Calibrated spatial uncertainty for Earth observation foundation models via Matérn-motivated latent stochastic regularization

Daniel P. Johnson[1]

[1] Department of Geography, School of Liberal Arts, Indiana University Indianapolis, Indianapolis, IN, USA
Email: dpjohnso@iu.edu

Preprint statement. This manuscript is a non-peer-reviewed preprint submitted to EarthArXiv. The manuscript has been submitted to *Remote Sensing of Environment* for peer review. Subsequent versions of this manuscript may have differing content. If accepted, the final version of this manuscript will be available via a journal-assigned DOI linked from this preprint.

Calibrated spatial uncertainty for Earth observation foundation models via Matérn-motivated latent stochastic regularization

Daniel P. Johnson[a,*]

[a]*Indiana University Indianapolis, Indianapolis, IN, USA*

Abstract

Earth observation foundation models produce dense spatial embeddings that support transfer learning across sensors and regions, yet these representations carry no explicit spatial statistical model for covariance, smoothness, or uncertainty. Existing deep learning uncertainty methods produce per-pixel variance estimates that ignore spatial dependence, and their calibration is rarely assessed beyond in-distribution data. This study introduces a Matérn-motivated latent stochastic regularization layer for frozen foundation model embeddings. The layer evolves embeddings through a learned Itô stochastic differential equation whose linear-case Fokker–Planck stationary distribution recovers the Matérn precision structure of the classical Whittle SPDE; in the implemented nonlinear parameterization, this correspondence serves as a design principle rather than as an explicit spatial differential operator. The layer adds fewer than 800,000 parameters with no foundation model retraining. Applied to OlmoEarth embeddings over a 10 km × 10 km training patch in Indianapolis, USA, the layer achieves 90 % prediction interval coverage within 0.8 percentage points of nominal for both NDVI and land surface temperature (LST). Under four-fold within-site spatial block holdout, coverage degrades by 3.9 percentage points for NDVI and 7.7 for LST, compared with 6.1–12.8 for matched-capacity baselines. Under deployment to the full Marion County study area, the layer retains PICP@90 = 82.7 % for LST while matched MLP and CNN baselines drop to 73.2–73.3 % despite comparable point predictions and similar mean predicted variances.

*Corresponding author.

Email address: dpjohnso@iu.edu (Daniel P. Johnson)

A tract-level analysis at the 38 °C summer LST threshold shows this calibration advantage flags approximately 33,000 additional residents as warranting heat-exposure investigation under a precautionary decision rule.

Keywords: Spatial uncertainty quantification, Prediction interval calibration, Gaussian random fields, Matérn covariance, Foundation models, Spatial domain shift

---

1. Introduction

Foundation models for Earth observation have expanded the available toolkit for geospatial analysis. Trained on continental-scale satellite imagery, models such as OlmoEarth (Beukema et al., 2025), PrithviEO (Jakubik et al., 2023), and SatMAE (Cong et al., 2022) produce dense embedding vectors that capture spectral–spatial patterns and support transfer learning across sensors, regions, and tasks. These embeddings are increasingly used as feature extractors for land cover classification, change detection, crop yield estimation, and environmental monitoring. Yet the embedding spaces produced by these models are, from a spatial statistical standpoint, unstructured. Each embedding vector is computed independently, or at best within a local patch, without any mechanism governing how neighbouring embeddings should covary across space. Spatial coherence, where it appears, is an emergent byproduct of training data rather than a property enforced by the model. As a result, foundation model embeddings lack three capabilities fundamental to spatial inference: an interpretable covariance function characterizing the scale and smoothness of spatial dependence; principled uncertainty quantification that propagates coherently across locations; and a formal connection to the theory of spatial stochastic processes that would enable rigorous interpolation and risk assessment.

Uncertainty quantification (UQ) in remote sensing has a long methodological tradition that predates the current wave of deep learning models. Foody's (2002) critique of single-number accuracy reporting in land cover classification established that remotely sensed products require richer per-pixel uncertainty characterization than confusion-matrix summaries can provide. Atkinson and Tate (2000) formalized the connection between variogram analysis, geostatistical uncertainty propagation, and remote sensing; this framework was extended by Atkinson (2013) to explicitly require uncertainty estimates as part of any defensible downscaling operation. Gaussian

process (GP) regression has since become the dominant Bayesian approach to producing per-pixel predictive distributions on EO products, with Verrelst et al. (2013) demonstrating GP-based LAI and chlorophyll retrievals with explicit uncertainty maps, and Camps-Valls et al. (2016) providing a comprehensive survey of GP applications in EO. Recent work has extended these ideas to deep learning: Ma et al. (2021) applied Bayesian neural networks to corn yield estimation with explicit prediction-interval evaluation; Lang et al. (2023) produced a global 10-m canopy height model with ensemble-based uncertainty and explicit reliability diagrams; and Martínez-Ferrer et al. (2022) decomposed aleatoric and epistemic uncertainty for high-resolution biophysical retrieval. These approaches share a common limitation when applied to foundation model embeddings: they either require retraining the entire pipeline end-to-end or treat spatial locations as independent for uncertainty estimation purposes.

The independence assumption matters, because uncertainty in spatial prediction is never truly local. Meyer and Pebesma (2022) demonstrated that random cross-validation systematically underestimates prediction error for spatially structured data, and introduced the concept of "area of applicability" to characterize where an ML model's predictions are supported by the training distribution. Kuronen et al. (2025) applied conformal prediction to produce calibrated 90% intervals for forest attributes from Sentinel-2 imagery, explicitly evaluating coverage under spatial block holdout. Work in the domain adaptation tradition (Tuia et al., 2016) has long emphasized that RS models trained on one acquisition regime degrade in characteristic ways when deployed elsewhere. Taken together, this literature points to a specific gap: methods that produce spatially coherent, calibrated prediction intervals on foundation model outputs without end-to-end retraining, and that maintain calibration when those outputs are used across geographic extents that differ from the training data.

Classical spatial statistics provides exactly the theoretical machinery needed. Gaussian random fields with Matérn covariance (Matérn, 1960; Stein, 1999) offer interpretable parameterizations of spatial range and smoothness, while the SPDE formulation introduced by Lindgren et al. (2011) yields sparse precision operators that scale to large spatial domains. These tools operate on observed spatial fields, however, not on learned representations: applying SPDE–Gaussian Markov random field (GMRF) methods (Rue and Held, 2005) directly to 768-dimensional embedding vectors produced by a vision transformer is not straightforward.

The motivation for the present work is that this incompatibility is an architectural problem, not a fundamental one. Foundation model embeddings carry substantial emergent spatial structure as a byproduct of training on spatially autocorrelated imagery: nearby pixels produce similar embedding vectors because the underlying imagery is autocorrelated. What the embeddings lack is not spatial information but spatial *statistical machinery* — parameterized covariance, calibrated uncertainty, principled propagation of error across scales. The classical SPDE–GMRF framework provides exactly that machinery, refined through decades of theoretical development (Whittle, 1954; Lindgren et al., 2011; Rue et al., 2009) and operational use in geostatistics, environmental science, and disease mapping. Bringing the two together therefore amounts not to hybridizing incompatible paradigms but to importing a mature spatial-statistical machinery into a setting that already exhibits spatial structure but lacks the apparatus to exploit it rigorously. The Fokker–Planck connection between SPDEs and Matérn fields suggests a particular route for this import: rather than enforcing covariance structure on the embeddings directly, one can evolve them through a stochastic dynamical system whose stationary distribution recovers the desired Matérn precision.

The latent SDE layer developed here is a lightweight stochastic module applied post hoc to frozen foundation model embeddings, motivated by the Matérn stationary structure of classical spatial SPDEs but implemented as a pointwise stochastic transformation in embedding space rather than as an explicit spatial differential operator. The layer evolves embeddings through a learned Itô stochastic differential equation with coupled drift and diffusion operators (Øksendal, 2003); the connection to classical Matérn theory is established through the Fokker–Planck equation, whose stationary solution for linear dynamics with constant diffusion recovers the Whittle SPDE (Whittle, 1954) precision operator. In the implemented nonlinear parameterization, this correspondence serves as a design principle that motivates the architecture, not as a literal spatial Matérn field imposed on the embeddings. The present paper focuses on empirical evaluation of this construction on real satellite imagery, with an emphasis on calibration behavior under spatial generalization beyond the training extent; the latent SDE is well-posed under standard Lipschitz and growth conditions on the drift and diffusion networks (Da Prato and Zabczyk, 1992; Øksendal, 2003).

The layer is evaluated on OlmoEarth embeddings over a 10 km × 10 km urban training patch within Marion County (Indianapolis), Indiana, USA,

with separate models trained to predict NDVI and LST at 10 m resolution. Four-fold spatial block holdout within the training patch evaluates calibration robustness under within-patch spatial domain shift, and inference on the full Marion County study area evaluates calibration under deployment-scale shift.

The central finding is that the SDE layer achieves prediction performance competitive with the matched MLP and CNN baselines (and a learnable Gaussian smoother) while providing prediction interval coverage that degrades substantially less under spatial holdout: 3.9 percentage points below nominal for NDVI (versus 6.1–7.1 for baselines) and 7.7 percentage points for LST (versus 11.0–12.8).

2. Study Area and Data

2.1. Study Area

The study area is Marion County, Indiana, USA (approximately 1043 km^2; centered on the city of Indianapolis), encompassing downtown urban core, industrial corridors, mixed residential neighbourhoods, suburban single-family zones, and agricultural-fringe land cover regimes characteristic of a Midwest U.S. metropolitan domain. Within this study area, a 10.24 km × 10.24 km training patch (approximate centroid 39.77°N, 86.13°W; 1024 × 1024 pixels at 10 m resolution in EPSG:32616, UTM zone 16N) was selected for model training. The training patch was chosen for 100% valid OlmoEarth embedding coverage and to include diverse land cover within a single contiguous tile: dense urban core, mixed residential, industrial zones, the airport, interstate transportation corridors, riparian vegetation along the White River, and agricultural margins. The training patch occupies less than 2 % of Marion County by area, providing a substantial spatial generalization challenge when models are deployed across the full study area (Section 4.5).

2.2. Data

2.2.1. OlmoEarth Embeddings

Foundation model embeddings were generated by OlmoEarth v1.0 (Beukema et al., 2025), a vision transformer trained on Sentinel-2 imagery. The model produces 768-dimensional embedding vectors quantized to Int8 (range $[-127, 127]$, with -128 reserved for nodata) at 10 m spatial resolution. Embeddings for the study area were extracted from annual composites for 2020, downloaded and compiled into a single virtual raster covering the Indianapolis metropolitan region. Prior to training, embeddings were normalized by dividing by

127, yielding floating-point values in $[-1, 1]$. No further preprocessing (principal component analysis, whitening, or band selection) was applied; all 768 dimensions were retained.

2.2.2. NDVI Target

Normalised difference vegetation index (NDVI) was computed from Sentinel-2 Level-2A surface reflectance imagery accessed via Google Earth Engine (collection `COPERNICUS/S2_SR_HARMONIZED`). Images were filtered to the June–August 2020 window and cloud-masked using the Scene Classification Layer (SCL band), retaining pixels classified as vegetation, bare soil, or water. A per-pixel median composite was computed across all valid observations in the temporal window, with NDVI calculated as $(\mathrm{B8} - \mathrm{B4})/(\mathrm{B8} + \mathrm{B4})$ where B8 is the near-infrared band and B4 is the red band. The resulting 10 m raster was exported in EPSG:32616 and aligned to the embedding grid via nearest-neighbour resampling; both datasets share the native 10 m Sentinel-2 grid. NDVI values were used directly as prediction targets without further transformation.

2.2.3. LST Target

Land surface temperature (LST) was derived from Landsat 8 Collection 2 Level 2 Science Products accessed via Google Earth Engine (collection `LANDSAT/LC08/C02/T1_L2`, band `ST_B10`). Images were filtered to June–August 2020 and cloud-masked using the `QA_PIXEL` bitmask, retaining clear-sky pixels only. Surface temperature values were converted from scaled digital numbers to degrees Celsius following the USGS conversion: $\mathrm{LST} = \mathrm{ST\_B10} \times 0.00341802 + 149.0 - 273.15$. A per-pixel median composite was computed across valid observations. The Landsat 8 Collection 2 ST product is distributed on the 30 m Landsat reflective-band grid, but the thermal infrared sensor's native physical thermal support is approximately 100 m (Zhan et al., 2013). The product was resampled from the 30 m distribution grid to 10 m via bilinear interpolation using `rasterio.warp.reproject` to match the embedding grid. This resampling does not recover thermal information below the underlying 100 m native physical support; this is a known limitation when interpreting fine-grained LST predictions, although that native physical support is itself a substantial improvement over MODIS-derived thermal products typically used at these scales (Zhan et al., 2013). The relationship between the 10 m SDE predictions and the underlying 100 m native physical thermal support is analogous to vegetation-index-based ther-

mal sharpening (Kustas et al., 2003; Agam et al., 2007): classical sharpening exploits the empirical relationship between coarse-resolution thermal observations and a co-registered fine-resolution covariate (typically NDVI from Landsat or MODIS) to disaggregate thermal pixels onto the covariate's finer grid. The SDE layer performs an analogous disaggregation, but with the 768-channel OlmoEarth embedding at 10 m taking the role of the single fine-resolution covariate. The classical and embedding-based formulations differ in the dimensionality and learnedness of the covariate space (one vegetation index versus a 768-dimensional learned representation) but share the underlying geostatistical principle: a high-resolution covariate field contains spatial information that, properly used, permits inference of a target variable at the covariate's resolution rather than the target's native resolution.

