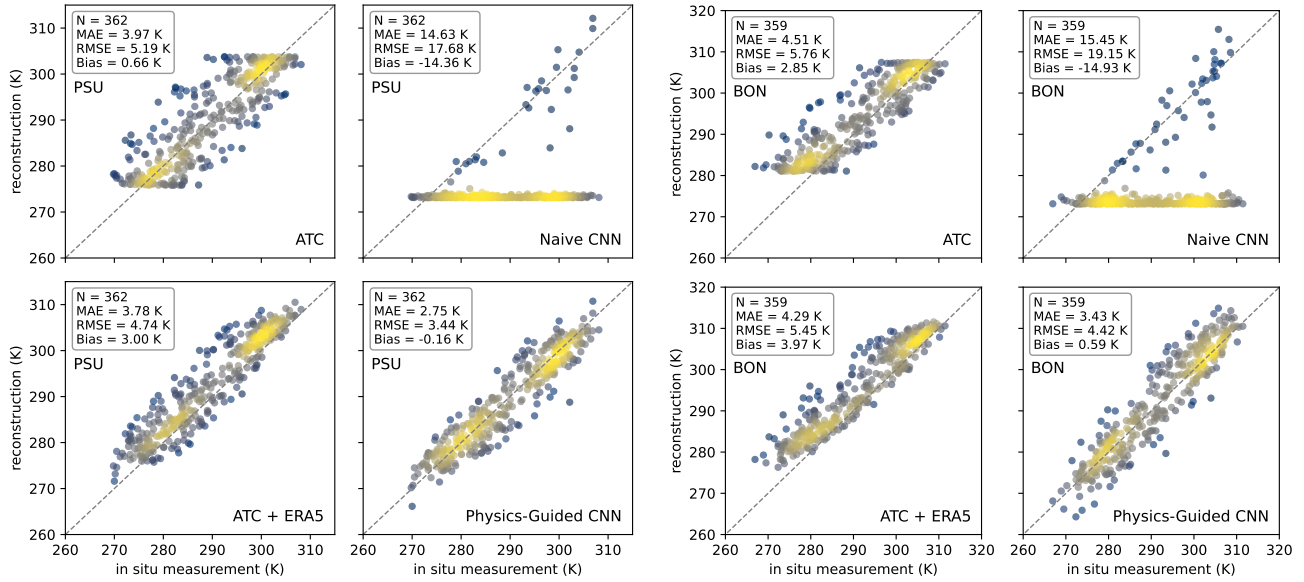


Figure 2. Comparison of reconstructed LST of Landsat (30 m) at PSU and BON sites from four models (ATC, Naïve CNN, ATC + ERA5, Physics-Guided CNN). Each subplot shows a density-colored scatter plot comparing *in situ* measurements (x-axis) with reconstruction (y-axis).

