Environmental Data Science (2026), 5: e32, 1–10
doi:10.1017/eds.2026.10052

CAMBRIDGE UNIVERSITY PRESS

APPLICATION PAPER

Downscaling land surface temperature data using edge detection and block-diagonal Gaussian process regression

Sanjit Dandapanthula[1], Margaret Johnson[2], Madeleine Pascolini-Campbell[2], Glynn Hulley[2] and Mikael Kuusela[1]

[1]Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA, USA
[2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA
Corresponding author: Sanjit Dandapanthula; Email: sanjitd@cmu.edu

Received: 31 May 2026 UTC; **Revised:** 31 May 2026 UTC; **Accepted:** 04 June 2026 UTC

Keywords: change of support; data fusion; ECOSTRESS; evapotranspiration; Landsat; spatial statistics

Abstract

Accurate and high-resolution estimation of land surface temperature (LST) is crucial in estimating evapotranspiration, a measure of plant water use and a central quantity in agricultural applications. In this work, we develop a novel statistical method for downscaling LST data obtained from NASA's ECOSTRESS mission, using high-resolution data from the Landsat 8 mission as a proxy for modeling agricultural field structure. Using the Landsat data, we identify the boundaries of agricultural fields through edge detection techniques, allowing us to capture the inherent block structure present in the spatial domain. We propose a block-diagonal Gaussian process (BDGP) model that captures the spatial structure of the agricultural fields, leverages independence of LST across fields for computational tractability, and accounts for the change of support present in ECOSTRESS observations. We use the resulting BDGP model to perform Gaussian process regression and obtain high-resolution estimates of LST from ECOSTRESS data, along with uncertainty quantification. Our results demonstrate the practicality of the proposed method in producing reliable high-resolution LST estimates, with potential applications in agriculture, urban planning, and climate studies.

Impact Statement

Land surface temperature (LST) is an important input for estimating evapotranspiration, a central quantity in agriculture. We develop a statistical method for downscaling the blurry but frequent LST data from NASA's ECOSTRESS mission to a finer scale, using high-resolution imagery from the Landsat satellite, both to delineate agricultural fields and to sharpen the LST estimates within them. By modeling each field as an independent unit, our method remains computationally tractable and provides principled uncertainty quantification, with potential applications in agriculture, urban planning, and climate science.

1. Introduction

Obtaining high-resolution estimates of land surface temperature (LST) given observations from remote sensing satellites is an important problem in environmental statistics; for instance, accurate estimation of LST is crucial for understanding the *evapotranspiration* (ET), a central quantity in Earth's water cycle

This research article was awarded Open Data and Open Materials badges for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

(Fisher et al., 2017). Spatial modeling is also used in this context for gap-filling between observed data, but we do not focus on this application in this work. Current attempts to build statistical models for LST data have not been able to sufficiently account for the complex spatial structure of the data, especially around agricultural fields where estimates of LST are most crucial. Traditional spatial models, such as stationary Gaussian processes, fail to capture the abrupt changes in LST that occur at the boundaries of agricultural fields, which can lead to oversmoothing and inaccurate predictions (Paciorek and Schervish, 2006; Kumar and Chakraborty, 2026). Therefore, there is a need for more sophisticated statistical methods that can effectively model the spatial structure present in LST data.

ET is the sum of evaporation and plant transpiration from the Earth's land surface to the atmosphere (Fisher et al., 2017). It is a key component of the water cycle and is a critical variable in the study of climate change, agriculture, and water resource management. ET is influenced by several factors, including LST, solar radiation, wind speed, and humidity. In particular, LST is a key driver of ET as it controls the amount of energy available for evaporation and transpiration. Accurate estimation of LST is therefore essential for predicting ET.

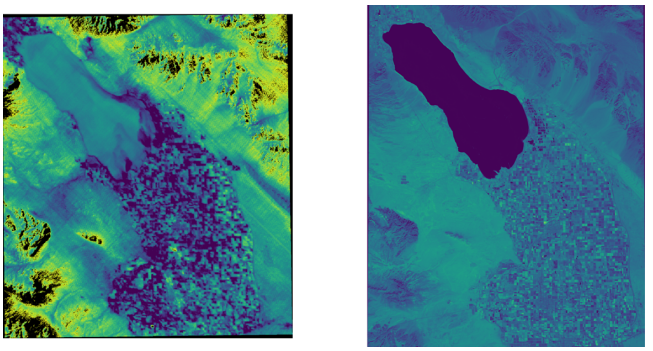
The ECOsystem Spaceborne Thermal Radiometer Experiment on Space Station (ECOSTRESS) radiometer was deployed on the International Space Station (ISS) in June 2018 to measure LST and emissivity, using an infrared sensor (Fisher et al., 2020). The *overpass return frequency*, or revisit rate, of a sensor is the frequency with which that sensor passes over a given location on Earth. The ECOSTRESS mission has a variable overpass return frequency (between 1 and 5 days) and currently provides the highest spatial resolution spaceborne estimates of LST at ~70 m resolution. However, the effective resolution of LST images from ECOSTRESS has been shown to be significantly lower in practice (Holmes et al., 2024), likely due to blur and artifacts resulting from sensor damage.