2.2.4. Nodata Handling

A small number of pixels (188 of 1,048,576 in the NDVI target) contained NaN values from incomplete cloud-mask coverage. These were set to zero and excluded from loss computation during training via a validity mask. The LST target contained no missing values within the training patch.

3. Methods

The latent SDE layer is a post-hoc module applied to OlmoEarth embeddings without any retraining of the foundation model.

3.1. Latent SDE Layer

The implemented layer is a finite-dimensional latent SDE acting pointwise on each 768-dimensional embedding vector. It is therefore not a classical spatial SPDE and does not impose a mesh-based Matérn precision matrix over geographic coordinates. The connection to Whittle–Matérn theory is used as an architectural design principle: in the linear spatial-field analogue developed in Section 3.2, the Fokker–Planck stationary distribution links drift–diffusion dynamics to a Matérn precision operator, and that correspondence motivates the coupled drift and diffusion architecture used in the embedding-space implementation. In the nonlinear embedding-space implementation used here, this correspondence does not provide a closed-form Matérn covariance or a posterior over a spatial random field. Three clarifications are essential to position the construction correctly. First, the implemented layer

is pointwise in embedding space: it evolves each pixel's 768-dimensional embedding vector through an SDE in feature-space pseudo-time, with no explicit spatial differential operator, finite-element mesh, or graph Laplacian over geographic coordinates. Spatial coherence in the output arises because neighboring pixels begin with similar foundation-model embeddings, and the learned drift and diffusion networks process those similar inputs into similar outputs. Second, the layer is not a Bayesian SPDE–Gaussian Markov random field method in the sense of Lindgren et al. (2011): no posterior is computed over a latent random field, and no INLA or Markov chain Monte Carlo inference is performed. The Fokker–Planck identification serves as a design principle for a neural architectural component trained by maximum likelihood, with the Matérn-shaped implicit prior acting as a regularizer rather than as an integrated posterior. Third, the layer produces predictive moments $(\hat{\mu}, \hat{\sigma})$ as deterministic point estimates from learned decoders applied to the SDE-transformed embeddings, not as posterior summaries. The naming convention "latent SDE layer" reflects the design heritage and stationary-distribution motivation; the empirical contribution is prediction interval calibration, demonstrated below.

The latent SDE layer operates on OlmoEarth embeddings $\mathbf{z}(s) \in \mathbb{R}^{768}$ indexed by spatial location s, evolving each embedding through a learned Itô stochastic differential equation (Øksendal, 2003) in pseudo-time τ:

$$\mathrm{d}\mathbf{z} = F_\varphi(\mathbf{z})\,\mathrm{d}\tau + \Sigma_\psi(\mathbf{z})\,\mathrm{d}\mathbf{W}, \tag{1}$$

where $F_\varphi : \mathbb{R}^{768} \to \mathbb{R}^{768}$ is the drift operator, $\Sigma_\psi : \mathbb{R}^{768} \to \mathbb{R}^{768}_+$ is the diagonal diffusion operator, and $\mathbf{W}$ is a standard Wiener process in $\mathbb{R}^{768}$. The drift and diffusion terms are used in the sense of Itô stochastic calculus: the drift governs deterministic evolution of the embeddings while the diffusion injects structured stochastic perturbations whose amplitude is learned from data. These are not physical advection or material diffusion processes; they are components of a stochastic process acting in embedding space.

Both operators are parameterized as two-layer multilayer perceptrons ($768 \to 256 \to 768$) with GELU activation. The drift is initialized near zero (output weights scaled by 0.01, biases zeroed) so that the layer initially acts as an approximate identity. The diffusion uses a softplus output nonlinearity with positivity floor $\varepsilon = 10^{-3}$ to ensure uniform ellipticity: $\Sigma_\psi(\mathbf{z}) = \mathrm{softplus}(g_\psi(\mathbf{z})) + \varepsilon$. The total parameter count for the full model is 790,018 (drift MLP: 394,240; diffusion MLP: 394,240; mean decoder: 769; variance decoder: 769).

A key design choice is that all operators act *pointwise* on individual embedding vectors: the drift, diffusion, and decoder networks process each pixel independently through shared parameters, with no explicit spatial interaction between neighbouring pixels in the forward pass. Spatial structure in the outputs arises because neighbouring pixels have similar embeddings, a consequence of the foundation model's training on spatially autocorrelated imagery, and the shared nonlinear SDE transformation maps similar inputs to similar outputs. The distinction from raw embeddings is that the SDE transformation is designed using the Fokker–Planck/Matérn correspondence described below as an architectural motivation, whereas the raw embeddings carry no comparable stochastic regularization principle.

3.2. Connection to Matérn Covariance via the Fokker–Planck Equation

The coupled drift–diffusion architecture is motivated by the stationary structure of the Itô SDE in Equation (1). For the SDE, the evolution of the probability density $p(\mathbf{z}, \tau)$ of the latent state is governed by the Fokker–Planck equation

$$\frac{\partial p}{\partial \tau} = -\sum_i \frac{\partial}{\partial z_i}(F_{\varphi,i}\, p) + \frac{1}{2}\sum_{i,j} \frac{\partial^2}{\partial z_i\, \partial z_j}(D_{\psi,ij}\, p), \tag{2}$$

where $D_\psi = \Sigma_\psi \Sigma_\psi^\top$ is the diffusion tensor (diagonal in the parameterization used here).

For the linear case, $F_\varphi(\mathbf{z}) = A\mathbf{z}$ with A a stable matrix and $\Sigma_\psi = B$ constant, the stationary solution of Equation (2) is Gaussian, $p_\infty(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \Sigma_\infty)$, with stationary covariance satisfying the Lyapunov equation (Kailath, 1980; Da Prato and Zabczyk, 1992)

$$A\Sigma_\infty + \Sigma_\infty A^\top + BB^\top = 0, \tag{3}$$

and corresponding precision $Q = \Sigma_\infty^{-1} = -2\sigma^{-2}A$ when $BB^\top = \sigma^2 I$.

To connect this finite-dimensional SDE motivation to the classical spatial case, consider the analogous linear stochastic evolution of a spatial field $u(s, \tau)$ in which the drift operator acts over geographic coordinates rather than over the 768 embedding coordinates of $\mathbf{z}$. The classical Whittle SPDE (Whittle, 1954) defines a Matérn Gaussian random field as the stationary solution to $(\kappa^2 - \nabla^2)^{\alpha/2}\, u(s) = \sigma\, \mathcal{W}(s)$, where ∇^2 is the spatial Laplacian and $\mathcal{W}$ is spatial white noise. Choosing the corresponding spatial drift

$A = -(\sigma^2/2)(\kappa^2 - \nabla^2)^\alpha$ in the linear parabolic SPDE yields a stationary precision $Q \propto (\kappa^2 - \nabla^2)^\alpha$, recovering the Matérn precision structure (Lindgren et al., 2011). The operator identification displayed in this paragraph applies to the spatial-field analogue, not directly to the implemented 768-dimensional pointwise embedding SDE; the latter is finite dimensional, with all operators acting over embedding coordinates. The role of the spatial-field analogue is to motivate the architectural choice of coupled drift and diffusion in the embedding SDE, by analogy with the classical case in which such a coupling is known to produce Matérn stationary structure.

The structural correspondence between a parabolic SPDE with Matérn-form drift and a stationary spatial field with Matérn precision has been developed in the spatial statistics literature. Carrizo Vergara et al. (2022) established a general framework for stationary SPDE-based random fields parameterized by spectral symbols, under which the Whittle Matérn field is recovered for appropriate symbol choices. Clarotto et al. (2024) subsequently placed the full-power operator $(\kappa^2 - \nabla \cdot H \nabla)^\alpha$ directly in the drift of a parabolic SPDE and proved that its stationary spatial marginal is Matérn, via spectral-symbol arguments. A complementary finite-dimensional construction appears in the state-space Gaussian-process literature, where Särkkä and Solin (2019) derived temporal SDEs whose stationary marginals correspond to Matérn covariances. The Fokker–Planck calculation given above is an operator-level specialization of this line of work: the gradient-flow parameterization $BB^\top = \sigma^2 I$ with commuting A simplifies the Lyapunov equation to a direct identification between the drift and the Matérn precision. This paper applies that identification as an architectural design principle for learned foundation-model embeddings, a role it has not previously played.

In the nonlinear regime, where F_φ and Σ_ψ are parameterized by neural networks, this exact correspondence no longer holds. Under standard Lipschitz and growth conditions on F_φ and Σ_ψ, the latent SDE is well-posed (Da Prato and Zabczyk, 1992; Øksendal, 2003); existence of a stationary invariant measure requires additional dissipativity or recurrence conditions which are not enforced here, so the Matérn correspondence is used as a design motivation for the architecture rather than as a closed-form guarantee in the learned nonlinear parameterization. The empirical contribution of the layer — calibration robustness under spatial deployment shift — is demonstrated in Section 4 below.

3.3. Numerical Integration and Training

Equation (1) is integrated using the Euler–Maruyama scheme (Kloeden and Platen, 1992) over $K = 10$ steps with step size $\Delta\tau = 0.01$:

$$\mathbf{z}_{k+1} = \mathbf{z}_k + \Delta\tau\, F_\varphi(\mathbf{z}_k) + \sqrt{\Delta\tau}\, \Sigma_\psi(\mathbf{z}_k) \odot \boldsymbol{\epsilon}_k, \qquad \boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{4}$$

where $\odot$ denotes elementwise multiplication. Because each embedding vector is processed independently, no spatial mass matrix or mesh is required. The integration is fully differentiable via the reparameterization trick (Kingma and Welling, 2014), enabling end-to-end gradient-based training.

The model is trained with a Gaussian negative log-likelihood (NLL) loss (Nix and Weigend, 1994) that jointly learns the predictive mean and a heteroscedastic variance:

$$\mathcal{L} = \frac{1}{|M|} \sum_{s \in M} \left[\frac{1}{2} \log \sigma^2(s) + \frac{(y(s) - \hat{\mu}(s))^2}{2\,\sigma^2(s)} \right], \tag{5}$$

where M is the set of valid pixels, $\hat{\mu}(s)$ is the predicted mean from a linear decoder applied to the SDE-transformed embeddings, and $\sigma^2(s) = \exp(\hat{v}(s))$ is the predicted variance from a parallel linear decoder operating on the same transformed embeddings. This objective is a strictly proper scoring rule for Gaussian predictive distributions (Gneiting and Raftery, 2007): it penalizes both overconfident predictions (small σ^2, large error) and overly conservative predictions (large σ^2) through the log-variance term, encouraging calibrated uncertainty estimates.

Training used the AdamW optimizer (Loshchilov and Hutter, 2019) with learning rate 10^{-4}, weight decay 10^{-5}, and cosine annealing over 50 epochs. Each epoch comprised 250 iterations, each drawing a random 64×64 pixel patch ($640\,\mathrm{m} \times 640\,\mathrm{m}$) from the training patch. Patch locations were sampled uniformly at random from all valid positions within the 1024×1024 training patch. Separate models were trained for NDVI and LST. For LST, target values were normalized to $[0, 1]$ by min–max scaling before training, with predictions converted back to degrees Celsius for evaluation.

Computational cost.. All experiments were conducted on a single NVIDIA A100 GPU on Indiana University's Big Red 200 supercomputer. Training (50 epochs $\times$ 250 iterations per target variable) and full-image inference both fit comfortably on one GPU; a single deterministic forward pass over the training patch ($\sim$1 million pixels) completed in approximately 10 seconds. Memory

requirements were dominated by the embedding tensor ($1024 \times 1024 \times 768$ float32 $\approx 3\,\mathrm{GB}$); the SDE layer's parameters (790,018) occupied less than 4 MB.

