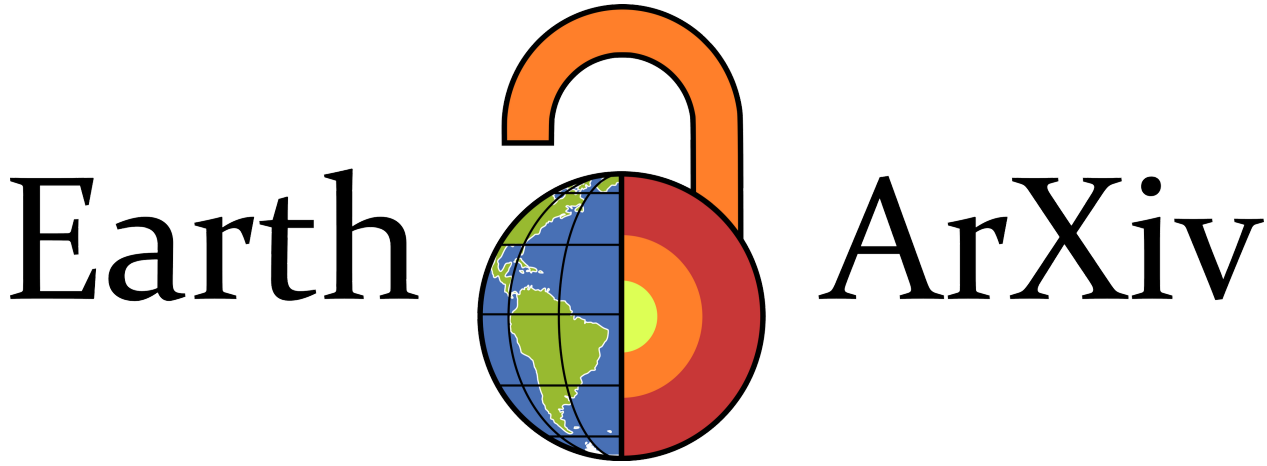

This manuscript has been submitted for publication in *Urban Climate*. Please note the manuscript has yet to be formally accepted for publication. Subsequent versions of this manuscript may have slightly different content. If accepted, the final version of this manuscript will be available via the 'Peer-reviewed Publication DOI' link on the right-hand side of this webpage. Please feel free to contact any of the authors; we welcome feedback.

Authors:

Zachary D. Calhoun (zachary.calhoun@duke.edu)
Mike Bergin (mike.bergin@duke.edu)
David Carlson (david.carlson@duke.edu)

Probabilistic interpolation of crowdsourced meteorological data for higher-resolution gridded estimates of surface air temperature

Zachary D. Calhoun[a], Mike Bergin[a], David Carlson[a,b]

[a] *Department of Civil and Environmental Engineering, Duke University, Hudson Hall 121, Durham, North Carolina, 27708, USA*

[b] *Department of Biostatistics and Bioinformatics, Duke University, Hudson Hall 121, Durham, North Carolina, 27708, USA*

Abstract

Crowdsourced air temperature data from networks like Weather Underground offer dense spatial coverage and are increasingly used to study the canopy urban heat island (CUHI) effect. However, these observations are noisy: siting conditions, environmental interference, and sensor failures introduce spatially and temporally varying bias. This complicates interpolation, limiting our ability to estimate neighborhood-level air temperature. While interpolation techniques such as kriging account for uncertainty, they do so under the assumption of homoscedasticity. Moreover, they struggle to scale beyond a few thousand observations, limiting their utility on crowdsourced data. To overcome these limitations, we develop a sparse variational Gaussian process model that accounts for heteroscedasticity, allowing us to efficiently interpolate air temperature fields with calibrated uncertainty quantification. To test our approach, we apply our model to six years of hourly data across Durham County, North Carolina, and compare predictions at held-out sensor locations with linearly-interpolated ERA5-Land. Our method improves estimates at held-out locations (MAE=0.57°C versus ERA5-Land MAE=3.20°C) and enables high-resolution analysis of CUHI patterns over space and time. We illustrate this by visualizing (1) how CUHI patterns vary with synoptic conditions, (2) differential impacts on heating and cooling demand, and (3) annual hours exceeding 35°C by neighborhood. Our method provides a scalable and statistically rigorous framework for transforming crowdsourced climate data into a gridded reanalysis product. Using this product, we can better quantify urban heat exposure and its

Email addresses: zachary.calhoun@duke.edu (Zachary D. Calhoun), michael.bergin@duke.edu (Mike Bergin), david.carlson@duke.edu (David Carlson)

impact on health and energy.

Keywords: urban heat island, citizen-science, uncertainty quantification, Gaussian process, interpolation, spatiotemporal modeling

---

1. Introduction

Extreme heat poses an increasingly consequential risk to human health and energy systems. This is especially the case in cities, where the urban heat island (UHI) effect both increases maximum temperatures and sustains these elevated temperatures through the night. The impact of this additional heat stress on health is severe. For example, during heat waves, urban populations have historically experienced increased mortality, especially among the most vulnerable sub-populations, such as the elderly (Tong et al., 2021; Laaidi et al., 2012). Energy demand is also crucially impacted by urbanization, as cooling demand is highest in urban neighborhoods, resulting in greater strain to electricity grids (Sharston and Singh, 2025). To improve resilience to extreme heat and adapt to a warming climate, cities should implement strategies that reduce the UHI effect.

A barrier to achieving this goal is in the availability of accurate urban air temperature data over space and time. Without this information, it is impossible to quantify the net benefits provided by cooling interventions. This data does not exist in many cities because official meteorological stations tend to be sited outside of urban centers, or there may only be a few stations to represent an entire city. Because the UHI effect varies sharply over space and time, higher density networks are required to capture this variability. Although reanalysis products such as ERA5-Land provide higher-resolution air temperature estimates (hourly, 9 km), they rely on sparse observational networks that underrepresent cities. As a result, urban temperatures are often underestimated, and the 9 km resolution is still too coarse to capture neighborhood-scale variability (Lee and Dessler, 2024; Zou et al., 2022; Raharjo et al., 2025).

To circumvent this problem, urban climatologists must collect their own data to adequately model the urban heat island effect. To do so, they often employ one of two methods: (1) mobile transects, in which data is collected along a predetermined route by vehicle-mounted sensors; or (2) installation of their own set of weather stations throughout the

urban environment. As a result, most urban climate datasets are either longitudinal, providing detailed temporal measurements at a few locations, or cross-sectional, offering spatial snapshots at a single point in time. This trade-off reflects inherent logistical challenges: deploying dense networks of high-fidelity sensors is prohibitively expensive, vulnerable to vandalism, and often unsustainable beyond a season or year (Muller et al., 2013). In either case, the result is a dataset collected for a specific city and typically for a restricted time period.

Crowdsourced meteorological networks, such as Weather Underground, offer an opportunity to better capture the spatiotemporal variation of urban heating. These networks provide hourly observations of temperature and other meteorological variables from a dense collection of low-cost sensors. Furthermore, in many cities, these networks are already sufficiently dense to allow for improved quantification of the UHI effect. By integrating these data into coarser resolution reanalysis products (i.e., ERA5-Land), we may produce estimates of neighborhood-scale urban meteorology that are both closer to observed conditions, and at higher resolution. However, the use of these data presents three major obstacles: (1) scalability, since data volumes grow rapidly; (2) preferential sampling, wherein sensor coverage is typically lowest in the hottest and most vulnerable neighborhoods; (3) heteroscedastic noise, as low-cost sensors are prone to biases and varying measurement error over time (Calhoun et al., 2024; Brousse et al., 2024). These obstacles make interpolation difficult, restricting our ability to quantify neighborhood heat exposure over time and its corresponding impact on health and energy.

In this manuscript, we consider whether sparse variational Gaussian processes (SVGPs) can provide a more scalable and flexible solution to this interpolation problem. Gaussian process (GP) regression excels at capturing uncertainty and may be adapted to model complex relationships through the definition of the covariance function. However, standard, or exact, GPs do not scale well beyond a few thousand data points. SVGPs redress this challenge by learning a small set of inducing points to approximate the exact GP and are efficiently trained over large datasets using stochastic gradient descent. As a result, SVGPs retain much of the flexibility of exact GP regression at a fraction of the computational cost. The machine learning community has further improved our ability to apply SVGPs through

the development of mature libraries for their implementation, such as the GPyTorch library used here (Gardner et al., 2018). Finally, we note that the SVGP approach allows for several adaptations that further improve their application to this problem. First, we exchange the standard Gaussian likelihood with a more robust likelihood, the Student's-t distribution. Second, we consider an alternative loss function that allows for heteroscedastic uncertainty quantification (Jankowiak et al., 2020). Lastly, we reduce the potential bias caused by preferential sampling by incorporating meteorology data from ERA5-Land and land-use data from the National Land Cover Database.

To demonstrate our model's utility, we develop a spatiotemporal model of air temperature in Durham County, North Carolina. Our model converges fast, in less than 5 minutes, to produce a month's worth of estimates over a 0.5° by 0.5° spatial extent. We show that the alternative loss function allows us to explicitly model uncertainty, providing insight into both the urban heat island effect as well as into microclimatic effects, too. Because our model trains quickly, we apply the technique to produce hourly estimates of air temperature from 2019-2024. As a result, we may use our model to visualize the various modes of urban heating, the effect of urbanization on cooling and heating degree days, as well as on the spatial variability of extreme heat exposure, or the hours above 35°C (95°F). Through this work, we aim to contribute a generalizable, scalable methodology for using noisy crowdsourced data to interpolate estimates of key urban climatological variables.

2. Related Work

2.1. On statistical modeling of urban air temperature

There is a vast literature of techniques employed to produce high spatiotemporal estimates of 2-meter air temperature (T_a). Until recently, much of the research has focused on estimating T_a as a 24-hour average (T_{avg}), or on the daily minimum and maximum temperatures (T_{min} and T_{max}). When aggregated as such, we note the relative success of regression-based methods (Vicente-Serrano et al., 2003; Kloog et al., 2014). However, nonlinear methods such as random forest or boosting have found success at achieving higher resolution estimates, too (Hough et al., 2020; Flückiger et al., 2022). These past works share much in common. First, they use many of the same data sources as features, such

as satellite observations of land surface temperature (LST, or T_s) from MODIS or Landsat, the normalized difference vegetation index (NDVI) to quantify vegetation, and elevation. Second, almost all techniques develop strategies to deal with a non-constant relationship between T_s and T_a: regression methods allow for diurnal random effects in a mixed-effects approach, while non-linear methods inherently capture this relationship. Lastly, they all achieve relatively high resolution, of 1km or less.