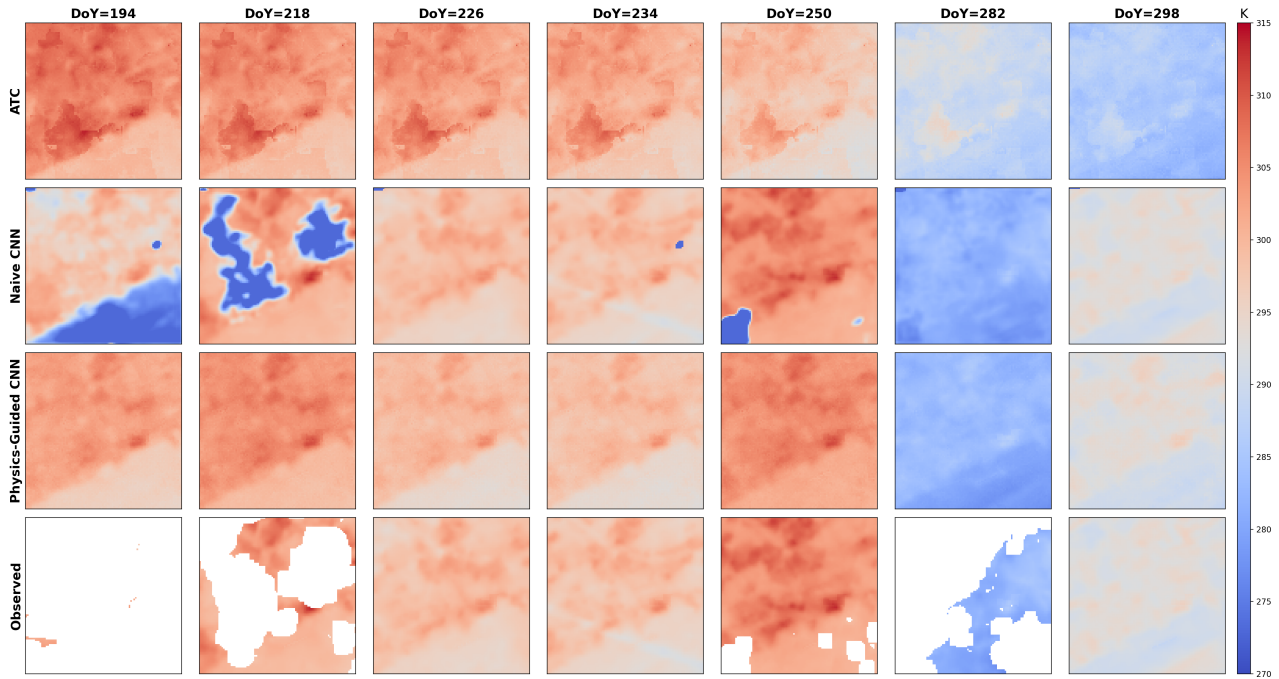


Figure 3. Reconstructed daily temperature maps at 30 m resolution from Landsat data.

We show the scatter plots of the ATC, naïve CNN, ATC+ERA5, and the proposed Physics-Guided CNN at the two sites in Fig. 2. The ATC reconstruction results are capped at a range of values, failing to capture temperature extremes at both the lower and higher ends. The naïve CNN

method can only fit the data on dates with observations; for the remaining dates, the predictions show an extreme overfitting trend. The sparse data density over the temporal domain renders the naïve CNN ineffective in this case. The ATC+ERA5 method alleviates some issues related to cap-

Table 2. Comparison of Landsat temperature data reconstruction across different methods against *in situ* measurements (BON).

Method	MAE (K)	RMSE (K)	Bias (K)	N	Resolution
4-step ATC [38]	–	4.95	1.29	20*	30 m
GEC-SEB [6]	–	5.51	0.67	222*	1,000 m
ATC	4.51	5.76	2.85	359	30 m
ATC + ERA5	4.29	5.45	3.97	359	30 m
Naïve CNN	15.45	19.15	-14.93	359	30 m
Proposed	3.43	4.42	0.59	359	30 m

*Existing methods are not daily for all-weather conditions and applied to selected days with observations only. For other days, existing methods reduce to ATC.

turing extremes, but a large bias still exists due to the lack of spatiotemporal correction. The proposed Physics-Guided CNN further reduces this bias by additionally accounting for spatiotemporal dependencies in the observational data.

Finally, we show the reconstructed LST surfaces from seven selected days using ATC, naïve CNN, and the proposed Physics-Guided CNN, along with the valid partial observations, in Fig. 3. The reconstruction surfaces are biased to be too high (DoY: 194, 226, 234) or too low (DoY: 298) for the ATC model, due to its inability to capture daily fluctuations. Unrealistic reconstruction results are sometimes obtained (DoY: 194, 218, 250) using the naïve CNN due to its poor generalization ability. In contrast, the proposed Physics-Guided CNN consistently outperformed the competitors and achieved better results.

4.3. Results on GOES-16 Data

For the two datasets on GOES-16 data reconstruction, we first show *in situ* validation through the scatter plots with the PSU and BON sites in Fig. 4. The observed pattern is similar to the Landsat reconstruction. For the PSU site, the proposed method achieved the best performance (RMSE = 3.72 K, N = 362). Notably, even the valid, cloud-free satellite observations (RMSE = 4.63 K, N = 145) yield worse results compared to the reconstruction. Similar pattern is observed over the BON site, with the Physics-Guided CNN achieving the best result (RMSE = 3.07 K) compared to the cloud-free observations (RMSE = 4.84 K), the ATC model (RMSE = 5.62 K), the Naïve CNN (RMSE = 7.64 K) and the ATC + ERA5 method (RMSE = 3.65 K).