3.4. Baseline Architectures

Five baseline architectures provide comparison points that isolate different aspects of the SDE layer's contribution. All baselines were trained with the same AdamW optimizer, learning-rate schedule, and patch-sampling protocol as the SDE layer.

Raw baseline.. A single linear decoder ($768 \to 1$; 769 parameters) applied directly to the embeddings, trained with mean-squared-error loss. This baseline has no mechanism for uncertainty estimation and tests whether the SDE layer adds predictive value beyond what is already encoded in the embeddings.

Baseline+NLL.. A two-decoder architecture (mean + log-variance; 1,538 parameters) applied directly to the embeddings without SDE processing, trained with NLL loss. This isolates the contribution of the loss function from that of the SDE architecture: any NLL-trained model can learn to produce marginally calibrated intervals on its training distribution.

Drift-only and diffusion-only ablations.. Two ablation models isolate the SDE components. The drift-only model ($\Sigma_\psi = 0$; 395,778 parameters) retains the deterministic drift integration but removes all stochastic perturbation; the diffusion-only model ($F_\varphi = 0$; 395,778 parameters) retains the stochastic perturbation but removes the deterministic drift. Both include the NLL variance head and use $K = 10$ integration steps.

Matched-capacity MLP.. A matched-parameter MLP (790,018 parameters) applying two residual MLP blocks ($768 \to 256 \to 768$, GELU activation) in a single forward pass with no SDE integration, followed by the same mean and log-variance decoders. This baseline controls for the possibility that the SDE's gains arise from added network capacity rather than stochastic dynamics.

Shallow CNN.. A shallow convolutional network (315,938 parameters) projecting embeddings to 96 channels, applying two 3×3 convolutional layers with residual connections, and projecting back to 768 dimensions. This introduces explicit spatial coupling between neighbouring pixels, a capability the

pointwise SDE layer lacks in its forward pass, and tests whether conventional spatial architectures achieve similar calibration.

Gaussian spatial smoother.. A learnable separable Gaussian kernel (kernel size 7, σ learned during training; 395,779 parameters) applied to the embedding grid, followed by a small refinement MLP with residual connection. This tests whether simple spatial smoothing of embeddings achieves comparable performance to the SDE layer.

3.5. Evaluation Protocol

3.5.1. In-Domain Evaluation

For in-domain evaluation, models were trained on the full training patch and evaluated on the same pixels. While this evaluation does not measure generalization, it establishes the upper bound on each architecture's predictive capacity and marginal calibration. Because all models operate pointwise on embedding vectors (or on local convolutional neighbourhoods in the case of the shallow CNN), there is no information leakage between pixels during the forward pass.

3.5.2. Four-Fold Spatial Block Holdout

Spatial generalization was evaluated via four-fold spatial block holdout. The 1024×1024 training patch was partitioned into four 512×512 quadrants (NW, NE, SW, SE). For each fold, models were trained on three quadrants and evaluated on the held-out quadrant; results are reported as the mean across the four folds. This design follows recent recommendations for evaluating spatial ML models under distribution shift (Meyer and Pebesma, 2022; Kuronen et al., 2025), which have established that random cross-validation systematically underestimates out-of-sample prediction error for spatially structured data. Block holdout forces each model to generalize across contiguous unseen regions within the training patch.

3.5.3. Operational-Extent Evaluation

For the strongest test of deployment generalization, models trained on the full training patch were applied to the entire Marion County, Indiana extent (approximately $1043\,\mathrm{km}^2$, 10 428 052 valid pixels), of which the primary patch occupies less than 2 %. This addresses a distinct dimension of domain shift, namely the spatial scale and land-cover heterogeneity of operational deployment, and complements the within-site block holdout. Embedding

preprocessing for the Marion extent followed the same procedure as for the primary patch (Section 2.2); inference was tiled in non-overlapping 64×64 patches with predictions masked to the Marion County boundary using the 2020 TIGER/Line shapefile.

3.5.4. Evaluation Metrics

Point prediction accuracy.. Root mean squared error (RMSE) and coefficient of determination (R^2) computed over all valid pixels in whichever evaluation regime is reported (training patch, held-out quadrant, or full Marion County extent, as specified per result).

Spatial structure.. Moran's I (Moran, 1950) on prediction residuals, computed with rook contiguity (four nearest neighbours) on the 10 m grid. Moran's I ranges from -1 (perfect dispersion) through 0 (spatial randomness) to $+1$ (perfect autocorrelation); lower residual Moran's I indicates more spatial signal captured by the model.

Empirical variograms.. Variograms (Matheron, 1963) computed on predictions and targets using exact lag distances along horizontal and vertical grid axes (no diagonal pairs) at integer pixel lags from 1 to 50 (10 m to 500 m). The semivariance at lag h is $\hat{\gamma}(h) = \frac{1}{2|N(h)|} \sum_{(s_i,s_j)\in N(h)} (z(s_i) - z(s_j))^2$, where $N(h)$ is the set of pixel pairs separated by exactly h pixels along a grid axis. This diagnostic evaluates whether the model reproduces the second-order spatial structure of the target field, following the geostatistical tradition established in the remote sensing uncertainty literature (Atkinson and Tate, 2000).

Prediction interval coverage.. Uncertainty calibration is evaluated via prediction interval coverage probability (PICP). For a nominal coverage level $1 - \alpha$ (so that $\alpha = 0.1$ corresponds to a 90 % interval), the predicted interval is $\hat{\mu}(s) \pm z_{1-\alpha/2}\, \hat{\sigma}(s)$, where $z_{1-\alpha/2}$ is the standard normal quantile (with $z_{0.95} \approx 1.645$ for the 90 % interval) and $\hat{\sigma}(s) = \sqrt{\exp(\hat{v}(s))}$ is the predicted standard deviation from the variance head. Empirical coverage is the fraction of valid pixels whose observed value falls within the interval. Coverage is reported at both the 90 % and 50 % levels; a well-calibrated model achieves empirical coverage close to the nominal level across the full range of quantiles. This evaluation protocol follows the forecast-verification tradition established by Hamill (2001) for assessing ensemble-based probabilistic predictions.

Reliability diagnostics.. Beyond interval coverage at fixed levels, calibration is assessed across the full predictive distribution via reliability plots: binning pixels by predicted quantile and comparing empirical versus expected coverage (Figures 1, 2). Deviations from the diagonal indicate miscalibration at specific quantile levels. For spatially structured predictions, per-quadrant PICP is also reported under the block holdout design to reveal spatial heterogeneity in calibration.

4. Results

4.1. In-Domain Prediction and Calibration

The first evaluation tested whether the SDE layer produces accurate predictions and calibrated uncertainty on the training patch when trained on the full 1024×1024 grid. Separate models were trained to predict NDVI and LST from OlmoEarth embeddings using the training protocol described in Section 3.3.

For NDVI, the SDE layer improved prediction from $R^2 = 0.909$ (RMSE $=$ 0.069) with raw embeddings to $R^2 = 0.944$ (RMSE $= 0.054$), a 22% reduction in RMSE (Table 1). The largest error reductions occurred at transition zones between impervious surface and vegetation (Figure 3c).

For LST, the SDE layer improved prediction from $R^2 = 0.814$ (RMSE $= 1.79\,^{\circ}\mathrm{C}$) to $R^2 = 0.861$ (RMSE $= 1.55\,^{\circ}\mathrm{C}$), a 13% reduction in RMSE (Table 1). The largest gains occurred in mixed residential neighbourhoods where 10 m embedding structure captures fine-scale surface heterogeneity associated with variation in the resampled Landsat-derived LST field (Figure 4c).

Both models produced well-calibrated prediction intervals. For NDVI, the 90% and 50% prediction intervals achieved empirical coverage of 90.8% and 51.3% respectively, within 0.8 and 1.3 percentage points of nominal. For LST, coverage was 90.2% and 50.3%, within 0.2 and 0.3 percentage points of nominal. Monte Carlo sampling of the stochastic forward pass (50 draws) yielded near-identical predictions across draws, indicating that the learned variance field absorbs the stochastic structure introduced during training; the training-time role of the diffusion term is examined in Section 4.2.

Reliability diagnostics confirmed calibration across the full range of quantile levels, not merely at 90% and 50%. Empirical coverage closely tracks the nominal level across quantile levels from 10% to 95%, with maximum absolute deviation below 2 percentage points at any level. Per-pixel predicted uncertainty correlates with absolute prediction error ($r = 0.574$ for NDVI;

$r = 0.369$ for LST), confirming that the model assigns higher uncertainty where errors are larger (Figures 1c, 2c).

4.2. Ablation Analysis: Drift, Diffusion, and Loss Function

Three ablation experiments isolate the contributions of the SDE architecture's components (Table 2).

First, the drift-only model (diffusion disabled) retained most of the predictive improvement over the raw baseline ($R^2 = 0.938$, RMSE = 0.056), confirming that the learned drift operator is the primary mechanism for spatial regularization. However, its prediction interval coverage undershot nominal levels (PICP@90 = 89.2% versus 90.8% for the full model, PICP@50 = 48.9% versus 51.3%), indicating that the variance head alone cannot achieve full calibration without the stochastic regularization provided by the diffusion term during training.

Second, the diffusion-only model (drift disabled) degraded prediction below the raw baseline ($R^2 = 0.902$ versus 0.909; RMSE = 0.071 versus 0.069) despite having 395,778 parameters, over 500× the baseline's 769. This result rules out the interpretation that the SDE layer improves prediction simply by adding network capacity: stochastic perturbations without deterministic spatial structure add noise rather than signal.

Third, the Baseline+NLL control (1,538 parameters; same two-decoder architecture as the SDE but applied to raw embeddings, trained with NLL loss) achieved approximate nominal calibration (PICP@90 = 90.5%, PICP@50 = 49.7%), demonstrating that the NLL loss is sufficient for marginal calibration regardless of architecture. However, this model achieved calibration by inflating its prediction intervals: its mean predicted standard deviation ($\hat{\sigma} = 0.059$) was 28% larger than the full SDE layer's ($\hat{\sigma} = 0.046$), compensating for worse predictions (RMSE = 0.071 versus 0.054) with wider intervals. The interpretation of these ablation results in terms of calibration versus informativeness, and the role of the diffusion term as a training-time regularizer, is developed in Section 5.2.

4.3. Spatial Structure Diagnostics

A complementary assessment evaluated whether the SDE layer reproduces the second-order spatial structure of the target fields. Predictions from a linear decoder applied to raw embeddings exhibited Moran's $I = 0.468$, comparable to the target NDVI field itself ($I = 0.470$), and their empirical variogram rose monotonically toward a sill near 0.039 at lags of 300 m–500 m.

Table 1: Prediction performance and uncertainty calibration on the training patch. RMSE and R^2 are computed over all valid pixels; PICP denotes prediction interval coverage probability (empirical coverage at nominal levels of 90% and 50%). The raw baseline has no uncertainty mechanism.

Target	Model	RMSE	R^2	PICP@90	PICP@50
NDVI	Raw baseline	0.069	0.909	—	—
	SDE layer	0.054	0.944	90.8%	51.3%
LST	Raw baseline	1.79 °C	0.814	—	—
	SDE layer	1.55 °C	0.861	90.2%	50.3%

Table 2: Ablation analysis for NDVI prediction. "Drift only" disables diffusion ($\Sigma_\psi = 0$); "Diffusion only" disables drift ($F_\varphi = 0$). "Baseline+NLL" applies the two-decoder mean and log-variance architecture to raw embeddings with no SDE processing. Mean $\hat{\sigma}$ is the average predicted standard deviation; at equivalent calibration, lower values indicate tighter, more informative intervals.

Model	Params	RMSE	R^2	PICP@90	PICP@50	Mean $\hat{\sigma}$
Full SDE	790,018	0.054	0.944	90.8%	51.3%	0.046
Drift only	395,778	0.056	0.938	89.2%	48.9%	0.046
Diffusion only	395,778	0.071	0.902	90.5%	49.5%	0.059
Baseline+NLL	1,538	0.071	0.903	90.5%	49.7%	0.059
Raw baseline	769	0.069	0.909	—	—	—

After application of the SDE layer, predictions track the target variogram more closely than baseline predictions across all lag distances (Figure 5a), reproducing both the curvature near the origin and the effective range at which semivariance stabilizes. Minor divergence between the empirical SDE and target variograms at lag distances above approximately 300 m is consistent with the well-known finite-sample property of empirical variogram estimation: as lag grows, the number of valid pixel pairs available for estimation drops sharply, increasing the sampling variance of the semivariance estimator at large lags (Matheron, 1963; Stein, 1999). The LST variogram exhibited the same pattern.

The uncertainty estimates themselves exhibited spatial coherence: predicted standard deviations were largest in areas of high land cover heterogeneity and along boundaries between distinct surface types, and smallest in spectrally homogeneous regions such as dense forest or open water (Figures 1a, 2a). The interpretation of these spatial diagnostics in relation to classical geostatistical structure is developed in Section 5.3.