While these methods are sufficient for modeling temperature trends over time, they are insufficient for modeling the urban heat island effect. This is because the urban-rural temperature difference is most salient from the evening into the early morning, as the increased thermal inertia of urban areas results in reduced nighttime cooling (Li et al., 2024; McBroom et al., 2024). This time frame does not coincide with the afternoon, when T_{max} occurs, or morning (i.e., T_{min}). Moreover, many of these methods do not incorporate features that correspond to the urban environment. To address these gaps, hourly temperature should be modeled, and land-use added as a feature. Additionally, because higher urban T_a is a result of more than just land-use, spatiotemporal interpolation techniques can effectively capture more complex dynamics, while providing uncertainty quantification, too (Oyler et al., 2015). In other words, we can quantify the extent to which the model fails to accurately estimate temperature.

Recently, there have been attempts to model hourly air temperature (Carrión et al., 2021; Chen et al., 2022). These prior works incorporate land-use into the model with data from the National Land Cover Database (NLCD), or using the Local Climate Zone (LCZ) classification (United States Geological Survey, 2024; Stewart and Oke, 2012). Again, these methods leverage non-linear methods like Random Forest or XGBoost to produce estimates, and achieve root mean squared errors (RMSEs) ranging from 1.2-1.6°C. In both cases, the authors highlight the better predictive ability of surface characteristics (e.g., land-use) at predicting nighttime temperature, which suggests that these approaches are most applicable then, and perform worse during the daytime. The underlying cause for this poor performance is likely because these models fail to account for advection. As the sun rises, wind speeds increase, resulting in stronger advection during the daytime. This process transports heat as much as 15km downwind of urban areas, and has been shown in both high-density observational

networks as well as in physical models (Bassett et al., 2016; Belcher et al., 2015). This strengthens the case for spatiotemporal interpolation, in which we implicitly capture the effect of advection in our model given sufficiently dense observations.

2.2. Citizen science for urban heat modeling

Crowdsourced meteorological networks provide an unprecedented density of temperature observations in urban areas, and continue to grow in popularity (Chen et al., 2021). In these networks, citizen scientists purchase low-cost meteorological stations to place in their yard, and transmit their data to the network over Wi-Fi. These personal weather stations (PWSs) typically have sensors that measure air temperature, dew point/humidity, and atmospheric pressure with some stations also measuring precipitation, incoming solar radiation, and wind speed/direction. Popular networks accepting these data include Netatmo and Weather Underground. As public websites, anybody is then able to view the local weather conditions according to nearby sensors.

Due to ease of access and increasing station density, these networks are progressively being used by the research community. Initially, researchers focused on validating the quality of these data, leading to the development of several quality control techniques (Meier et al., 2017; Napoly et al.; de Baar and van Der Schrier, 2025; Fenner et al., 2021; Båserud et al., 2020). These procedures are necessary because low-cost weather stations are inherently noisy, due to lack of consistent siting information and, in some cases, faulty radiation shields. Since the development of these quality control procedures, crowdsourced meteorological data has found a variety of uses. They have been used to contrast surface urban heat islands with canopy urban heat islands (Venter et al.), reduce bias in official weather research and forecasting models (Brousse et al., 2023), and better estimate urban temperatures more generally (Castro Medina et al., 2024; Romero Rodríguez et al., 2024). Beyond these applications, crowdsourced data have been found to provide an invaluable resource for better observing urban heat advection (Brousse et al., 2022), to have the potential to improve the spatial resolution of gridded meteorological products (van Beekvelt et al., 2024), and to enable more complex applications, such as the improved estimation of urban turbulent surface fluxes (van der Meer et al., 2025).

Many of these methods were developed with the Netatmo network, which has better coverage in Europe, yet whose proprietary weather stations have neither a radiation shield nor fan to reduce the effects of solar radiation on temperature observations. In contrast, the Weather Underground network has better coverage in the United States, but the weather station model varies by user. Despite the variety of sensors used, the stations recommended on their website mostly include aspirated radiation shields, and research has confirmed that the Weather Underground data tends to be more reliable than Netatmo(Agrawal et al., 2025). In both networks, it has been found that the distribution of sensors tends to vary by neighborhood income (Brousse et al., 2024; Calhoun et al., 2024). This can be problematic in the case of measuring urban heat, as the warmest temperatures are often in the lowest income communities (Hsu et al., 2021). If there are few or no observations in these warmest areas, then we call the data preferentially sampled (Diggle et al., 2010). If insufficiently accounted for, we risk underestimating the heat exposure in lower income communities.

A recent study has developed a Bayesian hierarchical model integrating Weather Underground data into a spatiotemporal framework of urban temperatures (Marquès and Messier, 2025). Like our work, their approach combines reanalysis data, surface characteristics, and statistical interpolation. However, their reliance on latent Gaussian Markov Random Fields fit with the integrated nested Laplace approximation (INLA) restricts them to space–time separable covariance functions and still incurs substantial computation. These choices limit the model's ability to represent complex dynamics, and they do not scale easily to larger spatial domains. In contrast, we explore variational approximations that both reduce computational burden and expand model flexibility, enabling richer covariance structures, variable noise modeling, and reanalysis at broader spatial extents. Practically, variational approximations are better supported in Python, whereas INLA is better supported in R.

2.3. Sparse variational Gaussian process regression

In this manuscript, we focus on the application of sparse variational Gaussian process (SVGP) regression (Titsias, 2009). To motivate the need for this approach, we first provide background on Gaussian process regression, then explain how variational approximations allow for scalability, robustness, and modeling heteroscedastic noise. For a more complete

overview of Gaussian processes, we refer the reader to Rasmussen and Williams (2006).

2.3.1. Gaussian process regression

Gaussian process regression (GPR) produces probabilistic estimates based on the learned correlation structure of the training data as a function of distance, much like kriging, a term likely to be more familiar to readers of this journal. While in some cases, kriging and Gaussian process regression produce the same estimates, GPR does so by maximizing the likelihood objective directly. In contrast, kriging implementations often involve fitting a specified correlation function to a variogram, then solving a system of equations to learn the optimal weighted average of observations for each interpolated point. While variograms are a useful tool for understanding spatial autocorrelation, this additional step presents a challenge when the data have an unknown, non-constant mean or the correlation structure is complex. In this respect, Gaussian process regression is a more generalizable technique; while varieties of kriging exist to deal with non-constant means (e.g., universal kriging or regression-kriging), the fitting procedure for GPR does not change when we alter the model definition. Practically, we can fit complex mean and covariance functions jointly as a result.

The standard Gaussian process regression definition is given by Equation 1, in which we are modeling a random field f at a vector-valued location $\mathbf{x}$ with a (noisy) observation y. Note that the Gaussian process component is defined completely by Equation 1a through the mean and covariance functions $m(\cdot)$ and $k(\cdot, \cdot)$, respectively. However, we rarely have noiseless observations, so noise is modeled explicitly as a normally distributed random variable with variance ε^2 in Equation 1b. For consistent notation, we denote a scalar with a lowercase unbolded letter (e.g., y), a vector as lowercase bolded (e.g., $\mathbf{x}$), a matrix as uppercase bolded (e.g., $\mathbf{K}$), and a function as lowercase with parentheses (e.g., $m(\cdot)$).

$$p(f \mid \mathbf{x}) = \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \tag{1a}$$

$$p(y \mid f) = \mathcal{N}(f, \varepsilon^2) \tag{1b}$$

The term $\mathbf{x}'$ indicates that the covariance is defined with respect to other locations. In the trivial case, where $\mathbf{x} = \mathbf{x}'$, then the covariance is equal to the variance, and our knowledge of f is completely defined by Equation 1b. However, if we model the random field f at two distinct locations, $\mathbf{x}$ and $\mathbf{x}'$ such that $\mathbf{x} \neq \mathbf{x}'$, then $k(\mathbf{x}, \mathbf{x}')$ evaluates the covariance between

$f(\mathbf{x})$ and $f(\mathbf{x}')$ as a function of distance between the two locations. This parameterization of covariance allows for prediction at unobserved points, which we show in Appendix D.

If we only consider this model at a discrete set of N points $\mathbf{X}$, then the Gaussian process is a multivariate normal distribution with mean $\boldsymbol{\mu} = m(\mathbf{X})$ and covariance $\mathbf{K} = k(\mathbf{X}, \mathbf{X}')$. We denote this discretized definition over a vector of noisy observations $\mathbf{y}$ having realized values of the field $\mathbf{f}$ (i.e., $p(\mathbf{f}) = p(f \mid \mathbf{X})$, although we write $p(\mathbf{f}) = p(\mathbf{f} \mid \mathbf{X})$ to make the relationship to $\mathbf{X}$ clear). So, $p(\mathbf{f} \mid \mathbf{X}) = \mathcal{N}(\boldsymbol{\mu}, \mathbf{K})$, and $p(\mathbf{y} \mid \mathbf{f}) = \mathcal{N}(\mathbf{f}, \varepsilon^2\mathbf{I})$, where $\mathbf{I}$ denotes the identity matrix. Because $p(\mathbf{f} \mid \mathbf{X})$ and $p(\mathbf{y} \mid \mathbf{f})$ are both normal, then the marginal likelihood $p(\mathbf{y} \mid \mathbf{X})$ is normal, too, so called because we integrate out (i.e., *marginalize*) $\mathbf{f}$, as defined in Equation 2.

$$p(\mathbf{y} \mid \mathbf{X}) = \int p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X}) d\mathbf{f} \tag{2}$$

GPR models are fitted by maximizing this marginal likelihood in its log form with respect to the parameters of the mean and covariance functions. That is, if we define parameters $\boldsymbol{\theta}$ as the parameters to be learned, and $\boldsymbol{\mu}_{\boldsymbol{\theta}} = m(\mathbf{X}; \boldsymbol{\theta})$ and $\mathbf{K}_{\boldsymbol{\theta}} = k(\mathbf{X}, \mathbf{X}'; \boldsymbol{\theta})$, we obtain the marginal log likelihood given by Equation 3. Note that, by definition, $\mathbf{y}$ is defined at $\mathbf{X}$, so $p(\mathbf{y}) = p(\mathbf{y} \mid \mathbf{X})$, but that $\mathbf{X}$ impacts $\mathbf{y}$ through $\mathbf{f}$. Again, we explicitly condition on $\mathbf{X}$ in our definitions for clarity and consistency with the relevant machine learning literature.

$$\begin{aligned} \log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}) = &- \frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_{\boldsymbol{\theta}})^\top (\mathbf{K}_{\boldsymbol{\theta}} + \varepsilon^2\mathbf{I})^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\boldsymbol{\theta}}) \\ &- \frac{1}{2}\log\det(\mathbf{K}_{\boldsymbol{\theta}} + \varepsilon^2\mathbf{I}) \\ &- \frac{N}{2}\log 2\pi \end{aligned} \tag{3}$$

This objective is maximized using gradient descent on the negative marginal log likelihood, by differentiating $-\log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta})$ with respect to the partial derivatives $\partial\boldsymbol{\theta}$. We highlight this objective because it demonstrates the main computational challenge with Gaussian process regression: inverting the $N \times N$ covariance matrix $\mathbf{K}_{\boldsymbol{\theta}} + \varepsilon^2\mathbf{I}$ has cubic time complexity, $O(N^3)$. In the machine learning literature, a vast number of solutions have been proposed to overcome this this problem, including techniques to approximate the kernel (e.g., Williams and Seeger (2000); Wilson and Nickisch (2015)), accelerating computation with GPUs (e.g., Gardner et al. (2018)), and variational approximations (e.g., Titsias (2009)). However, on a single GPU, exact methods typically become computationally intractable when $N > 10,000$.