The better agreement between the reconstruction results with the *in situ* data compared to the valid-observed cloud-free data is likely due to better agreement between LST *in situ* measurements and overhead observations under cloudy conditions [29]. Additionally, the imperfect match between the *in situ* data and GOES-16 is likely caused by differences in field of view. The *in situ* site has a field of view of approximately 70×70 m, while the GOES-16 observations have a spatial resolution of about 2 km. Due to the averaging effect from coarser-resolution data, GOES-16 measurements tend to be lower than the *in situ* values—a direction that is expected. Nonetheless, the results have demonstrated that

Table 3. Result comparison on the GOES-16 datasets.

Method	Training Data			Test Data		
	MAE (K)	RMSE (K)	Bias (K)	MAE (K)	RMSE (K)	Bias (K)
First GOES-16 Dataset (East Coast including PSU)						
ATC	3.63	4.91	0.05	3.75	5.02	0.05
Naïve CNN	3.67	15.85	0.59	3.68	15.91	0.60
ATC + ERA5	1.81	2.74	-0.27	1.83	2.70	-0.21
Proposed	1.77	2.64	-0.18	1.80	2.67	-0.18
Second GOES-16 Dataset (Midwest including BON)						
ATC	4.08	5.43	-0.16	4.21	5.55	-0.16
Naïve CNN	3.72	6.83	-2.04	3.71	6.83	-2.03
ATC + ERA5	2.12	2.98	-0.34	2.09	2.82	-0.27
Proposed	1.76	2.48	-0.27	1.79	2.51	-0.27

the proposed method has achieved indistinguishable reconstruction compared with the observational data.

We further evaluate the proposed method using the held-out 20% test data, as shown in Table 3. The proposed method achieves the best results among all, reducing the MAE metric by 52% compared to the ATC method, from 3.75 K to 1.80 K. By observing the difference between training and testing evaluation metrics, we notice that the physics-guided approaches (ATC, ATC + ERA5, and the proposed method) tend to have good generalization ability, with comparable results between the training and testing data (e.g., for the proposed method, the MAE is 1.77 K for training and 1.80 K for testing). As we are optimizing the F1 loss (MAE), although the naïve CNN approach has a low loss, it fails to generalize and has very high RMSE values compared to other methods.

To further demonstrate this difference, we show the reconstructed temperature maps from seven selected days, for the first GOES-16 dataset, in Fig. 5, covering a wide range of temperature values across all four seasons, and comparing the three methods (ATC, naïve CNN, Physics-Guided CNN). Among these, the proposed Physics-Guided CNN consistently demonstrates the best visual performance. For instance, in Fig. 5e, comparison with the near-complete valid observations reveals that the ATC model tends to overestimate temperatures in the northern region, while the naïve CNN produces unrealistic patterns, primarily due to its heavy reliance on spatiotemporal dependencies. In contrast, the Physics-Guided CNN generates reconstructions that most closely align with the valid observations. On days without valid observations (e.g., Fig. 5f), the naïve CNN fails entirely to reconstruct the temperature field, as this becomes an extrapolation task—a scenario in which conventional deep learning models that rely heavily on statistical patterns often fall short. In contrast, physics-guided approaches, including both the ATC model and the Physics-Guided CNN, successfully produce plausible reconstructions. In other cases, the Physics-Guided CNN consistently outperforms the ATC model and other compositors.