4.4. Calibration Robustness under Spatial Block Holdout

On the training quadrants, all models achieved approximate nominal calibration (Table 3), with PICP@90 ranging from 90.2% to 90.8%. The matched MLP and CNN baselines, which lack explicit spatial structural constraints, matched or exceeded the SDE layer on point prediction: the shallow CNN achieved $R^2 = 0.959$ for NDVI versus the SDE's 0.944, and the matched-capacity MLP achieved $R^2 = 0.955$. These in-domain results establish that spatially unstructured architectures can match or exceed the SDE on point prediction and marginal calibration when evaluated on data drawn from the training distribution.

The calibration picture changes substantially under block holdout (Table 4). The SDE layer's 90% prediction interval coverage on held-out NDVI blocks dropped from 90.7% (train) to 86.7% (test), a degradation of 3.9 percentage points. The CNN degraded from 90.4% to 83.4% (7.0 pp), the matched MLP from 90.3% to 83.2% (7.1 pp), and the Gaussian smoother from 90.3% to 84.2% (6.1 pp). For LST, where absolute degradation was larger for all models reflecting the harder prediction task, the SDE's coverage dropped 7.7 pp versus 11.0–12.8 pp for the CNN, MLP, and smoother. Note that the R^2 degradations are not directly comparable across targets: LST operates on a physical temperature scale (°C) with substantially higher

variance than the dimensionless NDVI scale. The SDE exhibits the smallest absolute R^2 gap between training and test quadrants for both targets (+0.032 for NDVI, +0.085 for LST), with baselines ranging from +0.037 to +0.038 for NDVI and +0.090 to +0.101 for LST. The mechanism underlying this calibration robustness is examined in Section 5.2.

4.5. Marion County Evaluation

Whereas Section 4.4 addresses the within-patch generalization question, this subsection reports calibration behaviour when models trained on the 10 km training patch are applied across the full Marion County study area (Section 3.5), a setting in which inference is required at distances of up to 20 km from the training extent and across land cover regimes underrepresented in the training data.

Table 5 reports prediction interval coverage on the training patch and on the full Marion evaluation, together with the patch-to-county delta. All four architectures achieve nominal coverage within 1 percentage point of the 90 % target on the training patch. Under deployment to the full county, the SDE layer's PICP@90 degrades by 6.4 percentage points for NDVI and 7.4 percentage points for LST. The matched-capacity MLP and shallow CNN baselines degrade by 8.0–8.4 percentage points for NDVI and 16.5–16.8 percentage points for LST. For LST in particular, the SDE retains PICP@90 = 82.7%, while MLP and CNN drop to 73.2–73.3 %. Point prediction R^2 shows the opposite ordering: MLP and CNN achieve slightly higher R^2 than the SDE on Marion (ΔR^2 between +0.008 and +0.016 for both NDVI and LST), so the calibration result is not attributable to the SDE making more accurate point predictions. The Gaussian smoother baseline degrades less on LST PICP@90 (10.3 points) but suffers a 22.5-point degradation on NDVI PICP@90, falling to 67.8 % empirical coverage at the 90 % nominal level; this inconsistent calibration behaviour across targets is consistent with the smoother's lack of learned drift structure.

Beyond PICP@90, the Winkler interval score (Gneiting and Raftery, 2007) jointly penalizes prediction interval width and miscoverage, providing a calibration metric that discourages trivial widening: a model that achieves nominal coverage by inflating its variance pays a width penalty, while a model that achieves coverage with tighter, well-placed intervals does not. On the Marion County extent, the SDE layer achieves the lowest interval score for LST (9.67 °C) compared to 10.55 for the matched CNN and 10.57 for the matched MLP, despite producing intervals approximately 29 %

Table 3: In-domain performance on the training patch. All models except the Raw baseline (trained with MSE loss) were trained with NLL loss on the full 1024×1024 area; the Raw baseline provides a point-prediction reference and does not produce calibrated intervals. PICP@90 denotes 90 % prediction interval coverage; Mean $\hat{\sigma}$ is the average predicted standard deviation (lower values indicate tighter intervals at equivalent calibration).

Model	Params	R^2	RMSE	PICP@90	Mean $\hat{\sigma}$
Shallow CNN	315,938	0.959	0.046	90.4%	0.039
Matched MLP	790,018	0.955	0.048	90.3%	0.041
Gaussian smoother	395,779	0.949	0.051	90.2%	0.044
Full SDE	790,018	0.944	0.054	90.8%	0.046
Baseline+NLL	1,538	0.903	0.071	90.5%	0.059
Raw baseline	769	0.909	0.069	—	—

Table 4: Spatial block holdout results: mean ± SD across four folds (NW, NE, SW, SE). Models trained on three 512×512 quadrants and evaluated on the held-out quadrant. ΔR^2 is the train–test gap (larger values indicate more overfitting to the training quadrants). ΔPICP@90 is the train–test coverage degradation in percentage points. The SDE layer exhibits the smallest ΔR^2 and smallest |ΔPICP| for both targets.

Model	Train R^2	Test R^2	ΔR^2	Test RMSE	Test PICP@90	ΔPICP@
NDVI (target units: dimensionless [0, 1])						
Full SDE	0.942 ± 0.010	0.909 ± 0.040	+0.032	0.065 ± 0.008	$86.7 \pm 4.2\%$	−3.9 p
Shallow CNN	0.962 ± 0.006	0.925 ± 0.038	+0.037	0.058 ± 0.010	$83.4 \pm 4.2\%$	−7.0 p
Matched MLP	0.957 ± 0.008	0.919 ± 0.039	+0.038	0.061 ± 0.009	$83.2 \pm 3.9\%$	−7.1 p
Gaussian smoother	0.952 ± 0.008	0.914 ± 0.040	+0.038	0.063 ± 0.009	$84.2 \pm 3.5\%$	−6.1 p
LST (target units: °C)						
Full SDE	0.850 ± 0.017	0.765 ± 0.039	+0.085	1.89 ± 0.30	$82.1 \pm 5.4\%$	−7.7 p
Shallow CNN	0.910 ± 0.003	0.820 ± 0.032	+0.090	1.66 ± 0.30	$76.8 \pm 7.9\%$	−12.8 p
Matched MLP	0.900 ± 0.008	0.799 ± 0.039	+0.101	1.75 ± 0.29	$76.5 \pm 7.2\%$	−12.8 p
Gaussian smoother	0.886 ± 0.008	0.794 ± 0.043	+0.092	1.77 ± 0.32	$78.4 \pm 6.1\%$	−11.0 p

wider on average (6.11 °C versus 4.75 °C). The matched baselines incur miscoverage penalties of 5.78–5.80 versus 3.56 for the SDE (Table 5, Marion IS column). For NDVI, the matched MLP achieves a marginally lower interval score (0.246 versus 0.259 for SDE), with all three learned-feature architectures incurring comparable miscoverage penalties (0.111–0.132). The Gaussian smoother is the worst performer on interval score for both targets (0.428 NDVI; 9.71 LST). The implications of these interval-score results for the question of whether SDE calibration is bought through trivial interval widening are developed in Section 5.2.

4.5.1. Comparison to Split Conformal Prediction

A natural comparison for any post-hoc calibration method is split conformal prediction (Vovk et al., 2005; Lei et al., 2018), which produces prediction intervals with marginal coverage guarantees on the calibration distribution by construction. To assess whether those guarantees survive deployment shift to Marion County, two split conformal variants were applied to each matched-capacity baseline (MLP, CNN, Gaussian smoother). For each baseline, in-domain predictions on the training patch were combined with the corresponding target to compute residuals $|y(s) - \hat{\mu}(s)|$, and the empirical $1 - \alpha = 0.90$ quantile of those residuals (with the standard finite-sample correction) was taken as the calibration radius $\hat{q}$, yielding the conformal interval $\hat{\mu}(s) \pm \hat{q}$. The two variants differ in how the calibration set is drawn: "conformal_standard" uses a random 20 % of training-patch pixels, while "conformal_spatial" uses one 512×512 spatial block of the training patch. To avoid sensitivity to the choice of spatial block, the spatial-block variant was run separately with each of the four quadrants (NW, NE, SW, SE) as the calibration set; results are reported as the mean and range across the four quadrants. The spatial-block construction is a spatially blocked calibration split motivated by the spatial cross-validation principles of Meyer and Pebesma (2022) and is in the same family as the conformal prediction approaches applied to remote sensing by Kuronen et al. (2025).

Table 6 reports Marion-extent performance for both variants. For LST, all four spatial-block calibrations underperform the SDE on PICP@90, with mean PICP@90 across the four quadrants of 0.66–0.72 (versus 0.83 for SDE) and a range of 0.59 to 0.78 across all twelve spatial-block configurations. Standard conformal performs comparably (PICP@90 = 0.69–0.73). Mean conformal interval scores on LST across the four quadrants (10.78–11.50) are worse than those of the uncalibrated baselines (10.55–10.57). The four spatial

blocks are not interchangeable: in each baseline, the western quadrants (NW, SW) yield larger $\hat{q}$ and higher PICP than the eastern quadrants (NE, SE), reflecting real spatial structure in the in-domain residual distribution. No single quadrant choice produces conformal calibration competitive with the SDE on LST.

For NDVI, where deployment shift is milder, spatial-block conformal calibration improves PICP@90 from the uncalibrated baseline level of 0.82–0.83 to mean values of 0.87–0.87 across the four blocks (range 0.83 to 0.92), with corresponding mean interval scores of 0.30–0.31. The matched MLP baseline retains the lowest interval score on NDVI in every comparison (uncalibrated 0.246, SDE 0.259, conformal mean 0.301–0.310). Interpretation of these conformal results in terms of why conformal calibration fails to recover SDE-level robustness under LST deployment shift is developed in Section 5.1.

To characterize the structure of the calibration loss across the full probability range rather than at the single 90 % point, Figure 6 shows reliability diagrams for all four models on the Marion extent. Empirical coverage is plotted against nominal coverage at 5 % increments from 0.05 to 0.95. The diagonal $y = x$ indicates perfect calibration; curves below the diagonal indicate that prediction intervals are too narrow (overconfident). For LST, the SDE curve tracks the diagonal within 5–8 percentage points across the full range, while MLP and CNN curves sit a uniform 13–18 percentage points below. The miscalibration pattern is roughly flat as a function of nominal level, indicating systematic underestimation of variance rather than failure specifically at the distribution tails. For NDVI, the three learned-feature models (SDE, MLP, CNN) cluster more tightly, while the Gaussian smoother falls visibly below this cluster across the full nominal range, again consistent with its degraded performance under operational-extent deployment.

Figure 7 provides a spatial view of the SDE's operational-extent LST prediction over Marion County. Panel A shows the observed Landsat 8 summer 2020 median composite LST. Panel B shows the SDE point prediction $\hat{\mu}$, which captures the urban heat island structure with a downtown core and industrial corridors clearly distinguished from suburban and riparian areas. Panel C shows the SDE-derived uncertainty $\hat{\sigma}$; uncertainty is highest at land cover transitions and at unusual urban features such as large impervious surfaces, water bodies, and rail corridors, consistent with the model's having seen fewer such pixels in the training patch. Panel D shows the per-pixel exceedance probability $P(\mathrm{LST} > 38\,^{\circ}\mathrm{C})$ derived directly from the calibrated $(\hat{\mu}, \hat{\sigma})$ pair; this output is the spatial-statistical primitive underlying the tract-

level operational-consequence analysis reported in Section 4.6.

The SDE's mean $\hat{\sigma}$ remains approximately stable from training patch to Marion deployment: NDVI $\hat{\sigma}$ moves from 0.046 to 0.045 and LST $\hat{\sigma}$ from 1.42 °C to 1.86 °C. The matched MLP and CNN baselines exhibit similarly stable mean $\hat{\sigma}$ (NDVI: 0.038–0.039 on both patch and Marion; LST: 1.44 °C–1.46 °C on Marion). The SDE's Marion-extent LST $\hat{\sigma}$ is approximately 28 % larger than the matched baselines', a gap considerably smaller than would be required to explain the 9-percentage-point PICP@90 advantage by uniform widening alone. Detailed analysis of whether SDE calibration robustness can be attributed to trivial interval widening — and the role of the structured stochastic regularization mechanism — is developed in Sections 5.2 and 5.1.