2.3.2. Sparse variational approximations

Sparse variational approximations reduce computational complexity by re-defining the Gaussian process over a set of $M \ll N$ *inducing points*, with *learned* locations $\mathbf{Z}$ and values $p(\mathbf{u}) = p(f \mid \mathbf{Z})$ (Titsias, 2009). That is, $\mathbf{u}$ is defined as the random field evaluated at $\mathbf{Z}$. The values of the field $\mathbf{f}$ at $\mathbf{X}$ are then conditioned on $\mathbf{u}$, $p(\mathbf{f} \mid \mathbf{u}, \mathbf{X}, \mathbf{Z})$. The inducing points are then learned by finding the locations and values which best approximate $p(\mathbf{f} \mid \mathbf{X})$. The variational approximation does this by defining an alternative probability distribution for $p(\mathbf{u})$, $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$, with $\mathbf{m}$ and $\mathbf{S}$ free parameters to be learned alongside the mean and covariance functions. Using this definition, we obtain $p(\mathbf{f} \mid \mathbf{X}) \approx q(\mathbf{f} \mid \mathbf{X})$ through:

$$q(\mathbf{f} \mid \mathbf{X}) = \int p(\mathbf{f} \mid \mathbf{u}, \mathbf{X}, \mathbf{Z}) q(\mathbf{u} \mid \mathbf{Z}) d\mathbf{u} \tag{4}$$

This approximation allows us to derive an alternative formulation of the marginal log likelihood with respect to the new parameters of the model. We demonstrate how to arrive at this approximation of the marginal log likelihood in Appendix A. The end-result is the so-called evidence lower bound (ELBO), shown in Equation 5, where $\mathbb{E}$ refers to the expectation, and KL the Kullback-Leibler divergence. The intuition behind this objective is that, when maximized, the first term improves model fit, and the second term penalizes values for $\mathbf{u}$ that are far from $f(\mathbf{Z})$.

$$\log p(\mathbf{y} \mid \mathbf{X}) \geq \sum_{i=1}^{N} \mathbb{E}_{q(f_i|\mathbf{x}_i)}[\log p(y_i \mid f_i)] - \mathrm{KL}[q(\mathbf{u})||p(\mathbf{u})] \tag{5}$$

This objective has two notable consequences. First, because the first term sums over an expectation of an independent and identically distributed probability distribution, we can estimate this contribution to the loss function over smaller subsets of the data (referred to as *mini-batching* in machine learning). This removes constraints due to computational resources, so the model can scale to arbitrarily large datasets. Furthermore, mini-batching enables *stochastic* gradient descent, in which cheaper estimates of the loss function are used for quicker convergence. Second, it is cheap to sample from $q(f_i)$, so non-Gaussian likelihoods for $p(y_i \mid f_i)$ become computationally tractable using Monte Carlo methods (Hensman et al., 2015).

The computational bottleneck for sparse variational Gaussian processes depends on the number of inducing points, generally scaling cubically with respect to M, depending on the

mini-batch size. When $M = N$, the approximation becomes equivalent to exact Gaussian process regression. In our experiments, on a single GPU, memory and time constraints typically appear as M exceeds $5,000$.

2.3.3. An alternative loss function

An alternative interpretation of the sparse variational Gaussian process is suggested in Jankowiak et al. (2020), expressed by the loss function in Equation 6. As the variational approximation introduces new parameters $\mathbf{Z}$, $\mathbf{m}$, and $\mathbf{S}$, we may view this approximation as a non-linear highly parameterized regressor. If we move the logarithm outside of the expectation in the ELBO:

$$\log p(\mathbf{y} \mid \mathbf{X}) \geq \sum_{i=1}^{N} \log \mathbb{E}_{q(f_i|\mathbf{x}_i)}[p(y_i \mid f_i)] - \mathrm{KL}[q(\mathbf{u})||p(\mathbf{u})], \tag{6}$$

we obtain an alternative loss function with different behavior. This objective is shown to learn a variational approximation which better captures variable uncertainty, or heteroscedastic noise. This is because the alternative likelihood allows for predictive variances to have a larger effect on the variational parameters during model training. In this manuscript, we thus refer to the standard SVGP as ELBO, and with this alternative loss, the PLL (predictive log likelihood).

2.3.4. Computational considerations

Within the past decade, several packages have emerged to facilitate fitting sparse variational Gaussian processes. We choose the Python package GPyTorch, as it includes implementations of both standard and sparse variational Gaussian processes, support for non-Gaussian likelihoods, as well as both loss functions discussed here (Gardner et al., 2018). Furthermore, because this package is built on PyTorch, GPU compatibility and autodifferentiation are natively supported, allowing for efficient implementations of the stochastic gradient descent algorithm. Lastly, by integrating with PyTorch, the method described here may be readily extended as a component in more complex architectures (e.g., using neural networks).

3. Methods

3.1. Data collection and pre-processing

Three sources of data are used in this model: Weather Underground (WU) for temperature observations, ERA5-Land for a reference temperature, which serves to capture larger scale meteorological effects, and surface characteristics from the National Land Cover Database (NLCD). We summarize the data acquisition and pre-processing pipeline in Figure 1. Briefly, we collect the data over the desired region and time-period. The Weather Underground station locations $\mathbf{s}$ (latitude, longitude) inform the linear interpolation of ERA5-Land data, as well as the processing of the land-use data to construct the feature $\zeta(\mathbf{s}, t)$, representing land-use at location $\mathbf{s}$ at time t (in hours), and the ERA5-Land estimated 2-meter air temperature $\tilde{y}(\mathbf{s}, t)$. The interpolated ERA5-Land temperature is also incorporated into the quality control process to remove anomalous data from the citizen-science dataset. We provide additional details on these processes, below.

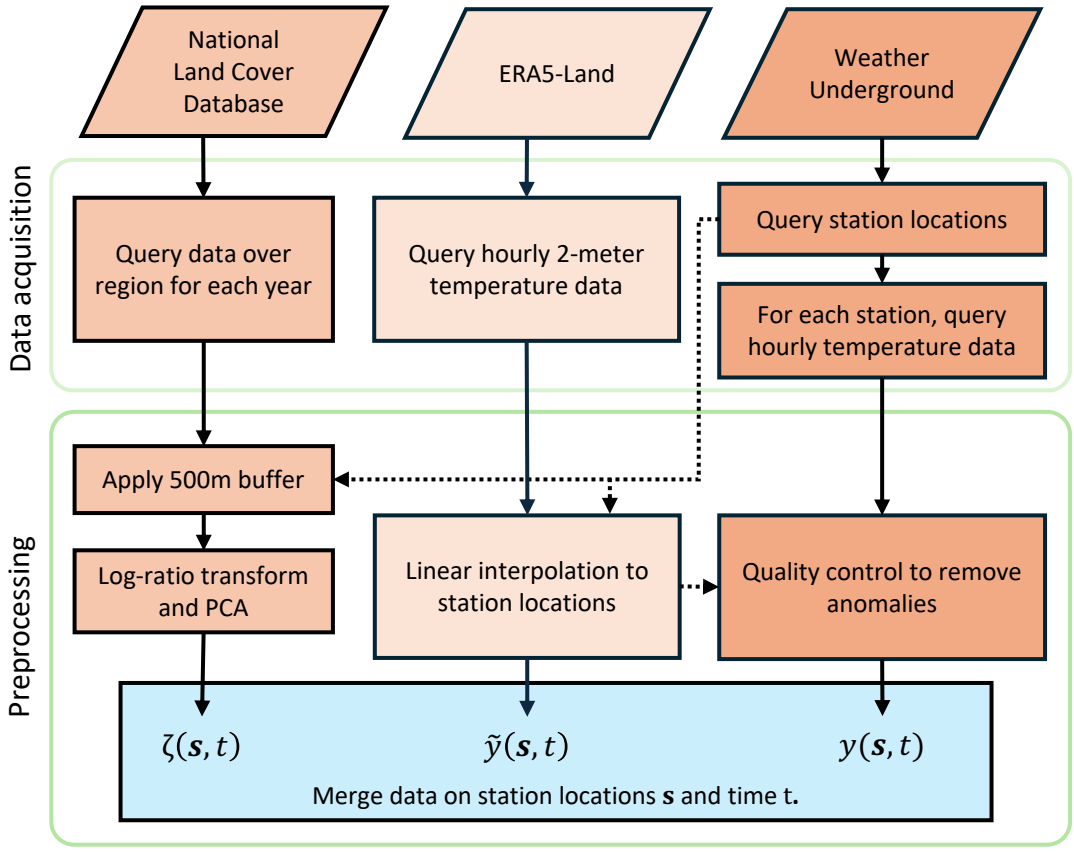

Figure 1: Summary of data acquisition and pre-processing pipeline.

Weather Underground. We collected hourly data from the Weather Underground API for sensors in Durham County, North Carolina, from January 2019 to December 2024. Sensors were located between latitudes 35.8°–36.3° N and longitudes 78.6°–79.1° W. During this period, the number of active sensors increased from roughly 100 in 2019 to over 300 in 2024. For initial quality control, we discarded observations flagged by Weather Underground's internal checks [1]. Most observations were timestamped near the end of the hour, so we removed those with a minute value less than 50 and rounded the remaining timestamps to the nearest hour to align with ERA5 data. For additional quality control, we adapted the approach of Fenner et al. (2021). Unlike their study, we had access to reference temperatures from ERA5-Land, which account for the adiabatic lapse rate in Durham's relatively uniform topography and therefore remove the need to adjust for elevation directly.

This quality control procedure proceeds through multiple steps to remove potential anomalies, in accordance with the adapted approach. First, for each hour of data, we subtract the ERA5-Land air temperature from the observed air temperature at the station's coordinates. This serves to correct implicitly for differences in elevation and for larger scale meteorological effects. Second, we use a z-scoring filter technique on these differences to remove stations which are clearly affected by radiation, with adapted critical values using a t-distribution if the number of sensors were fewer than 200. In rare cases where network issues caused fewer than 20 sensors to report for a given hour, we dropped that hour completely. Lastly, if a sensor fell outside the statistical filter more than 20% of the time, we drop that sensor completely, too. We note that, while this improves the quality of the data fed into the model, our method is robust to outliers. However, by excluding the obviously anomalous data, we achieve uncertainty quantification that is more aligned with true spatiotemporal variability.