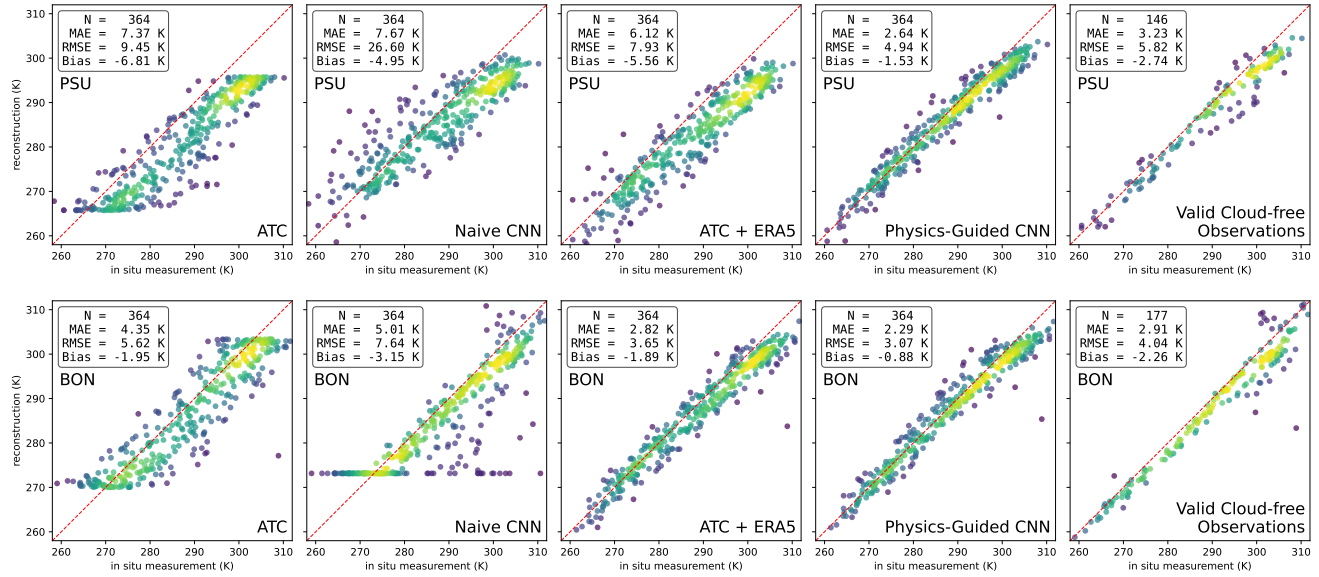


Figure 4. Comparison of GOES-16 LST reconstructions from ATC, Naive CNN, ATC + ERA5, Physics-Guided CNN, and cloud-free observations at PSU and BON. Each subplot shows a density-colored scatter plot of *in situ* (x-axis) vs. reconstructed values (y-axis).

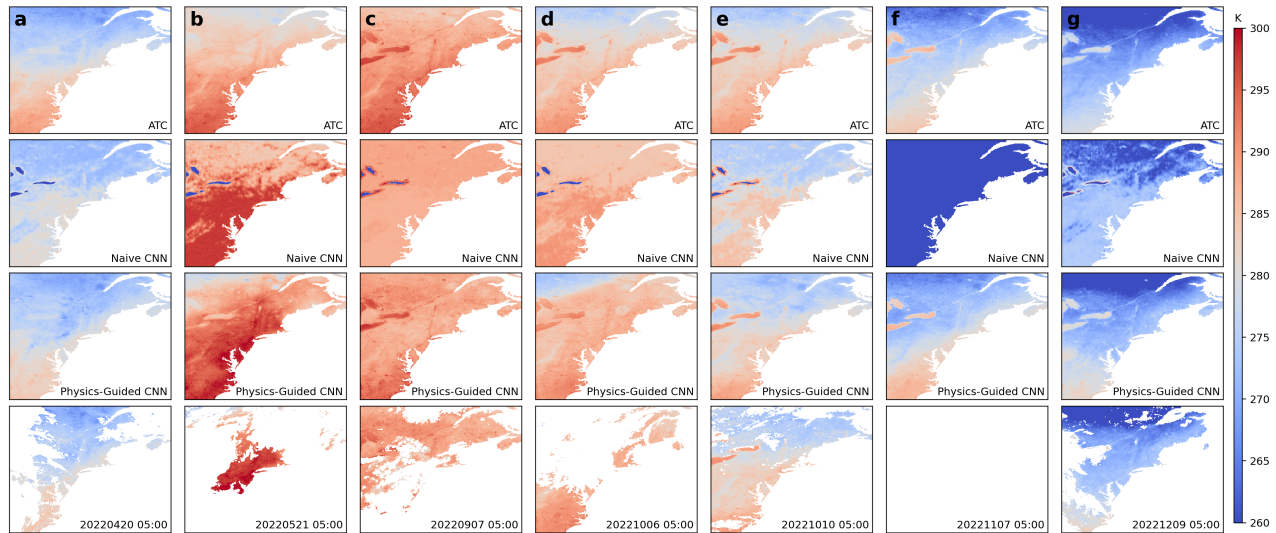


Figure 5. Reconstructed temperature maps of the GOES-16 data (East Coast including PSU).