4.6. Exceedance Probability Mapping and Operational Consequence

Because the SDE layer produces calibrated pixel-level predictive distributions $(\hat{\mu}(s), \hat{\sigma}(s))$, it enables direct computation of exceedance probabilities at any threshold: $P(\text{LST} > T \mid s) = 1 - \Phi((T - \hat{\mu}(s))/\hat{\sigma}(s))$. Figure 8 maps these probabilities across the full Marion County study area at three surface-heat screening thresholds (35, 38, 40 °C), classified into three decision-relevant bins: $P < 0.2$ (unlikely to exceed), $0.2 \leq P < 0.8$ (uncertain), and $P \geq 0.8$ (likely to exceed). These thresholds are chosen as illustrative levels for surface-heat exposure screening; LST is not air temperature, wet-bulb globe temperature, or any standard human-health index, and the thresholds are not proposed here as diagnostic of individual health risk. The middle band corresponds to pixels where calibrated uncertainty admits the possibility of exceedance even when the point estimate falls below threshold; an overconfident baseline would shrink this band and misclassify pixels as either confidently below or above threshold. The likely-to-exceed area contracts rapidly between 35 °C and 40 °C: at 35 °C, much of the developed county falls in the likely-to-exceed band; at 40 °C, only the densest impervious cores (downtown, large commercial parking, the airport apron) retain that classification.

To illustrate the operational consequence of the calibration result documented in Section 4.5, the SDE and matched-capacity baselines were each used to classify Marion County's 253 census tracts as exposed under two decision rules. The point-estimate rule flags a tract when its mean LST prediction $\bar{\mu}_{\text{tract}}$ exceeds a fixed threshold T; the precautionary rule flags a tract when the tract-mean pixel-level upper prediction bound, $\bar{\mu}_{\text{tract}} + 1.645\,\bar{\sigma}_{\text{tract}}$, exceeds T. The precautionary rule serves as a screening index that identifies

tracts where pixel-level uncertainty admits the possibility of exceedance even when the point estimate falls below threshold; this is the appropriate decision rule when underclassifying exposure carries population-health consequences. Tract-level summaries used pixel-mean $\bar{\mu}$ and pixel-mean $\bar{\sigma}$, with population from the 2020 American Community Survey 5-year estimates. Note that this construction is a precautionary screening rule rather than a formal 90 % prediction interval for the tract mean: a proper prediction interval for $\bar{Y}_{\text{tract}}$ would require modeling the spatial covariance of pixel-level prediction errors within each tract, which is reserved for a follow-on study focused on tract-level exposure inference.

Table 7 reports the result at $T = 38\,^{\circ}\text{C}$. Under the point-estimate rule, the four architectures flag 222–227 of Marion's 253 tracts, a difference of fewer than five tracts attributable to point-prediction variation between models. Under the precautionary rule, the SDE flags 28 additional tracts (representing approximately 116,000 residents, or 12.1 % of the Marion County population), while the matched-capacity MLP and CNN flag only 20–21 additional tracts (representing 81,000–83,000 residents, 8.4–8.7 %). The roughly 33,000-resident gap reflects neither prediction accuracy nor model architecture but the calibration honesty of the predictive intervals: the MLP and CNN intervals fail to widen sufficiently under operational-extent deployment to flag tracts that warrant further investigation.

5. Discussion

5.1. Summary and Position within the Remote Sensing Uncertainty Literature

The latent SDE layer provides a mechanism for encoding a Matérn-motivated spatial regularization on foundation model representations. Applied post hoc to OlmoEarth embeddings with fewer than 800,000 trainable parameters, the layer achieves prediction performance competitive with the matched MLP and CNN baselines (both of which achieve higher in-domain R^2) while providing calibrated uncertainty estimates whose robustness under spatial domain shift substantially exceeds that of conventional alternatives. On held-out spatial blocks within the training site, the SDE layer's 90 % prediction interval coverage remains 3.0–3.1 percentage points closer to nominal than the strongest baselines for NDVI, and 4.8–5.0 percentage points closer for LST. Under operational-extent deployment to the full Marion County metropolitan domain, the SDE retains PICP@90 = 82.7 % for

LST while the matched MLP and CNN baselines drop to 73.2–73.3 %, with the largest separation evident on the harder LST target where calibration honesty matters most for downstream environmental risk applications. The calibration advantage is strongest for LST, where spatial deployment shift produces substantial undercoverage in the matched baselines and where the SDE additionally achieves the lowest Marion-extent Winkler interval score; for NDVI, where deployment shift is milder, the learned-feature models perform more similarly and the matched MLP achieves a marginally lower interval score (0.246 versus 0.259) despite slightly poorer coverage preservation. The structured stochastic regularization provided by the diffusion pathway during training thus delivers its largest operational benefit on the regime where unaided baseline calibration most badly fails.

These findings contribute to a growing body of work on uncertainty quantification for remote sensing products. The classical geostatistical tradition (Atkinson and Tate, 2000; Atkinson, 2013; Foody, 2002) established that RS products require per-pixel uncertainty estimates and that this uncertainty must be spatially coherent to support downstream inference. Gaussian process approaches (Verrelst et al., 2013; Camps-Valls et al., 2016) have long provided one route to such estimates, but scale cubically with the number of locations and require kernel design choices that limit adaptivity to complex embedded representations. Recent work in deep learning UQ for EO, including Ma et al. (2021) for crop yield, Lang et al. (2023) for canopy height, and Martínez-Ferrer et al. (2022) for biophysical variables, has produced per-pixel uncertainty estimates from deep architectures, but typically requires end-to-end retraining and evaluates calibration only on in-distribution holdouts. Meyer and Pebesma (2022) made explicit that such in-distribution evaluation systematically underestimates spatial prediction error, and conformal prediction approaches (Kuronen et al., 2025; Vovk et al., 2005; Lei et al., 2018) have begun providing calibration guarantees under spatial block holdout. The Marion-extent comparison to split conformal calibration in Section 4.5.1 indicates that the marginal coverage guarantees of conformal calibration are not preserved under the magnitude of spatial shift evaluated here: across all twelve spatial-block configurations evaluated (three baselines $\times$ four calibration quadrants), conformal-calibrated MLP and CNN baselines produce Marion-extent PICP@90 values for LST below those of the uncalibrated baselines and well below the SDE. The latent SDE layer occupies a distinct position: it is post-hoc (applied to frozen embeddings without retraining), spatially structured (through the Matérn-motivated drift oper-

ator), and specifically evaluated for calibration robustness under within-site spatial shift and operational-extent deployment. Concurrent work by Xu and Markovich (2025) independently develops an SPDE-based uncertainty estimation framework with a Matérn-spectrum-derived covariance operator for graph neural networks, demonstrating improved out-of-distribution detection on graph-structured data. The two approaches share a common mathematical foundation in the Whittle SPDE formulation but operate in different settings: that work treats discrete graph domains with a prescribed Matérn kernel applied via the graph Laplacian spectrum, while the present approach addresses continuous two-dimensional spatial fields with a learned drift operator acting on frozen foundation-model embeddings. The convergent emergence of structure-informed SPDE machinery from these distinct problem settings provides independent support for the design principle.

The LST downscaling result deserves particular attention, with one important caveat. The Landsat 8 thermal band has a native acquisition resolution of 100 m, and the LST target used here was resampled from that native support to the 10 m embedding grid. The SDE layer's improvement from $R^2 = 0.814$ to 0.861 should therefore be interpreted as improved prediction of the resampled Landsat-derived target field, not as validated recovery of true 10 m thermal variation. With that caveat, the improvement pattern — operating through a 768-dimensional embedding rather than through a single vegetation index as in the classical thermal sharpening methods of DisTrad (Kustas et al., 2003) and TsHARP (Agam et al., 2007) — is consistent with the embeddings carrying sub-100 m surface-pattern information useful for predicting variation in the resampled Landsat-derived LST field beyond what VI-based sharpening exploits. Direct validation against independent finer-resolution thermal data (airborne thermal imagery, ECOSTRESS, or in situ measurements) would be required to substantiate that interpretation as a claim about true sub-100 m thermal recovery.

5.2. The Diffusion Term as Training-Time Regularizer

A key finding of the ablation analysis (Section 4.2) is that Monte Carlo sampling of the stochastic forward pass produces near-zero variance across draws at inference, yet the diffusion term is far from dispensable. Removing the diffusion term during training (the drift-only ablation) degraded prediction interval calibration from PICP@90 = 90.8% to 89.2%, and PICP@50 from 51.3% to 48.9%. The diffusion term's contribution is therefore not at inference, where the learned variance head has already absorbed the stochastic

structure, but at training time, where it serves as a structured stochastic regularizer. By perturbing the embeddings through Fokker–Planck-motivated noise during each forward pass, the diffusion forces the variance head to learn a robust heteroscedastic uncertainty surface that accounts for stochastic variation in the latent space. Without this training-time perturbation, the variance head tends to overfit to the deterministic embedding structure and underestimate uncertainty on out-of-distribution pixels.

The Fokker–Planck framework provides the design motivation for this training-time regularization. In the linear spatial-field analogue of Section 3.2, the stationary distribution is Matérn; in the nonlinear embedding-space implementation used here, the learned dynamics define an implicit stationary structure motivated by the Matérn case, but this structure cannot be characterized in closed form. The diffusion is diagonal in embedding space and applied pointwise at each pixel; it does not directly inject spatially correlated noise across pixels. What produces spatially coherent uncertainty fields is the combination of two properties: the foundation-model embeddings are themselves spatially autocorrelated as a consequence of pretraining on spatially autocorrelated imagery, and the same learned drift and diffusion networks are applied through shared parameters at every pixel. The perturbations are therefore embedding-dependent and, when applied through shared parameters to spatially autocorrelated embeddings, induce spatially coherent uncertainty fields at the output even though the perturbation mechanism itself is pointwise in embedding space. The empirical calibration robustness observed under spatial block holdout is consistent with this interpretation, though the precise theoretical connection between Fokker–Planck-motivated training perturbations and domain-shift robustness remains an open question for future investigation.

The ablation results (Section 4.2) establish a clear distinction between *calibration* and *informativeness* that is central to interpreting the empirical findings. Any model trained with the Gaussian negative log-likelihood objective, as a strictly proper scoring rule (Gneiting and Raftery, 2007), can learn to set its variance field so that nominal coverage is approximately achieved on the training distribution; this is a property of the loss function rather than of the architecture. The Baseline+NLL control demonstrated this directly, achieving in-domain PICP@90 = 90.5 % by inflating its variance estimates ($\hat{\sigma} = 0.059$, 28 % larger than the SDE's 0.046) to compensate for worse point predictions. What the SDE layer adds beyond this trivial-widening route is the spatial regularization that reduces NDVI prediction error by 22 %,

allowing the same calibrated coverage with substantially tighter, more informative intervals. For spatial decision-making, where the relevant question is not whether intervals are honest but whether they are narrow enough to be useful, this distinction is essential.

The same calibration-versus-informativeness logic addresses the natural concern that the SDE's robustness under spatial domain shift is achieved trivially by producing systematically wider intervals than the matched baselines. Three pieces of evidence from Sections 4.1–4.5 rule this out. First, the in-domain ablation shows the SDE achieves nominal calibration with tighter intervals and lower error than the trivial-widening Baseline+NLL control, so the in-domain calibration is not bought through inflated variance. Second, the SDE's mean $\hat{\sigma}$ remains approximately stable from training patch to Marion deployment (NDVI: 0.046 to 0.045; LST: 1.42 °C to 1.86 °C), and its Marion-extent LST $\hat{\sigma}$ is only 28 % larger than the matched baselines' — a gap considerably smaller than the magnitude required to explain the 9-percentage-point PICP@90 advantage by uniform widening alone. Third, the Winkler interval score (Table 5, Marion IS column) jointly penalizes width and miscoverage: on Marion-extent LST the SDE achieves the lowest interval score (9.67 °C) despite producing the widest intervals among matched architectures, because the baselines' miscoverage penalty under deployment shift exceeds their width advantage. This is strong evidence against the width-inflation explanation for the LST result. The remaining calibration advantage reflects spatial structure in the $\hat{\sigma}$ field: the SDE places uncertainty in locations where prediction error actually grows under spatial shift, whereas the baselines' $\hat{\sigma}$ fields are too flat to track shift-induced error growth.

The same mechanism explains why split conformal prediction (Section 4.5.1) fails to recover this calibration robustness when applied to the matched baselines. Conformal calibration sets a single global radius $\hat{q}$ from in-domain residuals, but the larger deployment-extent residuals on Marion (LST RMSE 2.07–2.10 versus in-domain 1.31–1.44) exceed this threshold, producing systematic undercoverage. The four-quadrant spatial-block conformal results clarify that this failure mode is not specific to any single calibration block: the in-domain residual quantile $\hat{q}$ varied substantially across NW, NE, SW, and SE calibration blocks (LST $\hat{q}$ range 1.71 to 2.39 °C for the matched MLP), reflecting real spatial structure in the in-domain residual distribution, but no quadrant choice produced PICP@90 competitive with the SDE on LST deployment. The SDE's heteroscedastic variance head, by contrast, learns spatial structure in prediction difficulty during training that transfers

more robustly under deployment-extent shift than any single post-hoc scalar calibration.