[1] Some of these internal checks are summarized at `https://support.weather.com/s/article/Are-personal-weather-stations-on-Wunderground-monitored-for-quality`. Generally, if a station reports values that vary too much from its neighbors, the API sets a field `qcStatus` to -1, while new stations have a `qcStatus` value of 0. We only keep observations where this value is 1, to indicate the observation passes this check. Because the technical details of the procedure are unclear, further quality control is needed.

ERA5-Land. ERA5-Land is a reanalysis dataset developed from the European Centre for Medium-Range Weather Forecasts (ECMWF) that combines weather models with Earth observations to provide hourly meteorological data at a 0.1° resolution (or, about 9km) (Muñoz-Sabater et al., 2021). This data source contains estimated 2-m air temperature, but also data on wind speed, wind direction, and dew point as well as a host of other variables. For the purposes of this study, we only use 2-m air temperature (T_a). We linearly interpolate ERA5-Land observations to the sensor coordinates to provide a reference temperature. When we merge the ERA5-Land and WU data, we find that the WU temperature data are generally centered on the ERA5-Land data. However, bias appears in the extremes. At warmer temperatures, the WU observations are higher than the ERA5-Land estimates, and at colder temperatures, WU observations are lower. By incorporating ERA5-Land T_a into our model, we effectively develop a downscaling technique that uses noisy personal weather station data to achieve higher spatial resolution.

NLCD. Lastly, we use the National Land Cover Database (NLCD) for an annually updated map of land-cover classifications at 30m spatial resolution (United States Geological Survey, 2024). Because this dataset was last updated in 2023, we assume minimal change between 2023 and 2024, imputing 2023's data for 2024. As commonly done in prior work, we average classes over a 500m radius around each station to yield an informative set of features to characterize the local environment (Schatz and Kucharik; Oke, 2006). However, in contrast with prior work (cf. Carrión et al. (2021)), we further pre-process these percentages for optimal usage in a Gaussian process. Namely, we prefer lower-dimensional covariates to maintain scalability, prevent overfitting, and reduce the impact of preferential sampling. So, we apply the centered-log transformation, then principal component analysis (PCA) to extract a single continuous feature (PC1) that explains the most variance in land-use (Pedregosa et al., 2011). Note, the centered-log ratio transformation is a common pre-processing step when applying principal component analysis to compositional data (i.e., percentages summing to 1) (Pawlowsky-Glahn and Egozcue, 2006).

To produce maps, we aligned NLCD percentages with the interpolated ERA5-Land dataset at a 0.005° resolution, then transformed the gridded percentages using the same centered-log ratio and PCA transformation to extract the first principal component over a

grid, as shown in Figure 2.

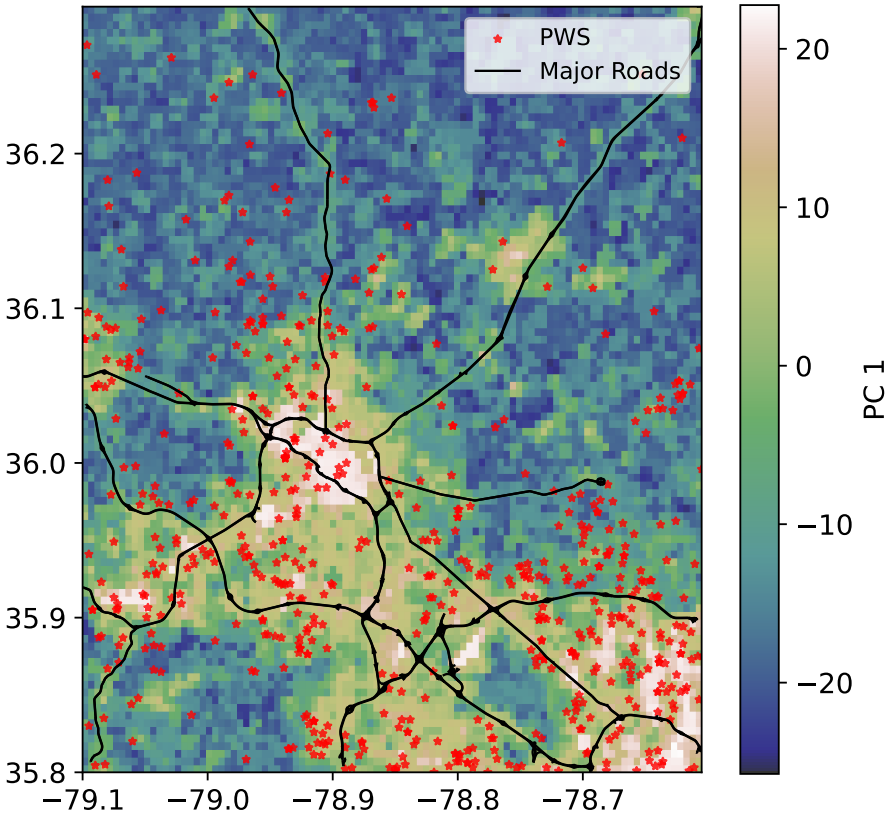

Figure 2: Principal component 1, extracted from NLCD . Sensor locations are denoted PWS, with major roads provided for reference. Pixels shown are 0.005° by 0.005°(approximately 500 by 500 meters). Higher PC1 values reflect more urbanized areas, with negative values reflecting rural areas.

3.1.1. Final dataset

Lastly, we concatenate all data to produce a set of N observations $(\mathbf{X}, \mathbf{y}) = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$ referring to the covariates, $\mathbf{x}_i$, and temperature, y_i for each observation, respectively. In addition to spatial coordinates ($\mathbf{s}$), time (t, with units hours), ERA5-Land 2m temperature ($\tilde{y}$), and the land-use principal component ζ_i, we define a feature $\boldsymbol{\phi}(t) = [\phi_1(t), \phi_2(t)]$ to capture correlations in the diurnal cycle of urban heating, consisting of the sine and cosine

transform,

$$\phi_1(t) = \sin\left(\frac{2\pi t}{24}\right), \quad \phi_2(t) = \cos\left(\frac{2\pi t}{24}\right).$$

Thus, the complete covariate set for each observation is: $\mathbf{x}_i = [\mathbf{s}_i, t_i, \boldsymbol{\phi}(t_i), \tilde{y}_i, \zeta_i]$. Observations are aggregated into files for each month-year for model fitting.

3.2. Model definition

To completely describe our modeling approach, we first define the mean and covariance functions. Next, we outline the variational approximation, the likelihood, and the loss function. We then provide details on the algorithm used for training. We provide a graphical summary of this model definition in Figure 3. Briefly, the mean function on land-use and ERA5-Land is necessary to set a reference temperature in the absence of useful information from the crowdsourced data; this reduces the potential impact of preferential sampling. The variational approximation improves scalability over exact Gaussian process regression, and the likelihood and loss function enable the model to robustly handle heteroscedastic noise.

3.2.1. The mean function

The mean function is a linear regression of the interpolated ERA5-Land T_a and the NLCD principal component, along with a bias term. This linear regression is initialized with a coefficient of 1 on ERA5-Land, and a small positive number on the principal component (i.e., 0.1). By initializing our model this way, we center the observations on the ERA5-Land data. This allows for deviations from the reanalysis data to be modeled through the covariance function.

3.2.2. The covariance function

The covariance function is defined with two goals in mind: (1) to capture the diurnally periodic effect of land-use on temperature; and (2) to capture random spatiotemporal variability from this diurnal pattern. The mathematical form of this covariance function is adapted from (Malings et al., 2017), which used a similar approach to model the surface urban heat island. In contrast with their approach, we have a reference temperature from ERA5-Land, which allows us to treat the hourly averaged temperature through the mean

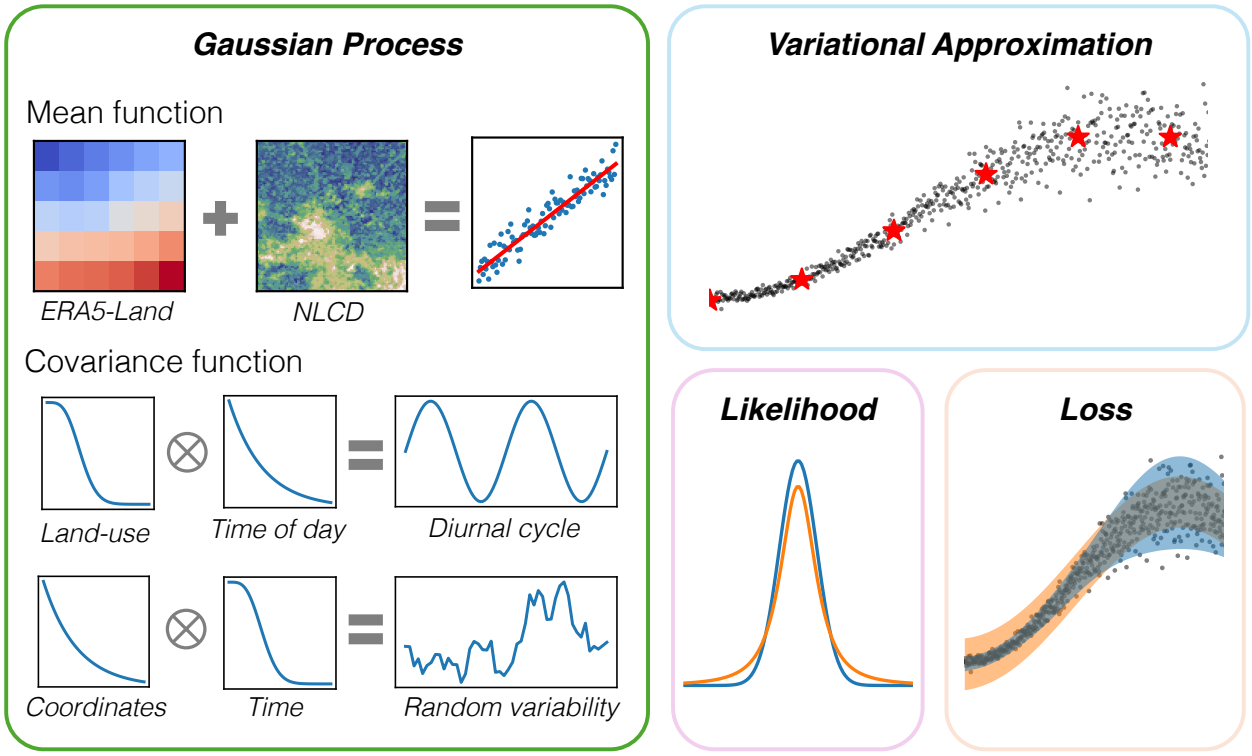

Figure 3: Model summary. The Gaussian process is defined with a mean and covariance function. The mean function centers model estimates on ERA5-Land, with an adjustment for the average effect of land-use on temperature. The covariance function captures the diurnal effect of land-use, along with random spatiotemporal variability. The variational approximation allows for scalability, by instead learning a set of inducing points (here, red stars) to approximate the observed data. These inducing points are learned over all covariates (i.e., space, time, land-use). We consider an alternative likelihood to improve robustness to outliers, and an alternative loss function to better quantify variable uncertainty.