5. Conclusion

In this study, we presented a physics-guided deep learning framework for land surface temperature data reconstruction, achieving, for the first time, end-to-end temperature reconstruction based on a single deep learning model. The proposed framework was tested on four datasets from two satellites with different resolutions (30 m daily for Landsat, 2 km hourly for GOES-16), achieving 2–3 K accuracy in hold-out validation for GOES-16, consistent with sensor’s accuracy, showing state-of-the-art performance on 30 m reconstruc-

tion for the Landsat datasets, and achieving better agreement with *in situ* measurements on the GOES-16 data due to its ability to estimate cloud-covered temperature. Central to the proposed framework is the encoding of a physics-guided model within a CNN and the usage of ERA5 reanalysis data to capture daily fluctuations. We used surface reflectance to reconstruct the residual surfaces. Recent Earth foundation models have shown promise to generate deep representations that better capture the Earth surface [25, 39]. Integrating the deep representations can further enhance temperature reconstruction under the proposed framework.

Acknowledgments

S.L. was supported by a USC Dana and David Dornsife College of Letters, Arts and Sciences / Graduate School Fellowship at the University of Southern California. L.Z. was supported by the National Institutes of Health (NIH) under grant numbers P20HL176204, P30ES007048 and R01ES031590. The views and conclusions expressed in this study are those of the authors and do not necessarily represent the official policies or endorsements of the NIH. The authors acknowledge the open data policies and practices of the National Oceanic and Atmospheric Administration (NOAA), the National Aeronautics and Space Administration (NASA), and the European Centre for Medium-Range Weather Forecasts (ECMWF).