5.3. Relationship to Classical Spatial Methods

The latent SDE is best understood as a regularization principle borrowed from the SPDE–GMRF literature, not an instance of it. The classical methods of Lindgren et al. (2011), Rue et al. (2009), and Rue and Held (2005) operate on observed spatial fields with explicit likelihood functions and provide approximate Bayesian posterior inference over interpretable covariance parameters via INLA or MCMC. The latent SDE layer is not Bayesian in this sense: it is a neural architectural component trained by maximum likelihood that produces deterministic point estimates of predictive mean and variance, with the Matérn-class covariance structure entering through the Fokker–Planck stationary distribution as an implicit prior on latent dynamics rather than through posterior integration over a random field. The two approaches are therefore complementary, not substitutable: classical methods for direct spatial modeling of observed fields under explicit observation models, and the latent SDE layer for encoding a Matérn-motivated stochastic regularization within learned representations where no observation model exists in the classical sense. The implicit nonstationarity learned by the drift and diffusion networks is analogous to, but distinct from, the explicit nonstationarity parameterizations of Paciorek and Schervish (2006) and the barrier model extensions of Fuglstad et al. (2015, 2019).

An equivalent perspective: classical SPDE–GMRF methods provide the best available spatial uncertainty framework for observed fields with well-defined spatial coordinates and straightforward observation models. Deep foundation models provide the best available representation learning framework for complex multi-spectral imagery. The latent SDE layer bridges these frameworks not by performing Bayesian inference on the embeddings but by adopting the stationary-distribution structure of the classical SPDE as an architectural constraint on the deep model's outputs. The benefit of this borrowing is empirical robustness under spatial generalization, documented in Sections 4.4 and 4.5, achieved at the cost of point-estimate predictive moments rather than fully integrated posterior moments.

5.4. Limitations

Several limitations bound the applicability of these findings.

The SDE layer does not achieve the highest point prediction accuracy among the architectures tested. Both the shallow CNN and the matched-capacity MLP exceed it on R^2 and RMSE within the training domain; the SDE's advantage is specific to calibration robustness under spatial domain shift. The magnitude of this advantage (3–5 percentage points of PICP) may vary with the degree of domain shift, the land cover heterogeneity of the training patch, and the characteristics of the target variable.

The experiments use annual composite embeddings, so temporal dynamics are not addressed. Extending the framework to temporal evolution, for example to hourly or daily thermal prediction using multi-temporal foundation model embeddings, would require coupling the spatial latent SDE to an explicit temporal dimension and is a natural direction for follow-on work.

The training patch covers a single $\sim$10 km $\times$ 10 km region of Indianapolis, with the four-fold block holdout demonstrating generalization across quadrants within this domain. The Marion County evaluation (Section 4.5) extends this to an approximately 1043 km^2 metropolitan study area encompassing land cover regimes (downtown urban core, industrial corridors, agricultural fringes) not represented in proportion within the training patch, and demonstrates that the calibration advantage observed within-site persists at this scale. However, all evaluations remain within a single metropolitan area in the temperate continental climate of the Midwest. Multi-site validation across a broader range of climates (e.g. humid subtropical, semi-arid), urban morphologies, and seasonal regimes remains necessary before the calibration advantages observed here can be assumed to generalize to other deployment contexts. The operational-consequence analysis in Section 4.6 similarly is illustrative rather than exhaustive: rigorous integration of population-weighted exposure with demographic vulnerability composites, and analysis at finer spatial units such as block groups, are reserved for a follow-on study focused on environmental exposure assessment.

The Int8 quantization of OlmoEarth embeddings may limit precision in the latent space, though the strong predictive performance suggests the 256-level discretization retains sufficient signal. The $K = 10$ Euler–Maruyama integration steps add computational overhead at inference (approximately $10\times$ the cost of a single linear decoder pass), though this remains modest relative to the cost of the foundation model itself.

The Fokker–Planck bridge to Matérn precision structure is exact under linear dynamics; in the nonlinear regime required for high-dimensional embeddings, the connection serves as a design motivation rather than a the-

oretical guarantee. Well-posedness of the latent SDE under Lipschitz drift and uniform ellipticity conditions is standard (Da Prato and Zabczyk, 1992; Øksendal, 2003), but the precise relationship between specific architectural choices (drift network depth, diffusion parameterization, integration step count) and the induced implicit covariance structure in the nonlinear regime is an open question. The empirical evidence presented here suggests that the Matérn-motivated structure is a useful constraint even when its formal conditions are only approximately satisfied.

5.5. Applications and Future Directions

The latent SDE layer is directly applicable to probability exceedance mapping for sub-neighbourhood environmental exposure assessment, where the ability to compute not only "this pixel is 38 °C" but "$P(\mathrm{LST} > 40\,^{\circ}\mathrm{C}) = 0.12$" is essential for equitable risk quantification. Because the calibration robustness demonstrated here is specific to spatial shift within a single urban domain, the immediate application is to within-site exposure mapping where a model trained on part of a city can be deployed over adjacent unmapped areas with retained calibration.

The framework's inductive bias may prove most valuable in sparse-observation settings where foundation model embeddings are available everywhere but target observations are limited, a regime where spatially unstructured architectures lack sufficient data to learn spatial structure empirically but the Matérn-motivated dynamics can propagate information principally through the embedding space.

Several methodological and applied extensions are natural directions for follow-on work. On the methodological side: a comprehensive characterization of how empirical calibration robustness depends on specific architectural choices (drift network depth, diffusion parameterization, integration step count) and the resulting implicit covariance structure remains open. A formal comparison between the implicit nonstationarity learned by the drift and diffusion networks and the explicit nonstationarity parameterizations of Paciorek and Schervish (2006) and Fuglstad et al. (2015, 2019) within the embedding-space setting would clarify when learned and explicit nonstationarity converge. A hierarchical Bayesian formulation jointly inferring drift, diffusion, and observation processes could recover fully interpretable posterior distributions over Matérn parameters. A systematic comparison with modern ML-based thermal sharpening approaches would situate the framework relative to that literature.

On the applied side: integration with operational pipelines such as Prithvi-WxC and PrithviEO would extend the approach to weather and multi-temporal monitoring. Rigorous propagation of pixel-level uncertainty to spatially-correlated tract aggregates, finer-resolution analysis at the block group level, and integration with demographic vulnerability composites are the focus of a planned follow-on study on environmental exposure assessment. Multi-city validation across diverse climates and urban morphologies, and extension to city-scale and regional applications with explicit integration of population exposure data and health outcomes, are natural targets for that program.

6. Conclusion

This paper introduced a Matérn-motivated latent stochastic regularization layer for foundation model representations. Applied post hoc to OlmoEarth embeddings over a 10 km × 10 km urban training patch within Marion County, Indianapolis, with fewer than 800,000 additional parameters, the layer achieves prediction performance competitive with the matched MLP and CNN baselines while providing prediction interval coverage that degrades substantially less under spatial domain shift. Under within-site four-fold block holdout, PICP@90 falls 3.9 percentage points below nominal for NDVI versus 6.1–7.1 for matched baselines, and 7.7 percentage points for LST versus 11.0–12.8. Under operational-extent deployment to the full Marion County metropolitan domain, the SDE retains PICP@90 = 82.7 % for LST while matched MLP and CNN baselines drop to 73.2–73.3 %, despite all three architectures producing comparable point predictions. A tract-level operational consequence analysis at the 38 °C threshold illustrates that this calibration honesty translates to approximately 33,000 additional Marion County residents identified as warranting heat-exposure investigation under a precautionary decision rule. The calibration advantage reflects the structured stochastic regularization provided by the Fokker–Planck-motivated diffusion term during training, which forces the variance head to learn a heteroscedastic uncertainty surface robust to spatial generalization within and beyond the training site. The framework is complementary to, rather than a replacement for, classical SPDE–GMRF methods: it encodes a Matérn-motivated stochastic regularization in the embedding space of a learned representation, where classical methods cannot directly apply, while preserving the design

heritage of Gaussian random field theory. Directions for extension are summarized in Section 5.5.

Code and Data Availability

Source code for the latent SDE layer, training and evaluation pipelines, baseline architectures, and all figures and tables in this paper is released at `https://github.com/dpjohnsoIU/calibrated-spatial-uq` under an MIT license. A permanent Zenodo archive of the version-of-record release will be deposited and the DOI added at the proof stage. The release includes data acquisition scripts (Google Earth Engine JavaScript and Python equivalents using the `earthengine-api` client) for reproducing the NDVI and LST target rasters from Sentinel-2 Level-2A and Landsat 8 Collection 2 Level 2 Science Products over the Indianapolis training patch, configuration files specifying all model hyperparameters and training schedules, and scripts for the reliability diagrams and exceedance probability maps.

The OlmoEarth v1.0 embeddings used as input are publicly available through the Allen Institute for AI (`https://allenai.org/olmoearth`); the tile identifiers and spatial extent of the Indianapolis mosaic used in this study are listed in the release. Sentinel-2 and Landsat 8 imagery are available through Google Earth Engine under open data agreements. The Indianapolis study patch coordinates (EPSG:32616) are provided in the release to enable exact reproduction of the analysis.

Declaration of Competing Interest

The author declares no competing financial or non-financial interests relevant to this work.

Acknowledgments

Computational resources were provided by Indiana University's Research Technologies division through the Big Red 200 supercomputer and associated Slate storage. This work was supported in part by internal funding through Indiana University's Vice President of Research; the content is solely the responsibility of the author and does not necessarily represent the official views of the funding agency.

CRediT Author Statement

Daniel P. Johnson: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review and editing, Visualization, Supervision, Funding acquisition.

Declaration of generative AI

Generative AI was used for grammatical corrections and for some coding assistance. All usage has been checked for accuracy and consistency.

Table 5: Prediction interval calibration on training patch versus Marion County operational extent. PICP@90 is empirical coverage of the nominal 90 % prediction interval. Calibration robustness Δ is the patch-to-Marion delta; values closer to zero indicate that calibration is preserved under spatial domain shift to the operational extent. Marion IS is the Winkler interval score on the Marion extent (lower is better; jointly penalizes width and miscoverage, discouraging trivial interval widening).

Target	Model	Patch PICP@90	Marion PICP@90	Δ PICP@90	Patch R^2	Marion R^2	Marion IS
NDVI	SDE	0.905	0.842	−0.064	0.938	0.936	0.259
	MLP	0.903	0.818	−0.084	0.955	0.946	0.246
	CNN	0.904	0.824	−0.080	0.959	0.942	0.253
	Smoother	0.902	0.678	−0.225	0.949	0.893	0.428
LST	SDE	0.901	0.827	−0.074	0.847	0.762	9.67
	MLP	0.898	0.733	−0.165	0.892	0.770	10.57
	CNN	0.899	0.732	−0.168	0.899	0.778	10.55
	Smoother	0.897	0.794	−0.103	0.880	0.772	9.71

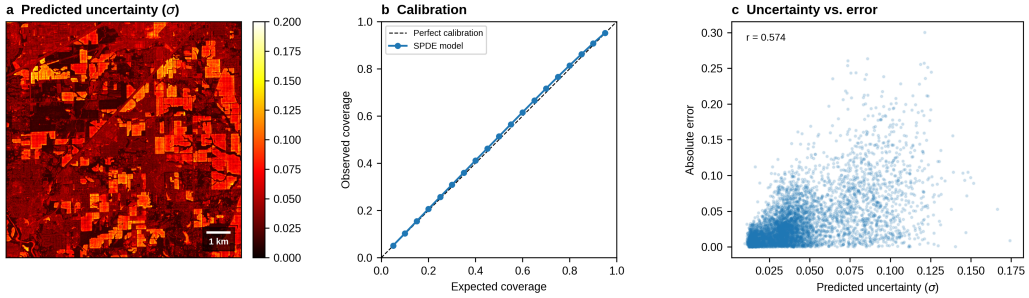

Figure 1: NDVI uncertainty estimates are calibrated and spatially coherent. **a**, Predicted uncertainty (σ) from the NLL-trained variance head. Higher uncertainty (brighter) occurs at land cover boundaries and in areas of high surface heterogeneity; lower uncertainty (darker) in homogeneous forest and agricultural areas. **b**, Reliability plot: observed prediction interval coverage versus expected coverage across quantile levels. Points fall close to the diagonal, indicating well-calibrated uncertainty (90% PI achieves 90.8% coverage; 50% PI achieves 51.3%). **c**, Predicted uncertainty versus absolute prediction error ($r = 0.574$), confirming that the model assigns higher uncertainty where it makes larger errors.