function rather than through the covariance function. Lastly, we emphasize that our model is fitted directly on the data, rather than the covariance parameters set beforehand, as in their approach. This direct fitting capability allows for easier application to other cities. The complete covariance function is defined as:

$$K = K_1(\boldsymbol{\phi}(t), \boldsymbol{\phi}(t'), \zeta, \zeta') + K_2(t, t', \mathbf{s}, \mathbf{s}')$$

K_1 captures the structured diurnal cycle of temperature as it varies by land-use type, and K_2 captures short-term deviations from this cycle (that are also not explained by the ERA5-Land dataset). We define K_1 as a multiplicative kernel to capture the interaction of land-use

and time of day on the diurnal cycle:

$$
\begin{aligned}
K_1 &= \sigma_1^2 K_a(\boldsymbol{\phi}(t), \boldsymbol{\phi}(t')) K_b(\zeta, \zeta') \\
K_a &= \alpha_1 \exp\left(- \frac{||\boldsymbol{\phi}(t) - \boldsymbol{\phi}(t')||_2}{\ell_a} \right) + \alpha_2 \\
K_b &= \left(1 + \frac{\sqrt{3}||\zeta - \zeta'||_2}{\ell_b} \right) \exp\left(- \frac{\sqrt{3}||\zeta - \zeta'||_2}{\ell_b} \right)
\end{aligned}
$$

In other words, locations with similar land-use characteristics are expected to exhibit similar temperature patterns at the same time of day. Note that σ_1^2 denotes the amount of variance explained by this term, K_a corresponds to an exponential kernel (i.e., Matérn with $\nu = 0.5$) plus added constants α_1 and α_2. K_b is a Matérn kernel with $\nu = 1.5$. $||\cdot||_2$ denotes the Euclidean distance. This builds on prior approaches; for example, Marquès and Messier (2025) use a linear model with land-use coefficients that vary by hour of day. Physically, this reflects how the surface energy balance exhibits diurnal periodicity driven by solar radiation and moderated by the thermal properties of different land surfaces.

Because the surface energy balance varies from day to day, we include a second term, K_2, to capture short-term fluctuations:

$$
\begin{aligned}
K_2 &= \sigma_2^2 K_c(t, t') K_d(\boldsymbol{s}, \boldsymbol{s'}) \\
K_c &= \left(1 + \frac{\sqrt{3}||t - t'||_2}{\ell_c} \right) \exp\left(- \frac{\sqrt{3}||t - t'||_2}{\ell_c} \right) \\
K_d &= \exp\left(- \sqrt{(\boldsymbol{s} - \boldsymbol{s'})^\top \mathbf{A}^{-2} (\boldsymbol{s} - \boldsymbol{s'})} \right) + \alpha_3 \\
\mathbf{A} &= \begin{bmatrix} \ell_{d1} & 0 \\ 0 & \ell_{d2} \end{bmatrix}
\end{aligned}
$$

Again, we have σ_2^2 to denote the strength of this short-term variance term, K_c as a Matérn kernel with $\nu = 1.5$, and K_d as an exponential kernel, with dimension-specific length-scales, as parameterized by $\mathbf{A}$. Note that the dimension-specific length-scales obviate the need to convert the coordinate reference system of the data; instead, we learn correlation in each direction. Additionally, the square root is merely the Euclidean distance with respect to the linear transformation $\mathbf{A}$. This kernel represents fluctuations that persist longer than a

single time step, as implied by the Matérn$_{\nu=1.5}$ kernel, while remaining spatially localized. Physically, this term captures short-lived meteorological variability, such as transient cloud cover or localized wind effects, that is not accounted for by the large-scale ERA5-Land fields.

In the SVGP framework, the combined kernel allows the inducing points to approximate both components of variability: the structured diurnal cycle when surface-energy balance effects dominate, and the short-term deviations driven by larger-scale meteorological variability.

3.2.3. Methods for robust Gaussian processes

We consider two techniques to improve the robustness of our model.

Student's t-likelihood. Given a learned Gaussian process f, we typically assume that our observations y_i are normally distributed around f_i, so $p(y_i \mid f_i) = \mathcal{N}(f_i, \varepsilon^2)$. Instead, we may decide to improve our model's robustness to outliers by considering the noise as following a Student's t-distribution, where an additional parameter for the degrees of freedom is learned during model fitting. The variational approximation allows us to account for this case.

Predictive Log Likelihood. We also test the use of the alternative variational objective, the predictive log likelihood (PLL) previously discussed (Jankowiak et al., 2020). We consider this objective with both the Normal and Student's t-distribution for the likelihood, $p(y_i \mid f_i)$.

3.2.4. Model fitting

We fit a separate model over a month's worth of data (i.e., one model for January 2019, one for February 2019, one for January 2020, etc.). This offers several advantages over fitting one model to the entire dataset. First, performance is limited by the number of inducing points. By fitting a set of independent models to smaller datasets, we can do so in parallel, using fewer inducing points per model and producing the gridded estimates over the whole dataset faster. Second, the model definition assumes a stationary mean and covariance function. By fitting monthly, we do not rely on stationarity assumptions over larger time-scales. This presents a trade-off, where we may prefer reducing the time-scale further to better capture non-stationary behavior. However, we reason that further reducing the model time period may inhibit calibration, as we expect noisier observations in the

afternoon than at night. With monthly time periods, we expect the model sees enough daily cycles to capture this relationship.

We summarize the fitting procedure in Algorithm 1. Given the monthly training data and algorithm hyperparameters, the algorithm first initializes inducing points by randomly sampling M points from the training set. While the inducing point locations are learned, the training procedure does not allow for them to drastically change (due to the learning rate, discussed below). However, the starting locations should represent the domain, both temporally and spatially. Random initialization accomplishes this objective. The model is then defined with the mean and covariance function, along with these inducing points and the variational distribution. In GPyTorch, we use the Cholesky form of the standard variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$, which defines the covariance matrix using the lower triangle for computational stability.

We then specify the likelihood and objective function, noting that the likelihood is a separate object from the SVGP object in GPyTorch, with both objects passed to the objective function and optimizer. In GPyTorch/PyTorch, the optimizer tracks gradients and updates parameters. We use a standard optimizer: the Adam optimizer with a learning rate of $\eta = 0.01$; however, we decrease the learning rate according to a Cosine Annealing schedule, finding that this stabilizes convergence. This is because the learning rate dictates how much we allow the model parameters to change during each iteration of the algorithm. As the algorithm progresses, decreasing the learning rate allows for the model parameters to converge to their local optima.

We iterate over the monthly dataset $T = 10$ times (i.e., 10 epochs) using a batch size of $b = 512$. On lines 8-12, we explicitly show how GPyTorch fits the objective similarly to fitting PyTorch models. That is, the optimizer resets parameter gradients at the beginning of each iteration, tracks the gradients during the forward pass and loss computation, then uses the gradients to update parameter values for the next iteration, scaled by the learning rate.

Algorithm 1 Training an SVGP with GPyTorch

Require: Training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n}$, number of inducing points m, number of iterations T, batch size b, learning rate η, likelihood p, objective $\mathcal{L}$

Initialize inducing points $\mathbf{Z} = \{\mathbf{z}_j\}_{j=1}^{m}$ ▷ random subset of X

Define `model` $\leftarrow \mathcal{SVGP}(m_{\boldsymbol{\theta}}(\mathbf{x}), k_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{x}'), \mathbf{Z}, q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S}))$

Define `likelihood` $\leftarrow p(y \mid f, \varepsilon^2)$ ▷ p: Gaussian or Student's-t

Define `objective` $\leftarrow \mathcal{L}$(`model, likelihood`) ▷ $\mathcal{L}$: ELBO or PLL

Define `opt` $\leftarrow$ `optimizer`$(\{\boldsymbol{\theta}, \mathbf{Z}, \mathbf{m}, \mathbf{S}, \varepsilon^2\},\ \eta)$ ▷ For updating parameters

for $t \in \{1, 2, ..., T\}$ **do**

  for mini-batch $(\mathbf{X}_k, \mathbf{y}_k) \in \mathcal{D}$ **do** ▷ k subsets of $\mathcal{D}$, each with size b

    `opt.zero()`

    $q(\mathbf{f}_k) \leftarrow$ `model`$(\mathbf{X}_k)$ ▷ Forward pass, $q(\mathbf{f}_k)$

    `loss` $\leftarrow -$`objective`$(\mathbf{y}_k, q(\mathbf{f}_k))$ ▷ Compute loss

    `loss.backward()` ▷ Backward pass, $\nabla\mathcal{L}$ w.r.t. $\boldsymbol{\theta}, \mathbf{Z}, \mathbf{m}, \mathbf{S}, \varepsilon^2$

    `opt.step()` ▷ Update parameters

  end for

end for

3.2.5. Prediction

Given the fitted model, we explain how the model predicts in Algorithm 2. The test points depend on the prediction objectives. For model validation, we predict on the data from the held-out locations. For producing interpolated fields, we construct a grid of coordinates. These coordinates are then used to interpolate temperature from ERA5-Land, and derive the principal component from the NLCD to create test points $\mathbf{X}_*$. Because this dataset is large, we batch predictions by the hour. For each mini-batch, note that the model separately generates a distribution from the random field, then the likelihood generates the "noisy" distribution. When the likelihood is Gaussian, GPytorch returns the distribution, which has the analytical form given in Appendix D. For non-Gaussian likelihoods, samples are generated from the random field according to the likelihood, from which we compute mean and quantile statistics.

Algorithm 2 Prediction with GPyTorch

Require: `model, likelihood`, test points $\mathbf{X}_*$ (e.g., over a uniform grid), and batch size b.

1: **for** mini-batch $(\mathbf{X}_{*,k}) \in \mathbf{X}_*$ **do**
2: $q(\mathbf{f}_{*,k}) \leftarrow \texttt{model}(\mathbf{X}_{*,k})$
3: $p(\hat{\mathbf{y}}_{*,k}) \leftarrow \texttt{likelihood}(q(\mathbf{f}_{*,k}))$ ▷ Samples from the learned distribution
4: **end for**

Ultimately, the constructed grid can be arbitrarily defined at the desired spatial resolution, as the model estimates continuous fields. However, the continuity of any single sample from the model is computationally limited. This is because sampling from a multivariate normal distribution requires inversion of the full covariance matrix. To produce spatially continuous samples, we advise mini-batching by time (and for temporally continuous samples, mini-batch by location).