References

- [1] Hicham Bahi, Lahouari Bounoua, Anas Sabri, Abderrazak Bannari, Anass Malah, and Hassan Rhinane. A new thermal fusion method to downscale land surface temperature to finer spatial resolution using sentinel-msi and landsat-oli/tirs imagery. *Remote Sensing Applications: Society and Environment*, 37:101519, 2025. 2
- [2] Christopher Beale, Hamid Norouzi, Zahra Sharifnezhadazi, Abdou Rachid Bah, Peng Yu, Yunyue Yu, Reginald Blake, Anna Vaculik, and Jorge Gonzalez-Cruz. Comparison of diurnal variation of land surface temperature from goes-16 abi and modis instruments. *IEEE Geoscience and Remote Sensing Letters*, 17(4):572–576, 2019. 5
- [3] Nils Bochow, Anna Poltronieri, Martin Rypdal, and Niklas Boers. Reconstructing historical climate fields with deep learning. *Science Advances*, 11(14):eadp0558, 2025. 2
- [4] Jie Cheng, Shunlin Liang, Yunjun Yao, and Xiaotong Zhang. Estimating the optimal broadband emissivity spectral range for calculating surface longwave net radiation. *IEEE Geoscience and Remote Sensing Letters*, 10(2):401–405, 2012. 5
- [5] Lirong Ding, Ji Zhou, Zhao-Liang Li, Jin Ma, Chunxiang Shi, Shuai Sun, and Ziwei Wang. Reconstruction of hourly all-weather land surface temperature by integrating reanalysis data and thermal infrared data from geostationary satellites (rtg). *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–17, 2022. 1
- [6] Wenhui Du, Zhao-Liang Li, Zhihao Qin, Jinlong Fan, Xi-angyang Liu, Chunliang Zhao, and Kun Cao. Reconstruction of cloudy land surface temperature by combining surface energy balance theory and solar-cloud-satellite geometry. *IEEE Transactions on Geoscience and Remote Sensing*, 2025. 5, 7
- [7] Filippo Giorgi and Roni Avissar. Representation of heterogeneity effects in earth system modeling: Experience from land surface modeling. *Reviews of Geophysics*, 35(4):413–437, 1997. 1
- [8] James Hansen, Reto Ruedy, Mki Sato, and Ken Lo. Global surface temperature change. *Reviews of Geophysics*, 48(4), 2010. 1
- [9] Yutong He, Dingjie Wang, Nicholas Lai, William Zhang, Chenlin Meng, Marshall Burke, David Lobell, and Stefano Ermon. Spatial-temporal super-resolution of satellite imagery via conditional pixel synthesis. *Advances in Neural Information Processing Systems*, 34:27903–27915, 2021. 2
- [10] Nobuyuki Hirahara, Motoharu Sonogashira, Hidekazu Kasahara, and Masaaki Iiyama. Denoising and inpainting of sea surface temperature image with adversarial physical model loss. In *Asian Conference on Pattern Recognition*, pages 339–352. Springer, 2019. 2
- [11] James R Irons, John L Dwyer, and Julia A Barsi. The next landsat satellite: The landsat data continuity mission. *Remote Sensing of Environment*, 122:11–21, 2012. 1
- [12] Aolin Jia, Shunlin Liang, Dongdong Wang, Kanishka Mallick, Shugui Zhou, Tian Hu, and Shuo Xu. Advances in methodology and generation of all-weather land surface temperature products from polar-orbiting and geostationary satellites: A comprehensive review. *IEEE Geoscience and Remote Sensing Magazine*, 2024. 2
- [13] Christopher Kadow, David Matthew Hall, and Uwe Ulbrich. Artificial intelligence reconstructs missing climate information. *Nature Geoscience*, 13(6):408–413, 2020. 2
- [14] Michael D King, Steven Platnick, W Paul Menzel, Steven A Ackerman, and Paul A Hubanks. Spatial and temporal distribution of clouds observed by modis onboard the terra and aqua satellites. *IEEE transactions on geoscience and remote sensing*, 51(7):3826–3852, 2013. 1
- [15] Kelly G Laraby and John R Schott. Uncertainty estimation method and landsat 7 global validation for the landsat surface temperature product. *Remote Sensing of Environment*, 216: 472–481, 2018. 5
- [16] Chan Li, Penghai Wu, Si-Bo Duan, Yixuan Jia, Shuai Sun, Chunxiang Shi, Zhixiang Yin, Huifang Li, and Huanfeng Shen. Lfsr: Low-resolution filling then super-resolution reconstruction framework for gapless all-weather modis-like land surface temperature generation. *Remote Sensing of Environment*, 319:114637, 2025. 2
- [17] Wan Li, Li Ni, Zhao-liang Li, Si-Bo Duan, and Hua Wu. Evaluation of machine learning algorithms in spatial downscaling of modis land surface temperature. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2299–2307, 2019. 1
- [18] Yuan Li, Zhi-Hao Lin, David Forsyth, Jia-Bin Huang, and Shenlong Wang. Climatenerf: Extreme weather synthesis in neural radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3227–3238, 2023. 2
- [19] Zhao-Liang Li, Hua Wu, Si-Bo Duan, Wei Zhao, Huazhong Ren, Xiangyang Liu, Pei Leng, Ronglin Tang, Xin Ye, Jinshun Zhu, et al. Satellite remote sensing of global land surface temperature: Definition, methods, products, and applications. *Reviews of Geophysics*, 61(1), 2023. 1, 2
- [20] Shengjie Liu and Lu Zhang. Deep feature gaussian processes for single-scene aerosol optical depth reconstruction. *IEEE Geoscience and Remote Sensing Letters*, 2024. 2
- [21] Shengjie Liu, An-Min Wu, and Hung Chak Ho. Spatial variability of diurnal temperature range and its associations with local climate zone, neighborhood environment and mortality in los angeles. *Urban Climate*, 49:101526, 2023. 1