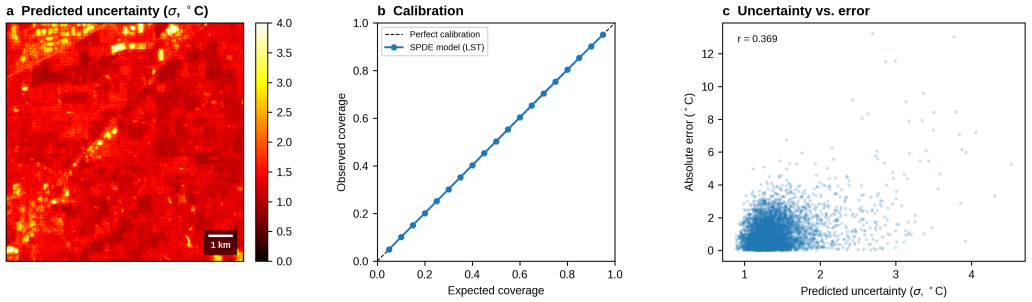

Figure 2: LST uncertainty estimates. **a**, Predicted uncertainty (σ, °C). The spatial pattern mirrors surface complexity: highest uncertainty in the heterogeneous urban core (mean $\sigma = 1.42\,°\text{C}$), lowest over homogeneous agricultural fields. **b**, Reliability plot (90% PI achieves 90.2% coverage; 50% PI achieves 50.3%). **c**, Uncertainty versus absolute error ($r = 0.369$); the weaker correlation relative to NDVI reflects the additional noise from resampling 100 m thermal data to the 10 m embedding grid.

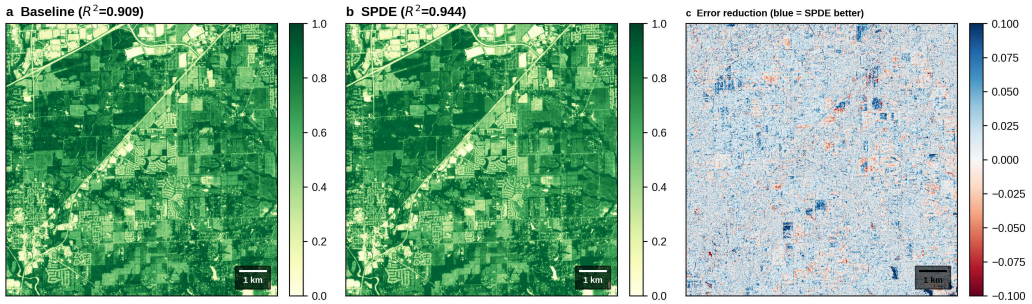

Figure 3: NDVI prediction comparison. **a**, Baseline prediction from raw embeddings ($R^2 = 0.909$). **b**, SDE prediction ($R^2 = 0.944$). **c**, Error reduction map: blue indicates pixels where the SDE layer has lower absolute error than the baseline; red indicates the opposite. The SDE layer produces the largest gains at land cover transitions: boundaries between impervious surface and vegetation, field edges, and riparian corridors.

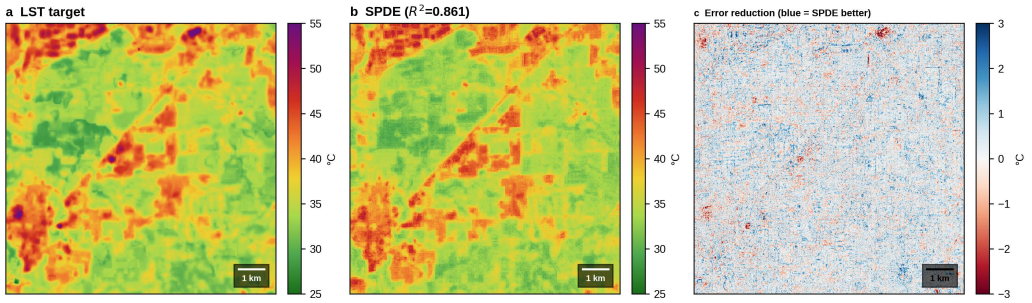

Figure 4: LST prediction comparison. **a**, LST target derived from Landsat 8 Collection 2 ST product (summer 2020 composite, distributed on the 30 m Landsat grid with 100 m native physical thermal support, resampled here to 10 m). **b**, SDE prediction ($R^2 = 0.861$). The model captures the urban heat island pattern and surface heterogeneity within the urban core across the resampled target field. **c**, Error reduction map: the SDE layer improves most in mixed residential areas where 10 m embedding structure captures surface heterogeneity associated with variation in the resampled Landsat LST field.

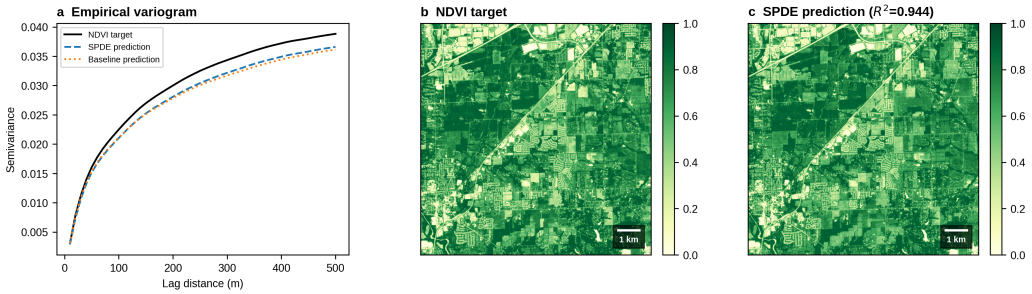

Figure 5: The SDE layer improves spatial coherence of NDVI predictions. **a**, Empirical variograms of the NDVI target (black), SDE predictions (blue dashed), and baseline predictions (orange dotted). Both models track the target's spatial structure, but SDE predictions follow the target variogram more closely across all lag distances. The target variogram rises toward a sill near 0.039 at lags of 300 m–500 m, consistent with land cover transition scales at 10 m resolution. **b**, NDVI target derived from Sentinel-2 summer 2020 composite. **c**, SDE prediction ($R^2 = 0.944$). The diagonal feature is Interstate 70; low-NDVI areas correspond to impervious surfaces in the urban core and airport area.

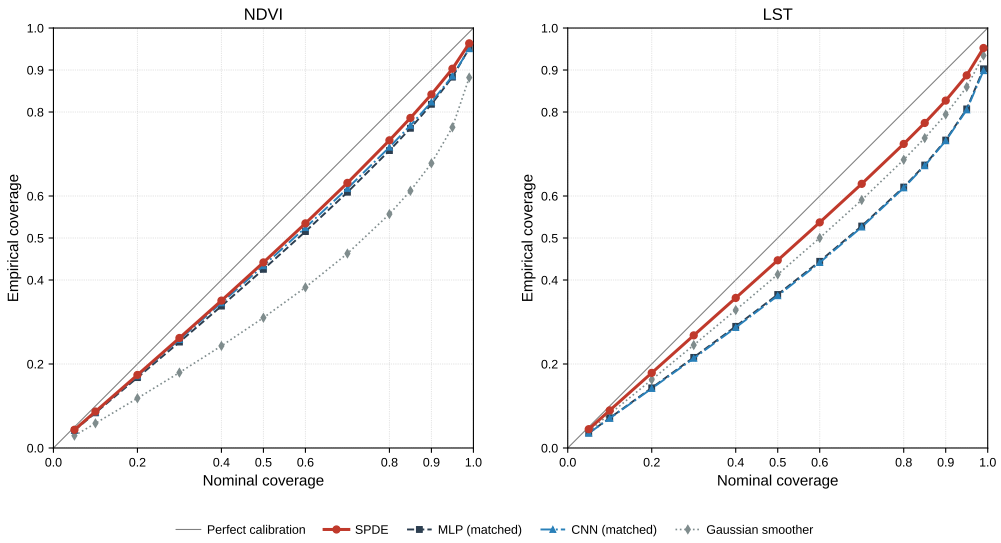

Figure 6: Reliability diagrams for the Marion County evaluation. Empirical coverage versus nominal coverage at 5 % increments for all four models, separately for NDVI (left) and LST (right). The diagonal indicates perfect calibration; curves below the diagonal indicate overconfident (too-narrow) prediction intervals. The SDE layer (red) tracks the diagonal more closely than the matched MLP and CNN baselines on both targets, with the largest separation visible on LST.

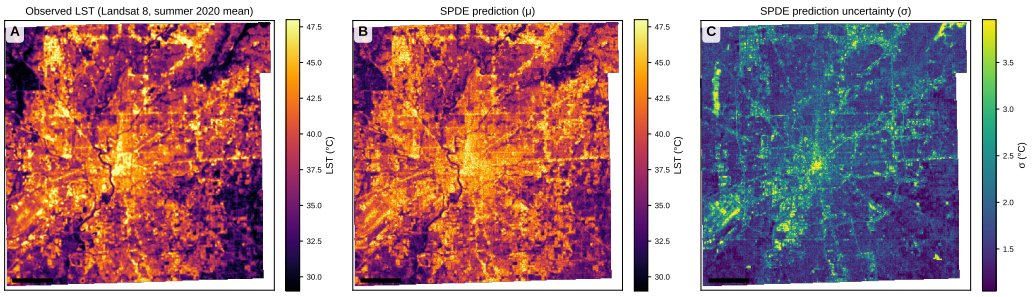

Figure 7: SDE LST inference over Marion County, summer 2020. **A**, Observed Landsat 8 median composite LST. **B**, SDE point prediction $\hat{\mu}$. **C**, SDE prediction uncertainty $\hat{\sigma}$. Panels A and B share a common color scale to permit direct visual comparison; uncertainty in panel C uses a separate scale. All panels are masked to the Marion County boundary; 5 km scale bar shown on each panel.

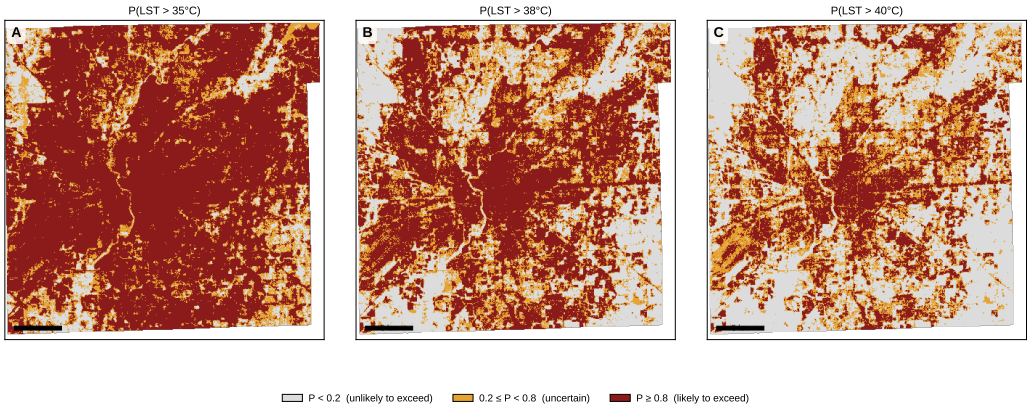

Figure 8: Calibrated exceedance probability maps over the full Marion County study area, derived from the latent SDE layer's calibrated $(\hat{\mu}, \hat{\sigma})$ pair as $P(\text{LST} > T) = 1 - \Phi((T - \hat{\mu})/\hat{\sigma})$ and classified into three decision-relevant bins: light grey ($P < 0.2$, unlikely to exceed), amber ($0.2 \leq P < 0.8$, uncertain), dark red ($P \geq 0.8$, likely to exceed). The amber band marks pixels where honest uncertainty admits the possibility of exceedance even when the point estimate falls below threshold; this is the spatial analogue of the precautionary decision rule used in the tract-level operational consequence analysis (Section 4.6). **A**, $T = 35\,^{\circ}\text{C}$ (moderate heat exposure threshold). **B**, $T = 38\,^{\circ}\text{C}$ (operational consequence threshold). **C**, $T = 40\,^{\circ}\text{C}$ (extreme heat threshold). The likely-to-exceed area contracts rapidly between $35\,^{\circ}\text{C}$ and $40\,^{\circ}\text{C}$, with the amber uncertain band concentrating along land cover transitions and at the suburban-urban interface. All panels are masked to the Marion County boundary; 5 km scale bar shown on each panel.