The Landsat program is a series of Earth-observing satellite missions jointly managed by NASA and the US Geological Survey (USGS) that has been providing high-resolution imagery of the Earth's surface since 1972 (Roy et al., 2014). Landsat satellites are equipped with sensors that capture data in multiple spectral bands, including visible, near-infrared, and thermal infrared, allowing for detailed analysis of land cover, vegetation health, and surface temperature. Landsat 8/9 instruments provide data in the visible to shortwave infrared (VSWIR) at a higher resolution (30 m) than ECOSTRESS, but with a lower overpass return frequency (16 days). Landsat 8/9 Thermal Infrared Sensors (TIRS) have fewer thermal bands than ECOSTRESS, resulting in potentially lower fidelity estimates of LST, and are at nominally lower spatial resolution (100 m).

In our work, we leverage the higher-resolution Landsat 8/9 VSWIR data as a proxy for modeling the agricultural field structure present in the ECOSTRESS LST data. This task, often referred to as *data fusion*, is of broad interest in the remote sensing community, in far more general settings than just LST data (Khaleghi et al., 2013; Samadzadegan et al., 2025). We plot sample images from both missions in Figure 1, which reflects the type of agricultural field structure that is the focus of our work. Note that the ECOSTRESS and Landsat images are not in the same units and are not taken at the same time; this figure is only meant to illustrate the difference in spatial resolution.

There is a long line of work studying downscaling procedures for land surface temperature data by direct regression of the high-resolution images on their low-resolution counterparts using random forest regression (Hutengs and Vohland, 2016; Li et al., 2018; Njuki et al., 2020) and neural networks (Bindhu et al., 2013; Wu et al., 2022; Li et al., 2025). By contrast, our statistical approach using Gaussian process regression is new in this literature. Some key advantages of our approach are the better interpretability of the underlying statistical model, weaker modeling assumptions, which may improve robustness, and the ability to obtain principled uncertainties for the resulting reconstruction.

More similar to our approach, Garrigues et al. (2007) proposed the mosaic process model for block-structured spatial data, which allows for different covariance structures in different regions of the spatial domain. However, the complexity of their model makes it difficult to scale to larger datasets, such as those encountered in remote sensing applications. Their model is also restricted to blocks that are bounded by straight lines through the spatial domain, which is not ideal for modeling agricultural fields with irregular



(a) ECOSTRESS (70 m resolution) (b) Landsat (30 m resolution)

Figure 1. *Comparison of resolution between ECOSTRESS and Landsat missions over the Salton Sea region in southern California.*

shapes. Furthermore, the tessellation in the mosaic process model is random, which may lead it to be underconstrained for spatial mapping.

Our main contributions in this work are as follows:

- We propose a novel statistical method for downscaling LST data from the ECOSTRESS mission using high-resolution data from the Landsat mission as a proxy for modeling agricultural field structure.
- We develop a block-diagonal Gaussian process (BDGP) model that captures the spatial structure of the agricultural fields and leverages independence of LST across fields for computational tractability.
- We demonstrate the practicality of the proposed method in producing high-resolution LST estimates along with uncertainty quantification.

We outline our methodology in Section 2; the edge detection method is described in Section 2.1, the Fourier regression approach for modeling the mean function is described in Section 2.2, the parameter estimation procedure is described in Section 2.3, and the kriging and uncertainty quantification method is described in Section 2.4. We present our results in Section 3 and conclude with a discussion in Section 4.

2. Methods

It is common to model spatial data using a Gaussian process (GP) (Gelfand et al., 2010). In order to account for the agricultural field structure present in LST data, we model LST using a block-diagonal Gaussian process with a squared exponential kernel. Suppose we have a partition, $\mathcal{R} := \{R_1, \ldots, R_N\}$, of $\mathcal{X}$ into N regions. Given a variance parameter $\boldsymbol{\sigma} := (\sigma_1, \ldots, \sigma_N) \in (0, \infty)^N$ and length scale parameter $\boldsymbol{\ell} := (\ell_1, \ldots, \ell_N) \in (0, \infty)^N$, we define the block-diagonal squared exponential covariance kernel with parameters $(\boldsymbol{\sigma}, \boldsymbol{\ell})$ as $k(x, x') = \sum_{r=1}^{N} k_r(x, x')$ where:

$$k_r(x, x') = \begin{cases} \sigma_r^2 \exp\left(-\|x - x'\|_2^2 / \left(2\ell_r^2\right)\right), & x, x' \in R_r \\ 0, & \text{otherwise.} \end{cases}$$

This kernel is stationary within each region. Given any mean function $\mu : \mathcal{X} \to \mathbb{R}$, we say that a block diagonal Gaussian process with parameters $(\mu, \boldsymbol{\sigma}, \boldsymbol{\ell}, \mathcal{R})$ (denoted $\text{BDGP}(\mu, \boldsymbol{\sigma}, \boldsymbol{\ell}, \mathcal{R})$) is a GP with a mean function μ and covariance kernel k of the form above.