- [22] Shengjie Liu, Siqin Wang, and Lu Zhang. Daily land surface temperature reconstruction in landsat cross-track areas using deep ensemble learning with uncertainty quantification. *arXiv preprint arXiv:2502.14433*, 2025. 1, 2
- [23] David T Lloyd, Aaron Abela, Reuben A Farrugia, Anthony Galea, and Gianluca Valentino. Optically enhanced super-resolution of sea surface temperature using deep learning. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–14, 2021. 2
- [24] Di Long, La Yan, Liangliang Bai, Caijin Zhang, Xueying Li, Huimin Lei, Hanbo Yang, Fuqiang Tian, Chao Zeng, Xianyong Meng, et al. Generation of modis-like land surface temperatures under all-weather conditions based on a data fusion approach. *Remote Sensing of Environment*, 246:111863, 2020. 2
- [25] Valerio Marsocci, Yuru Jia, Georges Le Bellier, David Kerekes, Liang Zeng, Sebastian Hafner, Sebastian Gerard, Eric Brune, Ritu Yadav, Ali Shibli, et al. Pangaea: A global and inclusive benchmark for geospatial foundation models. *arXiv preprint arXiv:2412.04204*, 2024. 8
- [26] Stephan Rasp, Michael S Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018. 2
- [27] Iain Rolland, Sivasakthy Selvakumaran, Shaikh Fairul Edros Ahmad Shaikh, Perrine Hamel, and Andrea Marinoni. Improving land surface temperature estimation in cloud cover scenarios using graph-based propagation. *Geophysical Research Letters*, 51(23):e2024GL108263, 2024. 2
- [28] 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). *Journal of Operational Meteorology*, 6(4):33–46, 2018. 5
- [29] Christelle Vancutsem, Pietro Ceccato, Tufa Dinku, and Stephen J Connor. Evaluation of modis land surface temperature data to estimate air temperature in different ecosystems over africa. *Remote Sensing of Environment*, 114(2): 449–465, 2010. 7
- [30] Chenghao Wang, Zhi-Hua Wang, Chuyuan Wang, and Soe W Myint. Environmental cooling provided by urban trees under extreme heat and cold waves in us cities. *Remote Sensing of Environment*, 227:28–43, 2019. 1
- [31] Martin Wegmann and Fernando Jaume-Santero. Artificial intelligence achieves easy-to-adapt nonlinear global temperature reconstructions using minimal local data. *Communications Earth & Environment*, 4(1):217, 2023. 2
- [32] Michael C Wimberly, Kirsten M de Beurs, Tatiana V Loboda, and William K Pan. Satellite observations and malaria: new opportunities for research and applications. *Trends in Parasitology*, 37(6):525–537, 2021. 1
- [33] Penghai Wu, Zhixiang Yin, Chao Zeng, Si-Bo Duan, Frank-Michael Göttsche, Xiaoshuang Ma, Xinghua Li, Hui Yang, and Huanfeng Shen. Spatially continuous and high-resolution land surface temperature product generation: A review of reconstruction and spatiotemporal fusion techniques. *IEEE Geoscience and Remote Sensing Magazine*, 9(3):112–137, 2021. 1, 2
- [34] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5505–5514, 2018. 2
- [35] Tao Zhang, Yuyu Zhou, Zhengyuan Zhu, Xiaoma Li, and Ghassem R Asrar. A global seamless 1 km resolution daily land surface temperature dataset (2003–2020). *Earth System Science Data Discussions*, 2021:1–16, 2021. 2
- [36] Xiaodong Zhang, Ji Zhou, Shunlin Liang, and Dongdong Wang. A practical reanalysis data and thermal infrared remote sensing data merging (rtm) method for reconstruction of a 1-km all-weather land surface temperature. *Remote Sensing of Environment*, 260:112437, 2021. 1
- [37] Zhenwei Zhang and Qingyun Du. Hourly mapping of surface air temperature by blending geostationary datasets from the two-satellite system of goes-r series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:111–128, 2022. 1
- [38] Xiaolin Zhu, Si-Bo Duan, Zhao-Liang Li, Penghai Wu, Hua Wu, Wei Zhao, and Yonggang Qian. Reconstruction of land surface temperature under cloudy conditions from landsat 8 data using annual temperature cycle model. *Remote Sensing of Environment*, 281:113261, 2022. 2, 5, 7
- [39] Xiao Xiang Zhu, Zhitong Xiong, Yi Wang, Adam J Stewart, Konrad Heidler, Yuanyuan Wang, Zhenghang Yuan, Thomas Dujardin, Qingsong Xu, and Yilei Shi. On the foundations of earth and climate foundation models. *arXiv preprint arXiv:2405.04285*, 2024. 8