Table 6: Conformal prediction baselines on the Marion County operational extent, compared to the uncalibrated learned-variance baselines and the SDE. Conformal_standard uses a random 20 % of training-patch pixels as calibration set. Conformal_spatial (mean±range) reports the mean and full range of PICP@90, IS, width, and $\hat{q}$ across four runs in which the NW, NE, SW, and SE 512×512 quadrants of the training patch were each used as a spatially-coherent calibration set. Conformal radius $\hat{q}$ is reported in target units; conformal intervals are $\hat{\mu}(s) \pm \hat{q}$ at all locations. Lower interval score is better. The SDE retains the highest PICP and lowest interval score for LST across all spatial-block calibration choices.

Target	Method	PICP@90	IS	Width	$\hat{q}$
NDVI	SDE	0.842	0.259	0.148	—
	MLP (uncalibrated)	0.818	0.246	0.128	—
	MLP + conformal_standard	0.879	0.299	0.154	0.077
	MLP + conformal_spatial (mean)	0.874	0.302	0.154	0.077
	MLP + conformal_spatial (range)	[0.83, 0.91]	[0.298, 0.309]	[0.130, 0.182]	[0.065, 0.091]
	CNN (uncalibrated)	0.824	0.253	0.121	—
	CNN + conformal_standard	0.879	0.307	0.145	0.073
	CNN + conformal_spatial (mean)	0.874	0.310	0.146	0.073
	CNN + conformal_spatial (range)	[0.83, 0.92]	[0.306, 0.317]	[0.122, 0.175]	[0.061, 0.088]
	Smoother + conformal_standard	0.719	0.482	0.163	0.082
	Smoother + conformal_spatial (mean)	0.715	0.487	0.163	0.081
	Smoother + conformal_spatial (range)	[0.66, 0.78]	[0.444, 0.532]	[0.140, 0.189]	[0.070, 0.094]
LST	SDE	0.827	9.67	6.11	—
	MLP (uncalibrated)	0.733	10.57	4.79	—
	MLP + conformal_standard	0.694	11.12	4.25	2.13
	MLP + conformal_spatial (mean)	0.684	11.40	4.19	2.10
	MLP + conformal_spatial (range)	[0.62, 0.75]	[10.19, 12.66]	[3.61, 4.77]	[1.81, 2.39]
	CNN (uncalibrated)	0.732	10.55	4.75	—
	CNN + conformal_standard	0.673	11.22	4.04	2.02
	CNN + conformal_spatial (mean)	0.663	11.50	3.99	2.00
	CNN + conformal_spatial (range)	[0.59, 0.74]	[10.10, 12.89]	[3.41, 4.62]	[1.71, 2.31]
	Smoother + conformal_standard	0.730	10.55	4.49	2.25
	Smoother + conformal_spatial (mean)	0.721	10.78	4.44	2.22
	Smoother + conformal_spatial (range)	[0.66, 0.78]	[9.83, 11.85]	[3.86, 4.99]	[1.93, 2.50]

Table 7: Marion County tract-level heat exposure classification at $T = 38\,^{\circ}\mathrm{C}$, summer 2020 Landsat 8 median composite LST. "Point" = tracts flagged by $\bar{\mu}_{\mathrm{tract}} > T$; "Upper bound" = tracts flagged by the precautionary screening index $\bar{\mu}_{\mathrm{tract}} + 1.645\bar{\sigma}_{\mathrm{tract}} > T$. Additional tracts and additional population are flagged by the precautionary screening rule but not by the point-estimate rule. Marion County total: 253 tracts, 957,000 residents.

Model	Point tracts	Upper bound tracts	Additional tracts	Add'l population (residents)	% county population
SDE	222	250	28	115,987	12.1 %
MLP	226	247	21	83,042	8.7 %
CNN	227	247	20	80,772	8.4 %

References

Agam, N., Kustas, W.P., Anderson, M.C., Li, F., Neale, C.M.U., 2007. A vegetation index based technique for spatial sharpening of thermal imagery. Remote Sensing of Environment 107, 545–558. doi:10.1016/j.rse.2006.10.006.

Atkinson, P.M., 2013. Downscaling in remote sensing. International Journal of Applied Earth Observation and Geoinformation 22, 106–114. doi:10.1016/j.jag.2012.04.012.

Atkinson, P.M., Tate, N.J., 2000. Spatial scale problems and geostatistical solutions: A review. The Professional Geographer 52, 607–623. doi:10.1111/0033-0124.00250.

Beukema, P., et al., 2025. OlmoEarth: Stable latent image modeling for multimodal earth observation. arXiv preprint arXiv:2511.13655 Allen Institute for Artificial Intelligence; code at https://github.com/allenai/olmoearth_pretrain.

Camps-Valls, G., Verrelst, J., Muñoz-Marí, J., Laparra, V., Mateo-Jiménez, F., Gómez-Dans, J., 2016. A survey on Gaussian processes for Earth observation data analysis: A comprehensive investigation. IEEE Geoscience and Remote Sensing Magazine 4, 58–78. doi:10.1109/MGRS.2015.2510084.

Carrizo Vergara, R., Allard, D., Desassis, N., 2022. A general framework for SPDE-based stationary random fields. Bernoulli 28, 1–32. doi:10.3150/20-BEJ1317.

Clarotto, L., Allard, D., Romary, T., Desassis, N., 2024. The SPDE approach for spatio-temporal datasets with advection and diffusion. Spatial Statistics 62, 100847. doi:10.1016/j.spasta.2024.100847.

Cong, Y., Khanna, S., Meng, C., Liu, P., Rozi, E., He, Y., Burke, M., Lobell, D.B., Ermon, S., 2022. SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery, in: Advances in Neural Information Processing Systems 35 (NeurIPS 2022). ArXiv:2207.08051.

Da Prato, G., Zabczyk, J., 1992. Stochastic Equations in Infinite Dimensions. Cambridge University Press, Cambridge.

Foody, G.M., 2002. Status of land cover classification accuracy assessment. Remote Sensing of Environment 80, 185–201. doi:10.1016/S0034-4257(01)00295-4.

Fuglstad, G.A., Simpson, D., Lindgren, F., Rue, H., 2015. Does non-stationary spatial data always require non-stationary models? Spatial Statistics 14, 505–531. doi:10.1016/j.spasta.2015.10.001.

Fuglstad, G.A., Simpson, D., Lindgren, F., Rue, H., 2019. Constructing priors that penalize the complexity of gaussian random fields. Journal of the American Statistical Association 114, 445–452.

Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association 102, 359–378. doi:10.1198/016214506000001437.

Hamill, T.M., 2001. Interpretation of rank histograms for verifying ensemble forecasts. Monthly Weather Review 129, 550–560. doi:10.1175/1520-0493(2001)129<0550:IORHFV>2.0.CO;2.

Jakubik, J., Roy, S., Phillips, C.E., Fraccaro, P., Godwin, D., Zadrozny, B., Szwarcman, D., Gomes, C., Nyirjesy, G., Edwards, B., Kimura, D., Simumba, N., Chu, L., Mukkavilli, S.K., Lambhate, D., Das, K., Bangalore, R., Oliveira, D., Muszynski, M., Ankur, K., Ramasubramanian, M., Gurung, I., Khallaghi, S., Li, H., Cecil, M., Ahmadi, M., Kordi, F., Alemohammad, H., Maskey, M., Ganti, R., Weldemariam, K., Ramachandran, R., 2023. Foundation models for generalist geospatial artificial intelligence. arXiv preprint arXiv:2310.18660 Prithvi geospatial foundation model.

Kailath, T., 1980. Linear Systems. Prentice Hall, Englewood Cliffs, NJ.

Kingma, D.P., Welling, M., 2014. Auto-encoding variational Bayes, in: International Conference on Learning Representations.

Kloeden, P.E., Platen, E., 1992. Numerical Solution of Stochastic Differential Equations. Springer, Berlin.

Kuronen, M., Räty, J., Packalen, P., Myllymäki, M., 2025. Uncertainty quantification for forest attribute maps with conformal prediction and k-nearest neighbor method. Remote Sensing of Environment 325, 114758. doi:10.1016/j.rse.2025.114758.

Kustas, W.P., Norman, J.M., Anderson, M.C., French, A.N., 2003. Estimating subpixel surface temperatures and energy fluxes from the vegetation index-radiometric temperature relationship. Remote Sensing of Environment 85, 429–440. doi:10.1016/S0034-4257(03)00036-1.

Lang, N., Jetz, W., Schindler, K., Wegner, J.D., 2023. A high-resolution canopy height model of the Earth. Nature Ecology & Evolution 7, 1778–1789. doi:10.1038/s41559-023-02206-6.

Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L., 2018. Distribution-free predictive inference for regression. Journal of the American Statistical Association 113, 1094–1111. doi:10.1080/01621459.2017.1307116.

Lindgren, F., Rue, H., Lindström, J., 2011. An explicit link between gaussian fields and gaussian markov random fields: The stochastic partial differential equation approach. Journal of the Royal Statistical Society: Series B 73, 423–498. doi:10.1111/j.1467-9868.2011.00777.x.

Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization, in: International Conference on Learning Representations.

Ma, Y., Zhang, Z., Kang, Y., Özdoğan, M., 2021. Corn yield prediction and uncertainty analysis based on remotely sensed variables using a Bayesian neural network approach. Remote Sensing of Environment 259, 112408. doi:10.1016/j.rse.2021.112408.

Martínez-Ferrer, L., Moreno-Martínez, Á., Campos-Taberner, M., García-Haro, F.J., Muñoz-Marí, J., Running, S.W., Kimball, J., Clinton, N., Camps-Valls, G., 2022. Quantifying uncertainty in high resolution biophysical variable retrieval with machine learning. Remote Sensing of Environment 280, 113199. doi:10.1016/j.rse.2022.113199.

Matérn, B., 1960. Spatial Variation: Stochastic models and their application to some problems in forest surveys and other sampling investigations. volume 49. Meddelanden från Statens Skogsforskningsinstitut, Stockholm.

Matheron, G., 1963. Principles of geostatistics. Economic Geology 58, 1246–1266.

Meyer, H., Pebesma, E., 2022. Machine learning-based global maps of ecological variables and the challenge of assessing them. Nature Communications 13, 2208. doi:10.1038/s41467-022-29838-9.

Moran, P.A.P., 1950. Notes on continuous stochastic phenomena. Biometrika 37, 17–23. doi:10.1093/biomet/37.1-2.17.

Nix, D.A., Weigend, A.S., 1994. Estimating the mean and variance of the target probability distribution, in: Proceedings of the IEEE International Conference on Neural Networks, pp. 55–60.

Øksendal, B., 2003. Stochastic Differential Equations: An Introduction with Applications. 6 ed., Springer, Berlin.

Paciorek, C.J., Schervish, M.J., 2006. Spatial modelling using a new class of nonstationary covariance functions. Environmetrics 17, 483–506. doi:10.1002/env.785.

Rue, H., Held, L., 2005. Gaussian Markov Random Fields: Theory and Applications. Chapman and Hall/CRC, Boca Raton.

Rue, H., Martino, S., Chopin, N., 2009. Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. Journal of the Royal Statistical Society: Series B 71, 319–392. doi:10.1111/j.1467-9868.2008.00700.x.

Särkkä, S., Solin, A., 2019. Applied Stochastic Differential Equations. Institute of Mathematical Statistics Textbooks, Cambridge University Press. doi:10.1017/9781108186735.

Stein, M.L., 1999. Interpolation of Spatial Data: Some Theory for Kriging. Springer, New York.

Tuia, D., Persello, C., Bruzzone, L., 2016. Domain adaptation for the classification of remote sensing data: An overview of recent advances. IEEE Geoscience and Remote Sensing Magazine 4, 41–57. doi:10.1109/MGRS.2016.2548504.

Verrelst, J., Rivera, J.P., Moreno, J., Camps-Valls, G., 2013. Gaussian processes uncertainty estimates in experimental Sentinel-2 LAI and leaf chlorophyll content retrieval. ISPRS Journal of Photogrammetry and Remote Sensing 86, 157–167. doi:10.1016/j.isprsjprs.2013.09.012.

Vovk, V., Gammerman, A., Shafer, G., 2005. Algorithmic Learning in a Random World. Springer, New York. doi:10.1007/b106715.

Whittle, P., 1954. On stationary processes in the plane. Biometrika 41, 434–449. doi:10.1093/biomet/41.3-4.434.

Xu, F., Markovich, T., 2025. Uncertainty estimation on graphs with structure informed stochastic partial differential equations, in: Advances in Neural Information Processing Systems (NeurIPS 2025). ArXiv:2506.06907.

Zhan, W., Chen, Y., Zhou, J., Wang, J., Liu, W., Voogt, J., Zhu, X., Quan, J., Li, J., 2013. Disaggregation of remotely sensed land surface temperature: Literature survey, taxonomy, issues, and caveats. Remote Sensing of Environment 131, 119–139. doi:10.1016/j.rse.2012.12.014.