We assume that at a given time, the underlying LST field is a BDGP $f^{\text{ES}} := \text{BDGP}\left(\boldsymbol{\mu}^{\text{ES}}, \boldsymbol{\sigma}^{\text{ES}}, \boldsymbol{\ell}^{\text{ES}}, \mathcal{R}\right)$ and that we observe samples from the blurry $\tilde{f}^{\text{ES}} := \left(f^{\text{ES}} * \kappa_{\text{blur}}\right) + f_{\text{sensor}}$, where $*$ denotes a two-dimensional

convolution. Here, κ_{blur} is the density of $\mathcal{N}\left(0, \sigma^2_{\text{blur}} I\right)$ (accounting for the change of support caused by the finite resolution of ECOSTRESS) and $f_{\text{sensor}} \sim \mathcal{N}\left(0, \sigma^2_{\text{sensor}} I\right)$ accounts for sensor noise within ECOSTRESS. We further assume that the Landsat data are generated from $f^{\text{LS}} := \text{BDGP}\left(\boldsymbol{\mu}^{\text{LS}}, \boldsymbol{\sigma}^{\text{LS}}, \boldsymbol{\ell}^{\text{LS}}, \mathcal{R}\right)$. We make the simplifying assumption that $\boldsymbol{\ell}^{\text{ES}} = \boldsymbol{\ell}^{\text{LS}}$ in order to ensure that the parameters are identifiable in this model.

Then, we outline our LST downscaling procedure as follows:

(1) *Edge detection*: Use edge detection methods on an averaged high-resolution Landsat image to identify boundaries between different agricultural fields, thereby estimating the partition $\mathcal{R}$.
(2) *Parameter estimation*: Fixing $\mathcal{R}$, fit the remaining parameters of the BDGP model to the Landsat and ECOSTRESS data by maximum likelihood estimation.
(3) *Kriging and uncertainty quantification*: Use the fitted BDGP model for ECOSTRESS to do kriging and obtain high-resolution predictions of the residual LST field ($f^{\text{ES}} - \mathbb{E}\left[f^{\text{ES}}\right]$), along with uncertainty quantification.

An overview of our downscaling pipeline is depicted in Figure 2.

2.1. Edge detection

First, we average six separate high-resolution Landsat images and use edge detection methods on the averaged image to identify boundaries between different agricultural fields. In our case, we use Meta's Segment Anything Model (SAM), which is a state-of-the-art vision transformer (ViT) based segmentation model providing an over-segmentation of the input image—the model may produce partially or fully overlapping segments (Kirillov et al., 2023). We experimented with several other methods, including the Canny, Sobel, and Laplace edge detection algorithms, but we found the SAM to work best in practice. Next, we post-process the over-segmented image to merge overlapping regions together using an algorithm described in Supplementary Appendix E. This post-processing step ensures that the final set of segments contains no overlaps. At this point, we have an estimated partition $\mathcal{R}$ of the spatial domain into agricultural fields.

2.2. Fourier regression

Next, we fit $\boldsymbol{\mu}^{\text{LS}}$ separately for each pixel $x = (i, j)$ using a Fourier regression with annual frequency:

$$\mu^{\text{LS}}_{ij}(t) = (\beta_0)_{ij} + (\beta_1)_{ij} \cos(2\pi t/365) + (\beta_2)_{ij} \sin(2\pi t/365),$$

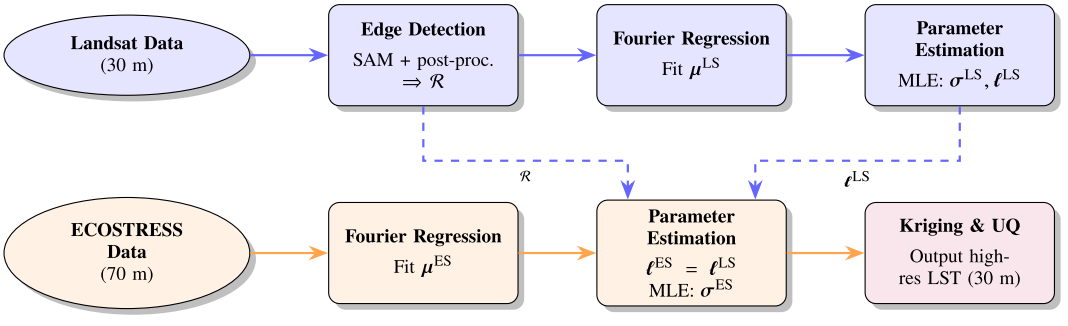


Figure 2. *Overview of the proposed downscaling pipeline. The Landsat data are shown in blue, the ECOSTRESS data are shown in orange, and the output of the algorithm is shown in purple.*

where t is in units of days. We use ordinary least squares to fit the coefficients. We only use an annual frequency here because the Landsat satellite has a *sun-synchronous* orbit, meaning that the satellite orbits over the same spot on Earth at the same time of day every 16 days. As a result, we are unable to model the diurnal frequency for these data. Next, we fit the blurred mean $\boldsymbol{\mu}^{\text{ES}} * \kappa_{\text{blur}}$ separately for each pixel (i,j) using a Fourier regression with annual and diurnal frequencies:

$$\left(\mu^{\text{ES}} * \kappa_{\text{blur}}\right)_{ij}(t) = (\beta_0)_{ij} + (\beta_1)_{ij}\cos(2\pi t/365) + (\beta_2)_{ij}\sin(2\pi t/365) + (\beta_3)_{ij}\cos(2\pi t) + (\beta_4)_{ij}\sin(2\pi t)$$

where t is in units of days. Note that we are fitting the blurred mean function here since we only observe the blurred data from ECOSTRESS. At last, we subtract the fitted mean functions from the observed data, so that we may model the residual data by zero-mean GPs. We see in Figure 1 (Supplementary Appendix C) that a single Fourier term is sufficient to capture the annual and diurnal cycles in the data, so that we can model the residuals by a mean-zero BDGP. Similarly, we depict the annual cycle fit for a random pixel in the Landsat data in Figure 2 (Supplementary Appendix C). The fitted cycles roughly capture the seasonal variations in land surface temperature, justifying our use of Fourier regression to model the mean function.