3.3. Experimental set-up

The goals of our experiments are: (1) to determine the optimal number of inducing points to achieve good performance; (2) to determine if the alternative losses and likelihoods improve robustness. To meet these objectives, we consider 500, 1000, 2000, and 5000 inducing points. For each number of inducing points, we consider 4 models with: (1) Gaussian likelihood and

ELBO loss; (2) Gaussian likelihood and PLL loss; (3) Student likelihood with ELBO loss; and (4) Student likelihood and PLL loss. All models are trained using GPyTorch on NVIDIA GeForce RTX 2080 GPUs (Gardner et al., 2018).

To validate, we test our model's ability to predict temperature at unobserved locations. To do this, we hold out 20% of the sensors for validation. We emphasize that this is not the same as random cross-validation, in which 20% of the data is withheld. This validation set is not seen by the model during fitting, and because the locations themselves are withheld, the metrics described below estimate the model error over the spatial domain. In contrast, random cross-validation, in which we randomly select validation points over time and space, would provide a biased estimate of performance because locations have temporally correlated observations. The same withheld locations are used across all models, allowing for direct comparison of results between models.

3.3.1. Metrics

Our metrics include the mean absolute error (MAE), mean squared error (MSE), negative log predictive density (NLPD), and quantile calibration error (QCE) at the 50%, 75%, and 95% confidence intervals. The purpose of using the NLPD and QCE metrics is to understand the ability of the model to quantify uncertainty. With all metrics here, lower values indicate better performance.

Formally, the negative log predictive density measures the model's estimated log likelihood of the true observations (Levi et al., 2022). That is, given our model's estimates $\mathbf{f}$ at data points $\mathbf{X}$, we calculate this metric with respect to our held-out data points $\mathbf{X}_*$ with true values $\mathbf{y}_*$:

$$\text{NLPD}(\mathbf{y}_* \mid \mathbf{f}, \mathbf{X}, \mathbf{X}_*) = -\log p(\mathbf{y}_* \mid \mathbf{f}, \mathbf{X}, \mathbf{X}_*) \tag{7}$$

In the case of a Gaussian likelihood, the above metric is trivial to calculate. With a Student's-t likelihood, however, we must approximate, due to the intractability of the non-conjugate likelihood. We do this by drawing S samples at each $\mathbf{x}_* \in \mathbf{X}_*$, yielding the set of estimates

$\mathbf{f}_*$. This procedure is given below:

$$\text{NLPD}(\mathbf{y}_* \mid \mathbf{f}, \mathbf{X}, \mathbf{X}_*) = -\log \int \text{Student}(\mathbf{y}_* \mid \mathbf{f}_*)\mathcal{N}(\mathbf{f}_* \mid \mathbf{f}, \mathbf{X}_*)d\mathbf{f}_* \tag{8}$$

$$\approx -\log\left(\frac{1}{S}\sum_{i=1}^{S}\text{Student}(y_* \mid \mathbf{f}_{*,i})\right) \tag{9}$$

$$= -\log\left(\sum_{i=1}^{S}\text{Student}(y_* \mid \mathbf{f}_{*,i})\right) + \log(S) \tag{10}$$

The quantile calibration error allows us to gain further insight into model calibration by specifying the error at several confidence interval values (Küppers et al., 2023). This calculation is merely the absolute value of the empirical coverage less the theoretical coverage, where the empirical coverage is the percentage of samples that fall within a given (two-sided) quantile interval. For example, a quantile calibration error of 1% at the 50% level would mean that 49% or 51% of the data is within the theoretical interval. As with the NLPD, the QCE calculation for the Gaussian likelihood is trivial. For the Student's-t likelihood, we again construct S theoretical t-distributions using the model estimates $\mathbf{f}_*$, then average the coverage over samples to calculate the empirical coverage.

4. Results

4.1. Evaluating model performance

We show model performance in Table 1 and its trade-off with scalability in Table 2. The complexity of SVGPs depends on the number of inducing points. Because this number also impacts scalability, we would like to understand the complexity-scalability trade-off. In Table 1, we show that across all months, the G-PLL model achieves the best performance overall, if we consider mean absolute error and the negative log predictive density as the most important metrics. These metrics suggest that this combination provides the best calibrated uncertainty quantification while being slightly more robust to any anomalous observations. All models improve MAE and MSE considerably, showing the value of using crowdsourced data.

However, if we observe the NLPD across all months, we see that the S-PLL model may be preferable during the summer months, as shown in Figure 4. While the G-PLL model excels

Model	MAE	MSE	NLPD	QC50(%)	QC75(%)	QC95(%)
G-ELBO	0.57 ±0.07	**0.59 ±0.13**	1.05 ±0.09	9.37 ±4.0	5.84 ±2.79	0.73 ±0.59
S-ELBO	0.57 ±0.07	**0.59 ±0.13**	1.12 ±0.12	4.15 ±2.58	**2.29 ±1.84**	**0.6 ±0.53**
S-PLL	0.57 ±0.07	0.63 ±0.15	0.96 ±0.15	**2.9 ±1.87**	2.92 ±2.14	1.16 ±1.27
G-PLL	**0.56 ±0.07**	0.60 ±0.13	**0.83 ±0.33**	3.76 ±2.02	3.71 ±2.27	1.4 ±1.53
ERA5-Land	3.20± 0.81	1.37± 0.17	-	-	-	-

Table 1: Experimental results. QC# represents the quantile calibration error. The standard deviation of the values over the months is provided after the ± symbol.

at capturing wintertime spatial variability and uncertainty, it fails to perform as well in the summer. A number of factors may cause this phenomenon. One, if the statistical filter did not adequately remove potentially biased sensors, then the Gaussian likelihood would result in a predictive variance that is too high. Second, like many places in the Southeast, summers are marked by bursts of localized heavy rain which cause rapid cooling. Because this model does not explicitly account for the effect of precipitation on temperature, we would have high uncertainty in locations where only a subset of the sensors experience rain. Lastly, there is more spatial variability in the summer in general, which would result in a higher NLPD.

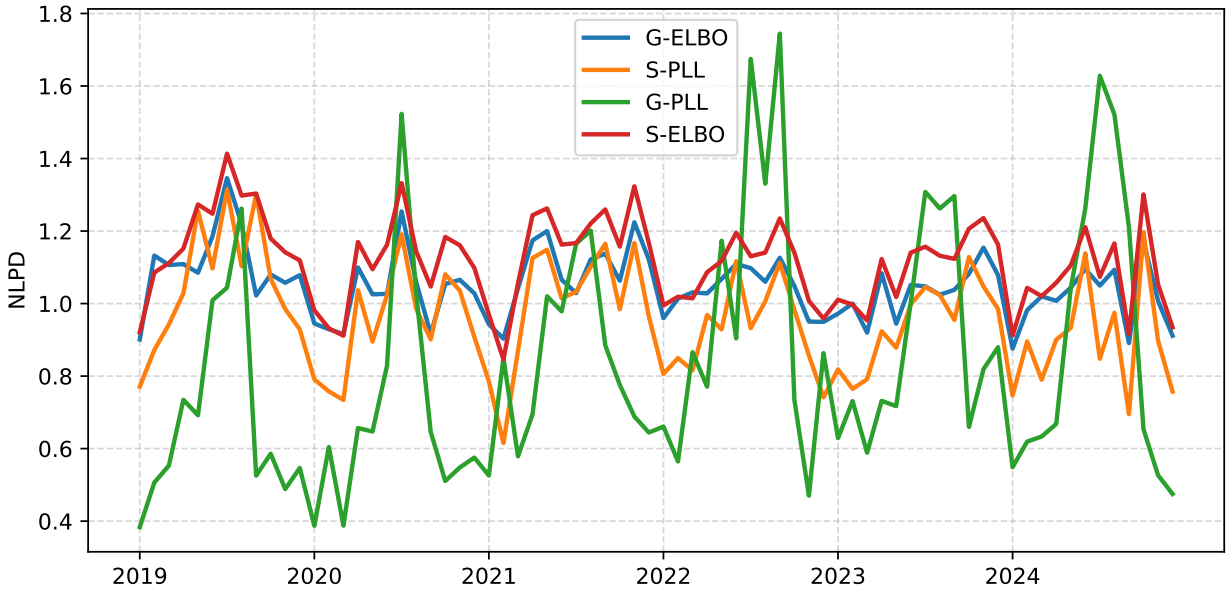

Figure 4: Negative Log Predictive Density by month on the held-out data.

4.1.1. On inducing points

With variational Gaussian processes, the number of inducing points required to generalize the model dictates the complexity of the data and the amount of time it takes to train. On the held-out data, we see in Table 2 that the mean absolute error increases little beyond 1000 inducing points, and that we actually increase the quantile calibration error as we use more inducing points. Furthermore, as we move beyond 2000 inducing points, training time increases significantly. Because we fit the model on each month separately, we set up an embarrassingly parallel training pipeline (i.e., we can train on all months at once, using separate GPUs in a high performance computing environment). With 500 inducing points, we can completely train on each month in 1.5 minutes. With 1000 inducing points, it takes about 2 minutes, with 2000 and 4000 taking up to 10 minutes and an hour, respectively.

IP	MAE	MSE	NLPD	QC50	QC75	QC95	Train Time (min)
500	0.60	0.68	1.00	**2.05**	**1.93**	**0.77**	**1.5**
1000	0.57	0.63	**0.96**	2.90	2.92	1.16	2
2000	**0.56**	0.60	**0.96**	4.42	4.95	2.47	10
4000	**0.56**	**0.59**	0.98	6.25	7.26	4.17	60

Table 2: S-PLL performance metrics on the held-out data with a varying number of inducing points.

4.1.2. Heteroscedastic uncertainty quantification

Given the similarity in 95% quantile calibration error between the robust model and the baseline SVGP, we further make the case for the more robust procedure by looking at the confidence intervals over time. In Figure 5a, we show the estimated prediction on a held-out sensor as a motivating example. While the mean predictions are nearly identical, we see that the more robust S-PLL model has reduced uncertainty at night, as shown in Figure 5b. Because this is when sensors tend to be less prone to bias due to solar radiation, we should prefer this model over the baseline model.