2.3. Parameter estimation

Since the uncertainties listed for the ECOSTRESS observations are all on the order of 0.1 degrees Kelvin, it is reasonable to choose $\sigma_{\text{sensor}} = 0.1$. From Holmes et al. (2024), we know that the point spread function of ECOSTRESS approximately has a full-width at half-maximum (FWHM) of 160 m in each direction, which is roughly twice the nominal pixel size of 70 m. From the FWHM, we can calculate the standard deviation by $\sigma_{\text{blur}} = (160/70)/\left(2\sqrt{2\ln(2)}\right) \approx 0.97$ in units of ECOSTRESS pixels. This computation is justified in Supplementary Appendix B.

Fixing $\mathcal{R}$, we now fit the remaining parameters of the BDGP model to the Landsat data. Note that each region can be fit independently here, which allows for easy parallelization and scalability to larger data. We do not include a temporal component in the covariance structure here, and instead sample far apart in time to reduce temporal correlation. As a result, we can treat these points as independent across time. We chose four sample dates t_1, t_2, t_3, t_4 that are ~3 months apart from each other. We write the likelihood of the observed data in the region R_r assuming that the observations are independent across time:

$$\mathcal{L}\left(\sigma_r^{\text{LS}}, \ell_r^{\text{LS}}; y\right) = \prod_{i=1}^{4} f_{\mathcal{N}(0,K_{R_r})}\left(y_{R_r,t_i}^{\text{LS}}\right),$$

where K_{R_r} is the covariance matrix of the BDGP restricted to the region r and y_{R_r,t_i}^{LS} is the observed residuals in the region R_r at time t_i. We estimate the parameters by maximizing this likelihood using the BFGS optimizer, adding a small nugget of $10^{-8}I$ to the covariance matrix for numerical stability.

Next, we fix $\mathcal{R}$ and assume that $\boldsymbol{\ell}^{\text{ES}} = \boldsymbol{\ell}^{\text{LS}}$. Although the Landsat and ECOSTRESS data measure different quantities, it is reasonable to think that the length scale of the residual field should be similar across both datasets since they are over the same geographical region with similar spatial variation. This assumption is required to ensure identifiability of the parameters in our model, due to ill-posedness caused by the change of support present in the ECOSTRESS data. Then, we fit $\boldsymbol{\sigma}^{\text{ES}}$ to the ECOSTRESS data by maximum likelihood estimation.

We now ignore the existence of blurring from other regions at the edges of each region, and fit each region independently; this is justified due to the small relative scale of the blur kernel compared to the regions. Under this simplifying assumption, a consistent estimator given a single observation is

$$\left(\widehat{\sigma}_r^{\text{ES}}\right)^2 = \left(\widehat{\text{Var}}\left(y_{R_r}^{\text{ES}}\right) - \sigma_{\text{sensor}}^2\right)\left(\frac{\left(\ell_r^{\text{ES}}\right)^2 + \sigma_{\text{blur}}^2}{\left(\ell_r^{\text{ES}}\right)^2}\right). \tag{2.1}$$

Here, we use $\mathrm{Var}\left(y_{R_r}^{\mathrm{ES}}\right)$ to denote the sample variance of the observed residuals in the region R_r, and we justify our use of this estimator in Supplementary Appendix A. We then aggregate our σ_r^{ES} estimates from each separate time point by taking the median across times.

2.4. Kriging and uncertainty quantification

Finally, we use the fitted BDGP model for ECOSTRESS to do kriging (on a 30 m grid) and obtain high-resolution predictions of the residual LST field. Although the regions are no longer independent after blurring, we may treat each region independently after considering a $4\sigma_{\mathrm{blur}}$-pixel neighborhood R_r' around each region R_r. First, enumerate the pixels as $x_1, \ldots, x_P$ and define k_* as the covariance vector between $f(x_*)$ and all blurry observations in R_r':

$$(k_*)_i = \mathrm{Cov}\left(f^{\mathrm{ES}}(x_*), \tilde{f}^{\mathrm{ES}}(x_i)\right) = (k(x_*, \cdot) * \kappa_{\mathrm{blur}}(\cdot))(x_i).$$

Next, let $K_{R_r'}$ be the covariance matrix of the BDGP restricted to the region R_r':

$$\left(K_{R_r'}\right)_{ij} = \mathrm{Cov}\left(\tilde{f}^{\mathrm{ES}}(x_i), \tilde{f}^{\mathrm{ES}}(x_j)\right) = \iint \kappa_{\mathrm{blur}}(x_i - u)\,\kappa_{\mathrm{blur}}(x_j - v)\,k(u, v)\,du\,dv.$$

Finally, let $y_{R_r'}^{\mathrm{ES}}$ be the observed residuals in region R_r'; namely, $\left(y_{R_r'}^{\mathrm{ES}}\right)_i = \tilde{f}^{\mathrm{ES}}(x_i)$ for $x_i \in R_r'$. Using this notation, the kriging formula (Cressie, 1991) gives us the conditional mean in the region R_r at location x_* as

$$\hat{f}_{R_r}^{\mathrm{ES}}(x_*) = k_*^{\top}\left(K_{R_r'} + \sigma_{\mathrm{sensor}}^2 I\right)^{-1} y_{R_r'}^{\mathrm{ES}}.$$

Repeating this computation for each region R_r gives us the full kriging prediction $\hat{f}^{\mathrm{ES}}$. By reconstructing each region separately, we can parallelize this step and maintain computational tractability as long as each region is not too large. In this work, we do not reconstruct the road region, since it is too large to reconstruct efficiently using our current implementation. Although the roads can be reconstructed by splitting into several overlapping subregions, the roads are not of primary interest in agricultural applications, so we leave this for future work. Additionally, because we still estimate the parameters for the road region, this does not affect the reconstructions in other regions.

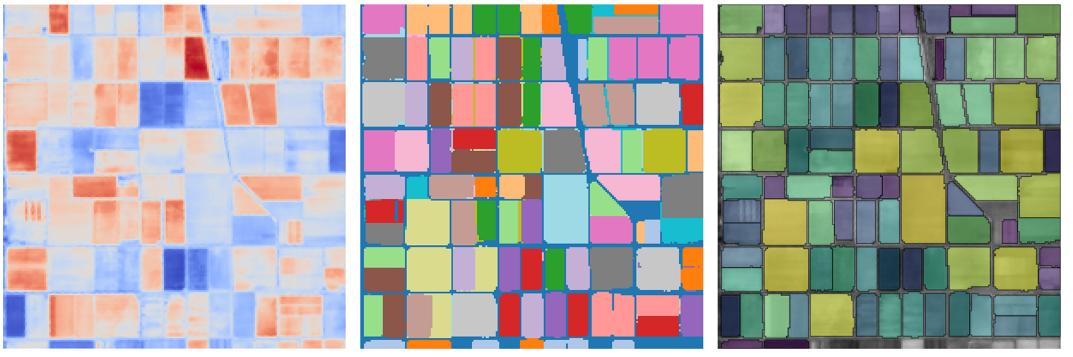
Next, to compute the variance of the kriging prediction at the location x_* in region R_r, we use the formula

$$\mathrm{Var}\left(\hat{f}^{\mathrm{ES}}(x_*)\right) = \mathrm{Var}\left(f^{\mathrm{ES}}(x_*)\right) - k_*^{\top}\left(K_{R_r'} + \sigma_{\mathrm{sensor}}^2 I\right)^{-1} k_*.$$

For instance, this formula can be found in Cressie (1991). Note that the resulting reconstruction is only of the *residual* field; to obtain the full LST field, we would need to add back a deconvolved mean function, which can be done more carefully using methods from spatial statistics. However, we leave this for future work.

3. Results

For our analysis, we consider an agricultural region of 6 km^2 near the Salton Sea in southern California (USA). We use data from the SR_B5 channel of the Landsat 8 satellite (Roy et al., 2014), which measures surface reflectance in the near-infrared spectrum of light. Our goal is to deconvolve an ECOSTRESS (Fisher et al., 2020) LST image taken in January 2021 using the Landsat data as a proxy for modeling the agricultural field structure. All code to run our experiments is provided in Dandapanthula (2026). We preprocess both datasets to be on a common grid with 30 m resolution. For all of our analysis, we manually choose images that are relatively cloud-free in our region of interest, and do not account for missingness. All of our parameter estimation and kriging procedures work with slight modifications to the code in the presence of some missing data. Similarly, large amounts of missingness can be handled by including a temporal component in the Gaussian process, but this is left for future work.



(a) Averaged Landsat image (b) SAM oversegmentation (c) Regions after post-processing

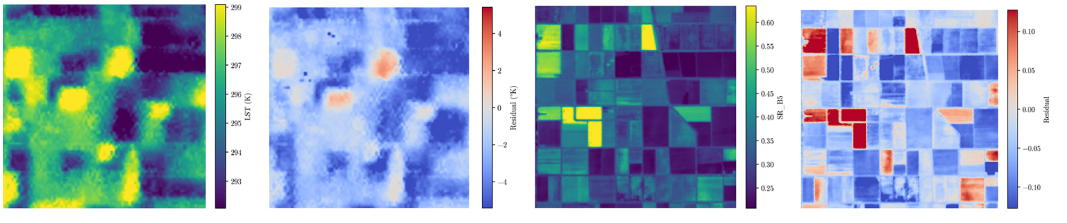
Figure 3. *Side-by-side comparison of (a) the averaged Landsat image, (b) the over-segmented image from SAM, and (c) the final regions after post-processing.*

We target an ECOSTRESS LST image taken in January 2021. The Landsat and ECOSTRESS mean functions are each fit using all available images of the region. The ECOSTRESS covariance parameter $\boldsymbol{\sigma}^{\text{ES}}$ and the Landsat covariance parameters $\left(\boldsymbol{\sigma}^{\text{LS}}, \boldsymbol{\ell}^{\text{LS}}\right)$ are each fit from four images, ~3 months apart, all taken before the target image to ensure that our fitted parameters generalize. Edge detection is performed on an average of six separate Landsat images, none of which coincides with the target date.