4.1.3. Is the predictive uncertainty informative?

As previously shown in Figure 5a, the S-PLL model adapts its predictive uncertainty (i.e., the estimated standard deviation) temporally, but it also adapts spatially, too. To

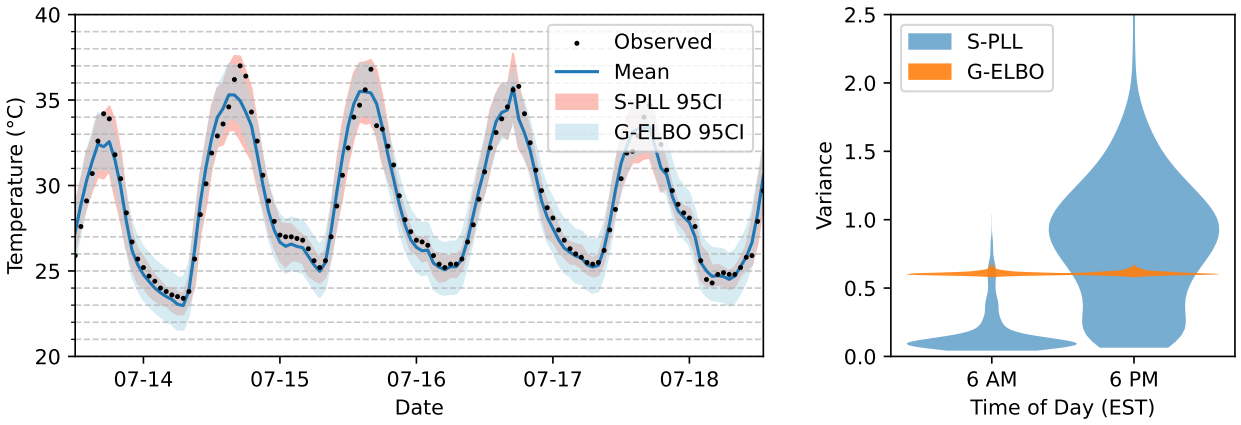

(a) (b)

Figure 5: (a) Comparison of S-PLL with G-ELBO on a held-out location. Only the S-PLL mean is shown for clarity, as the G-ELBO mean estimates are nearly identical. (b) The distribution of the predictive variances for the S-PLL model at 6am versus 6pm local time.

show why this is useful, we visualize the relationship between held-out station residuals and the predictive uncertainty in Figure 6. For a majority of estimates on held-out sensors, the residual error is low and the predictive uncertainty is low, too. However, when the residuals are higher, the model indicates that poor performance is likely. For example, if we analyze the worst performing prediction, where the model overpredicts temperature by 6°C, we find that there were convective thunderstorms in the area around the sensor, causing highly localized cooling. Because the training data varied significantly in this area during this time, the model learned that estimates would be highly uncertain close to these sensors.

To address this problem directly, we would need a high resolution precipitation dataset for convective thunderstorms, which we could insert into the mean function or as an additional term in the covariance function. In the absence of such data, the predictive uncertainty tells us when unobserved conditions make prediction difficult. However, given the low mean standard deviation of our model, such conditions are rare in the area studied.

This figure suggests that the model is well-calibrated, given the relationship between the residual intervals and the predictive uncertainty. For additional insight into model calibration, we compare and contrast the S-PLL model with the G-ELBO model in Appendix B. We show that the S-PLL model is significantly better calibrated than the G-ELBO model

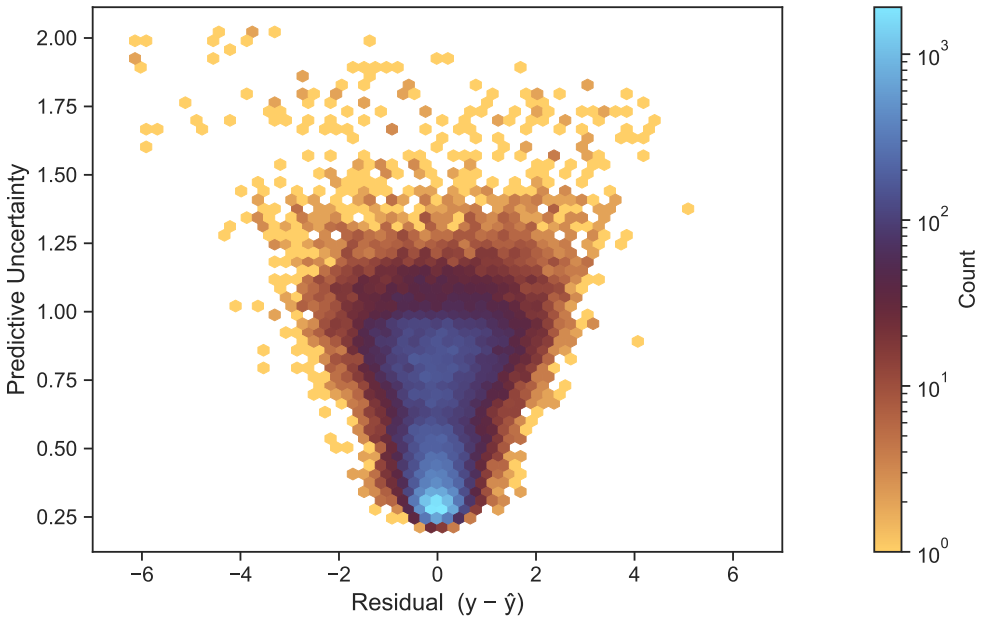

Figure 6: A hexagon plot of the residuals by the predictive uncertainty for July 2024 predictions on held-out sensors, where the predictive uncertainty is the standard deviation estimated by the model. Colors indicate density, with a majority of the observations having small residuals and low uncertainty. $\hat{y}$ denotes the mean estimate from the model.

both overall, and when limiting the comparison to morning or afternoon observations. However, despite the improvement over the baseline method, the S-PLL model remains slightly underconfident at 6am (the predictive intervals are too wide), and slightly overconfident at 6pm (intervals are too narrow).

4.2. Characteristics of urban heat in Durham

In the following sections, we highlight how the S-PLL model allows us to gain new insight into the urban heat island effect in Durham County, North Carolina. We do this by visualizing: (1) the impact on cooling and heating demand; (2) the representative modes of urban heat as a function of larger-scale conditions; and (3) the spatial variability of hours above 35°C and the number of tropical nights.

4.2.1. Heating and Cooling

In Figure 7, we observe the asymmetry between cooling degree days and heating degree days in 2024. Cooling degree days (CDDs) appear more strongly impacted by advection than

heating degree days (HDDs), as areas more frequently downwind (i.e., to the east/northeast of downtown) require more cooling but do not have a commensurate decrease in heating in the wintertime. If we assume that the energy burden of 1 CDD $\approx$ 1 HDD, then we see that the consequence of this asymmetry is that the greatest energy burden falls in the areas around the most urbanized regions. We provide this example to demonstrate the utility of our method: we capture hourly temperature at any point in the year, which enables us to calculate both HDDs and CDDs. Moreover, because our method is not regression based, but rather focused on spatial autocorrelation, we are able to implicitly capture the effect of advection.

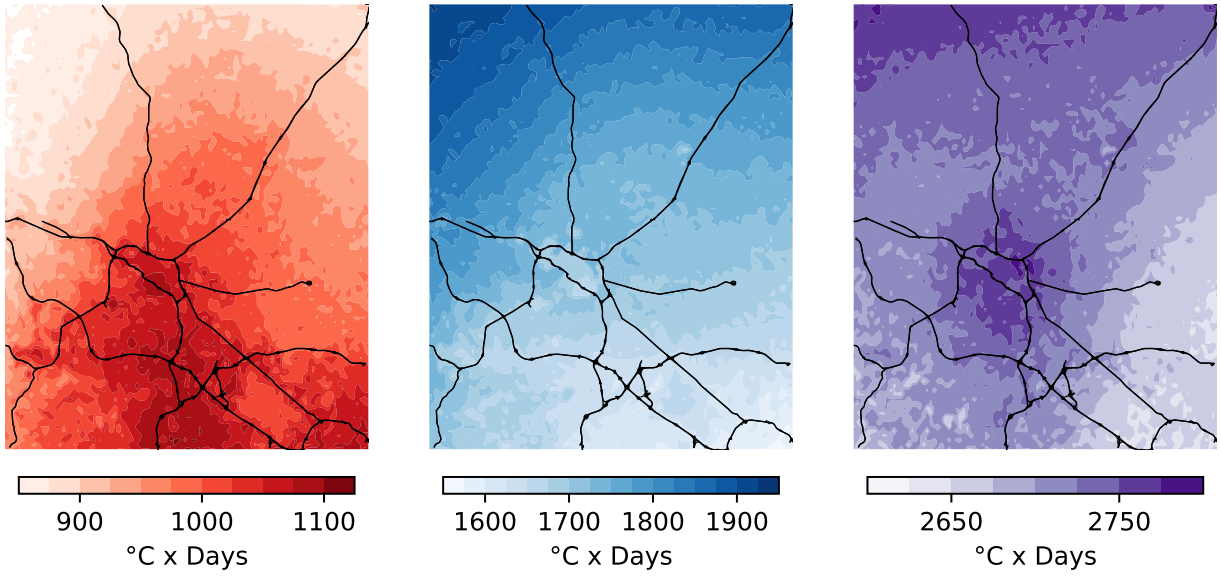

(a) Cooling Degree Days (b) Heating Degree Days (c) CDD + HDD

Figure 7: Cooling and heating degree days for 2024. We can see that the urbanization reflected in PC1 is captured in both cooling and heating degree days, but the effect of urbanization is much more pronounced on CDD.

4.2.2. UHI modes

We further demonstrate the ability of our method to highlight the varying spatial distributions (or modes) of urban heat. To do this, we collect the temperature estimates from June to August and, using ERA5-Land data, perform the KMeans algorithm on the hourly u10 and v10 wind components for each day. Together, these components give us the direction

and magnitude of wind over the course of the day, giving us a feature set of dimension 48 (24 hours of u10/v10 components). This allows us to group similar days based on their diurnal wind profiles. Using 8 clusters, we then classify each day according to its label assignment and cluster centers. For example, if the cluster center has a small u10 component, and a highly positive v10 component, we would classify those days as southerly, in that winds from the south dominate. We compare and contrast a few of these clusters, below.

The most common cluster is cluster 0, representing a dominating southerly wind, experienced on approximately 20% of days. We contrast this cluster with cluster 3, representing dominating northerly winds, on 13% of days. For each set of observations in a cluster, we then subtract the hourly minimum to obtain a measure of the urban heat island intensity (UHII). This allows us to average over the cluster to understand the spatial trend under each cluster's conditions. We visualize clusters 0 and 3, as well as their difference in Figure 8. For clarity, we focus on 7pm EST averages, when the UHII tends to be greater. When the southerly wind dominates, we see warmer temperatures to the northeast. With the northerly, we see higher relative southern temperatures. These highlight the meaningful role that advection plays in prolonging heat exposure downwind of urban areas.