We begin by showing the results of the edge detection method on the average of 6 Landsat images (starting in November 2021 and spread out by ~2 months) in our region of interest in Figure 3. Hence, we confirm that the final regions capture the agricultural field structure relatively well, and there are no overlaps between different regions after post-processing.

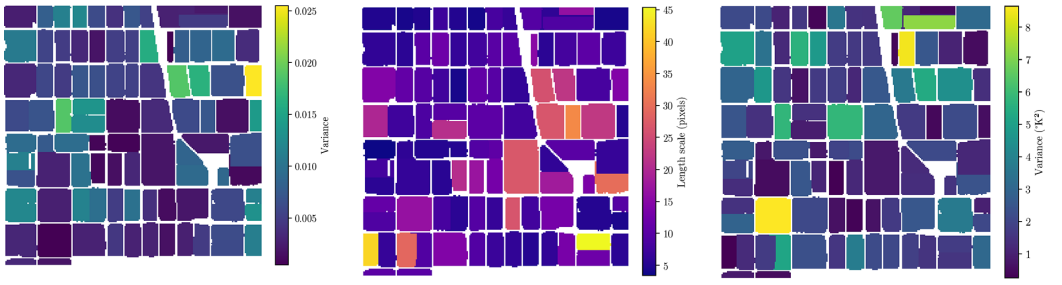
Next, we show our original ECOSTRESS image from January 2021 in Figure 4a (which we would like to deblur), along with the resulting residuals after subtracting the fitted mean function in Figure 4b. As expected, the residuals contain significant spatial structure that we can model using our BDGP approach. This day was particularly cold, leading to negative residuals across most of the image. To verify that our mean function is sensible for the Landsat data, we also show the residuals for a sample Landsat image in Figure 4d. Again, we see that the residuals contain significant spatial structure that we can model using our BDGP approach, and are close to zero.

Next, we show the fitted parameters for each region in Figure 5. We see that the length scales vary significantly across regions, justifying our use of a block-diagonal GP to model the agricultural field structure. Furthermore, the optimization procedure appears to have found reasonable parameter values for both datasets. Finally, we depict the kriging results along with the resulting uncertainty quantification in



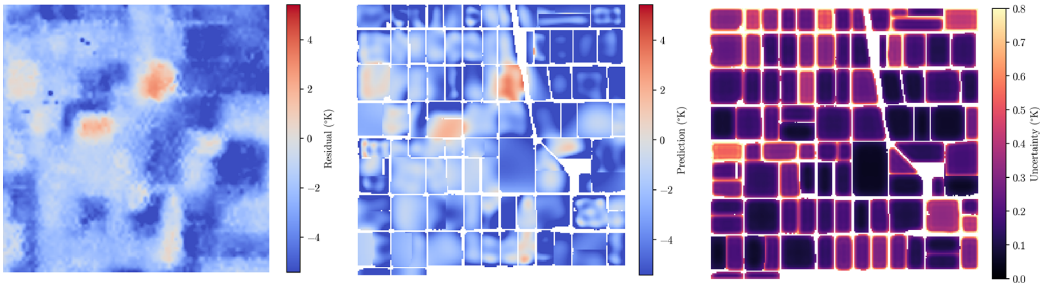
(a) ECOSTRESS original (b) ECOSTRESS residuals (c) Landsat original (d) Landsat residuals

Figure 4. *Side-by-side comparison of (a) the original ECOSTRESS image, (b) the ECOSTRESS residuals after subtracting the fitted mean function, (c) the original Landsat image, and (d) the Landsat residuals after subtracting the fitted mean function.*



(a) Landsat per-region variance (b) Per-region length scales (c) ECOSTRESS per-region variance

Figure 5. *Fitted parameters for each region: (a) Landsat per-region variance, (b) per-region length scales (shared between Landsat and ECOSTRESS), and (c) ECOSTRESS per-region variance.*



(a) Original ECOSTRESS residual (b) Kriging mean prediction (c) Kriging per-pixel 2σ

Figure 6. *Comparison of (a) the original ECOSTRESS residual image, (b) the kriging mean prediction, and (c) two per-pixel standard deviations for the kriging predictions.*

Figure 6. The kriging predictions effectively deblur many of the features present in the original ECOSTRESS image, and we can quantify the uncertainty in our predictions through the per-pixel variance. Note that the uncertainty for each region is higher near the edges; this is because the statistical assumption of independence between regions increases the uncertainty of the reconstruction near the edges.

We verify that our method works well by manually blurring and recovering a high-resolution Landsat image in Supplementary Appendix D.

4. Discussion

In this work, we proposed a novel statistical method for downscaling LST data from the ECOSTRESS mission using high-resolution data from the Landsat mission as a proxy for modeling agricultural field structure. By leveraging a block-diagonal Gaussian process model, we were able to capture the spatial structure of agricultural fields and perform scalable inference. Our results demonstrate that the proposed method effectively deblurs ECOSTRESS LST data, producing higher-resolution estimates along with uncertainty quantification. The fitted parameters varied significantly across regions, justifying our use of a block-diagonal GP to model the agricultural field structure. Furthermore, our toy example verified that the method can successfully reconstruct blurred images, recovering features present in the original data.

Despite the strengths of our proposed framework, several limitations remain. First, while we estimate parameters for the network of roads between agricultural fields, we do not reconstruct this region in the current work due to its size. Second, despite the strong performance of the segmentation method using the

Segment Anything Model (SAM), future studies might benefit from incorporating alternative segmentation algorithms or incorporating the uncertainty of the segmentation. Furthermore, the performance of our method is inherently limited by the quality of the ECOSTRESS data, which suffers from blur and artifacts resulting from physical damage to the sensor. This extends to our assumptions that the blur kernel and additive sensor noise are Gaussian; these assumptions may be inexact, given that Holmes et al. (2024) demonstrate that the ECOSTRESS point spread function is not perfectly Gaussian. Finally, to ensure identifiability, we rely on the simplifying assumption that the length scale parameters for ECOSTRESS and Landsat are identical, a condition that may not strictly hold in practice.

In future work, the Block-Diagonal Gaussian Process (BDGP) kernel could be augmented with a temporal length scale to capture time-varying dynamics. This would allow for gap-filling in time, which is of great practical value, given the large amount of missingness. Additionally, replacing the standard squared exponential kernel with a sharper Matérn kernel could more accurately model the roughness of the field. From a computational perspective, further optimizing the code would allow scalability to larger datasets. Another future direction is to model the mean of the high-resolution land surface temperature field, possibly using an auxiliary higher-resolution data source. Finally, the roads could be partitioned into smaller, more tractable regions so that they may be reconstructed as well.

The broader implication of this work is that by modeling a field structure using a block-diagonal GP, we obtain a pipeline for building a scalable statistical model for data fusion of sensors with complementary strengths. This approach accounts for both the change of support and field structure present in the data, allowing us to perform scalable statistical modeling on large datasets. For instance, this method could be applied to urban heat island modeling, European agricultural landscapes (with more irregular field shapes), or other remote sensing contexts where high-resolution proxies are available.

Supplementary material. The supplementary material for this article can be found at http://doi.org/10.1017/eds.2026.10052.

Acknowledgments. This research was partially carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

Author contribution. Conceptualization: S.D., M.J., and M.K.; Methodology: S.D. and M.K.; Data visualization: S.D.; Software: S.D.; Resources: M.J., G.H., M.P.-C., and M.K.; Supervision: M.J., M.P.-C., G.H., and M.K.; Writing—original draft: S.D.; Writing—review and editing: M.J. and M.K. All authors approved the final submitted draft.

Competing interests. The authors declare none.

Data availability statement. Code can be found at https://github.com/sanjitdp/downscaling-lst. Replication data can be obtained from https://appeears.earthdatacloud.nasa.gov/ using the .json request files contained in the code repository.

Ethics statement. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Funding statement. SD was sponsored by the CMU Department of Statistics and Data Science to attend Climate Informatics 2026.

Provenance statement. This article is part of the Climate Informatics 2026 proceedings and was accepted in Environmental Data Science based on the Climate Informatics peer review process.

References

Bindhu V, **Narasimhan B and Sudheer K** (2013) Development and verification of a non-linear disaggregation method (NL-DisTrad) to downscale MODIS land surface temperature to the spatial scale of Landsat thermal data to estimate evapotranspiration. *Remote Sensing of Environment 135*, 118–129. https://doi.org/10.1016/j.rse.2013.03.023.

Cressie N (1991) *Statistics for Spatial Data*. New York: John Wiley & Sons.

Dandapanthula S (2026) Downscaling-lst. https://doi.org/10.5281/zenodo.19348084; Available at https://github.com (accessed 31 March 2026).

Fisher JB, **Lee B**, **Purdy AJ**, **Halverson GH**, **Dohlen MB**, **Cawse-Nicholson K**, **Wang A**, **Anderson RG**, **Aragon B**, **Arain MA**, **Baldocchi DD**, **Baker JM**, **Barral H**, **Bernacchi CJ**, **Bernhofer C**, **Biraud SC**, **Bohrer G**, **Brunsell N**, **Cappelaere B**, **Castro-Contreras S**, **Chun J**, **Conrad BJ**, **Cremonese E**, **Demarty J**, **Desai AR**, **De Ligne A**, **Foltýnová L**, **Goulden ML**, **Griffis TJ**, **Grünwald T**, **Johnson MS**, **Kang M**, **Kelbe D**, **Kowalska N**, **Lim JH**, **Maïnassara I**, **McCabe MF**, **Missik JE**, **Mohanty BP**, **Moore CE**, **Morillas L**, **Morrison R**, **Munger JW**, **Posse G**, **Richardson AD**, **Russell ES**, **Ryu Y**, **Sanchez-Azofeifa A**,

Schmidt M, **Schwartz E**, **Sharp I**, **Šigut L**, **Tang Y**, **Hulley G**, **Anderson M**, **Hain C**, **French A**, **Wood E and Hook S** (2020) ECOSTRESS: NASA's next generation mission to measure evapotranspiration from the International Space Station. *Water Resources Research 56*(4). https://doi.org/10.1029/2019WR026058.

Fisher JB, **Melton F**, **Middleton E**, **Hain C**, **Anderson M**, **Allen R**, **McCabe MF**, **Hook S**, **Baldocchi D**, **Townsend PA**, **Kilic A**, **Tu K**, **Miralles DD**, **Perret J**, **Lagouarde JP**, **Waliser D**, **Purdy AJ**, **French A**, **Schimel D**, **Famiglietti JS**, **Stephens G and Wood EF** (2017) The future of evapotranspiration: Global requirements for ecosystem functioning, carbon and climate feedbacks, agricultural management, and water resources. *Water Resources Research 53*(4), 2618–2626. https://doi.org/10.1002/2016WR020175.

Garrigues S, **Allard D and Baret F** (2007) Using first- and second-order variograms for characterizing landscape spatial structures from remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing 45*(6), 1823–1834. https://doi.org/10.1109/TGRS.2007.894572.

Gelfand AE, **Diggle P**, **Guttorp P and Fuentes M** (2010) *Handbook of Spatial Statistics*. Boca Raton: CRC Press. https://doi.org/10.1201/9781420072884.

Holmes TR, **Poulter B**, **McCorkel J**, **Jennings D**, **Wu D**, **Efremova B**, **Shiklomanov A**, **Johnson W**, **Jhabvala M and Hook S** (2024) On-orbit spatial performance characterization for thermal infrared imagers of Landsat 7, 8, and 9, ECOSTRESS and CTI. *Journal of Geophysical Research: Biogeosciences 129*(2). https://doi.org/10.1029/2023JG007506.

Hutengs C and Vohland M (2016) Downscaling land surface temperatures at regional scales with random forest regression. *Remote Sensing of Environment 178*, 127–141. https://doi.org/10.1016/j.rse.2016.03.006.

Khaleghi B, **Khamis A**, **Karray FO and Razavi SN** (2013) Multisensor data fusion: A review of the state-of-the-art. *Information Fusion 14*(1), 28–44. https://doi.org/10.1016/j.inffus.2011.08.001.

Kirillov A, **Mintun E**, **Ravi N**, **Mao H**, **Rolland C**, **Gustafson L**, **Xiao T**, **Whitehead S**, **Berg AC**, **Lo WY**, **Dollár P and Girshick R** (2023) Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026. Los Alamitos, CA: IEEE.

Kumar S and Chakraborty S (2026) Scalable Random Wavelet Features: Efficient Non-Stationary Kernel Approximation with Convergence Guarantees. Preprint, arXiv:2602.00987.

Li W, **Ni L**, **Li ZL and Wu H** (2018) Downscaling land surface temperature by using random forest regression algorithm. In *IGARSS 2018–2018 IEEE International Geoscience and Remote Sensing Symposium*, Valencia, Spain. Piscataway, NJ: IEEE, pp. 2527–2530.

Li S, **Wan H**, **Yu Q and Wang X** (2025) Downscaling of ERA5 reanalysis land surface temperature based on attention mechanism and Google Earth Engine. *Scientific Reports 15*(1), 675. https://doi.org/10.1038/s41598-024-83944-w.

Njuki SM, **Mannaerts CM and Su Z** (2020) An improved approach for downscaling coarse-resolution thermal data by minimizing the spatial averaging biases in random forest. *Remote Sensing 12*(21), 3507. https://doi.org/10.3390/rs12213507.

Paciorek CJ and Schervish MJ (2006) Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics 17*(5), 483–506. https://doi.org/10.1002/env.785.

Roy D, **Wulder M**, **Loveland T**, **C.E W**, **Allen R**, **Anderson M**, **Helder D**, **Irons J**, **Johnson D**, **Kennedy R**, **Scambos T**, **Schaaf C**, **Schott J**, **Sheng Y**, **Vermote E**, **Belward A**, **Bindschadler R**, **Cohen W**, **Gao F**, **Hipple J**, **Hostert P**, **Huntington J**, **Justice C**, **Kilic A**, **Kovalskyy V**, **Lee Z**, **Lymburner L**, **Masek J**, **McCorkel J**, **Shuai Y**, **Trezza R**, **Vogelmann J**, **Wynne R and Zhu Z** (2014) Landsat-8: Science and product vision for terrestrial global change research. *Remote Sensing of Environment 145*, 154–172. https://doi.org/10.1016/j.rse.2014.02.001.

Samadzadegan F, **Toosi A and Dadrass Javan F** (2025) A critical review on multi-sensor and multi-platform remote sensing data fusion approaches: Current status and prospects. *International Journal of Remote Sensing 46*(3), 1327–1402. https://doi.org/10.1080/01431161.2024.2429784.

Wu J, **Xia L**, **Chan TO**, **Awange J and Zhong B** (2022) Downscaling land surface temperature: A framework based on geographically and temporally neural network weighted autoregressive model with spatio-temporal fused scaling factors. *ISPRS Journal of Photogrammetry and Remote Sensing 187*, 259–272. https://doi.org/10.1016/j.isprsjprs.2022.03.009.

Cite this article: Dandapanthula S, Johnson M, Pascolini-Campbell M, Hulley G and Kuusela M (2026). Downscaling land surface temperature data using edge detection and block-diagonal Gaussian process regression. *Environmental Data Science*, 5: e32. doi:10.1017/eds.2026.10052