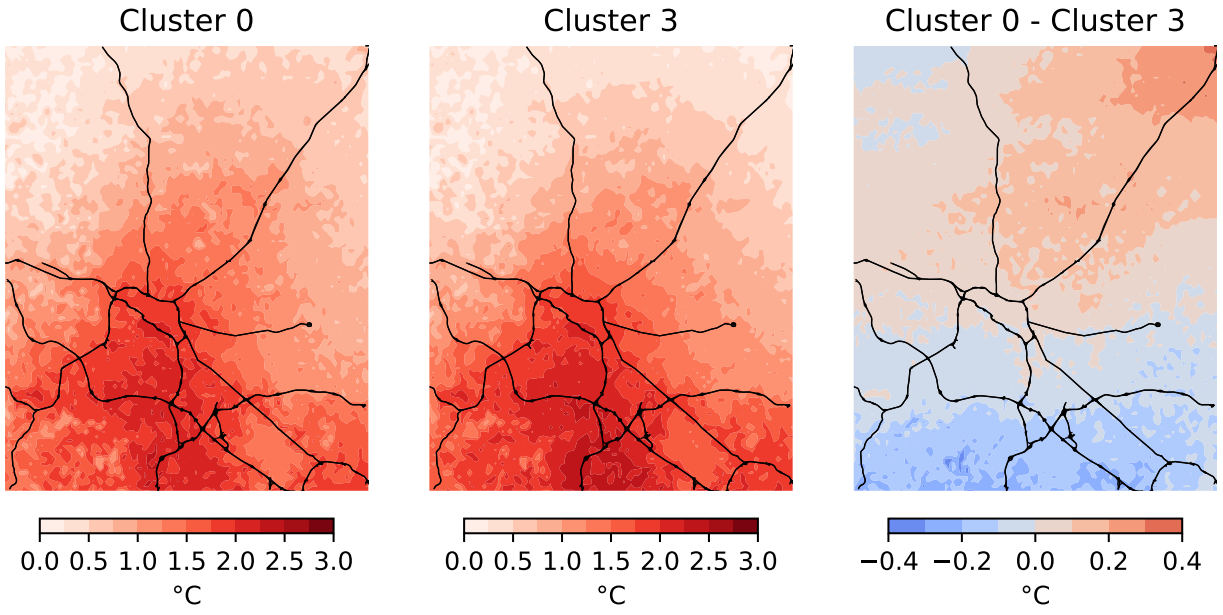

(a) Southerly (b) Northerly (c) Difference

Figure 8: The effect of southerly versus northerly winds

Perhaps the most dramatic difference is seen when comparing clusters 4 and 5 in Figure 9, which correspond to days with dominating southeasterly and southwesterly winds, respectively. In this case, we look at the difference in the urban heat island signature at 9pm EST. We see that the southeasterly dominating wind results in generally a stronger urban effect. We hypothesize this is likely because Raleigh is directly to the Southeast of Durham. As a result, advection is carrying warmer air from Raleigh into Durham, resulting in reduced cooling efficiency in the evening. Cluster 4 represents about 13% of the summer

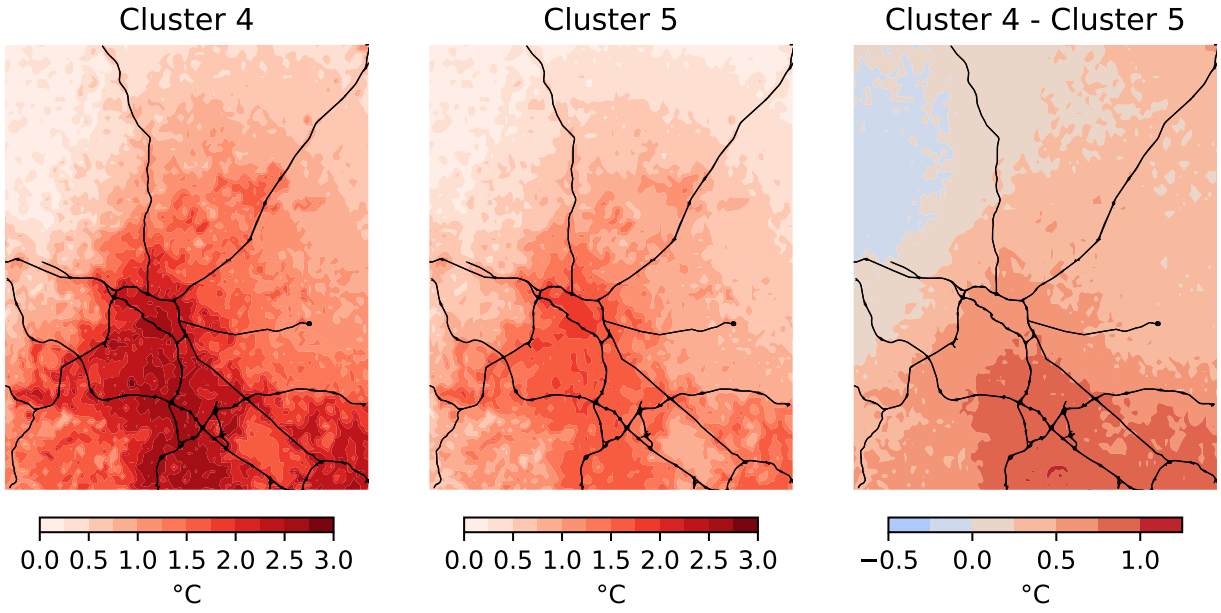

(a) Southeasterly (b) Southwesterly (c) Difference

Figure 9: The effect of southeasterly versus southwesterly winds.

days, whereas cluster 5 represents about 17% of the days.

4.3. Hours above threshold

Lastly, we use our model to evaluate the number of hours above 35°C (95°F), shown in Figure 10. To generate these estimates, we first generate 100 samples from our model over June, July, and August. For each sample, we calculate the number of hours exceeding 35°C. This gives us 100 estimates over the whole map for the number of hours exceeding 35°C, from which we calculate the mean (10a) and standard deviation (10b). We visualize the spatial density of these estimates in Figure 10c, which suggests that urban areas experience as

much as twice as many hours exceeding 95°C. The mean figure highlights the high disparity in extreme heat exposure in particular. However, the standard deviation estimate produced here also provides useful information on spatial variability. If we believe these sensors are still unreliable during the hottest parts of the day, then the standard deviation simply reflects that sensor noise. However, if we assume that the quality control procedure is effective, then this standard deviation highlights the potential for local interventions. This is because the standard deviation reflects the local variability not captured by the model, and thus attributed to local conditions.

5. Discussion

These experiments demonstrate that the sparse variational Gaussian process is an efficient and effective method for spatiotemporal modeling of urban heat: very few inducing points are needed, which allows for rapid model training, and using the predictive log likelihood loss function with the Student's t-likelihood, we obtain well calibrated estimates of urban temperature. While other approaches to spatiotemporal modeling suggest that more inducing points are better (e.g., Hamelijnck et al. (2021)), we find that this is not the case with noisy citizen-science data. In fact, the model may benefit from using fewer inducing points, as this reduces the potential for overfitting. Furthermore, while other methods such as Bayesian Hierarchical Models allow for well-calibrated probabilistic spatiotemporal models, they typically require significantly more training time. Even with 2000 inducing points, our method fits each month in less than 10 minutes.

Despite the rapid training time, there may be more opportunities to reduce computational complexity. Namely, we use 1000 inducing points where even fewer may be sufficient. When we tested with 500 inducing points, we noted that model performance only decreased slightly, but we hypothesize this could be due to poor initialization. With fewer inducing points, initialization has a more significant impact on model convergence. 2000 inducing points appears to be more than enough to avoid the initialization problem, but better initialization could allow for better performance with fewer inducing points (see Burt et al. (2020) for an overview of initialization techniques). Furthermore, natural gradient descent (NGD) allows for more rapid convergence and may enable fewer inducing points to obtain better

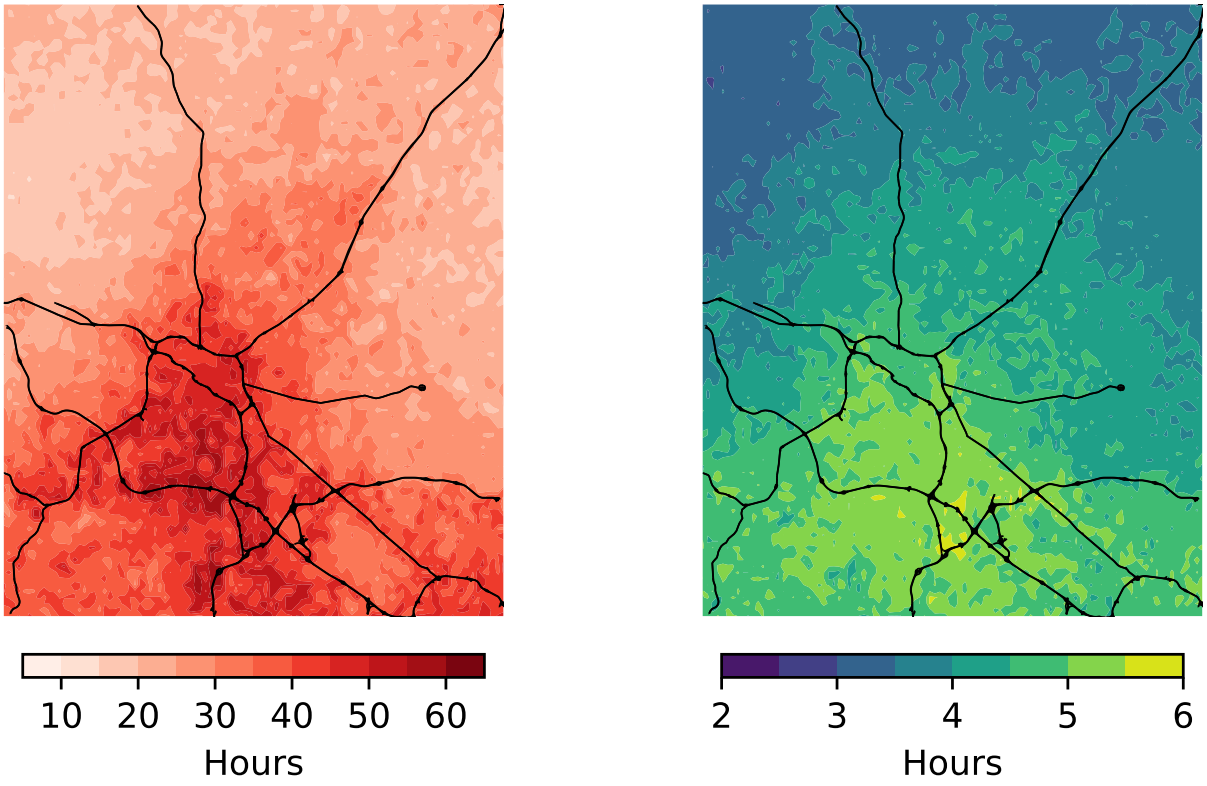

(a) Mean

(b) Standard Deviation

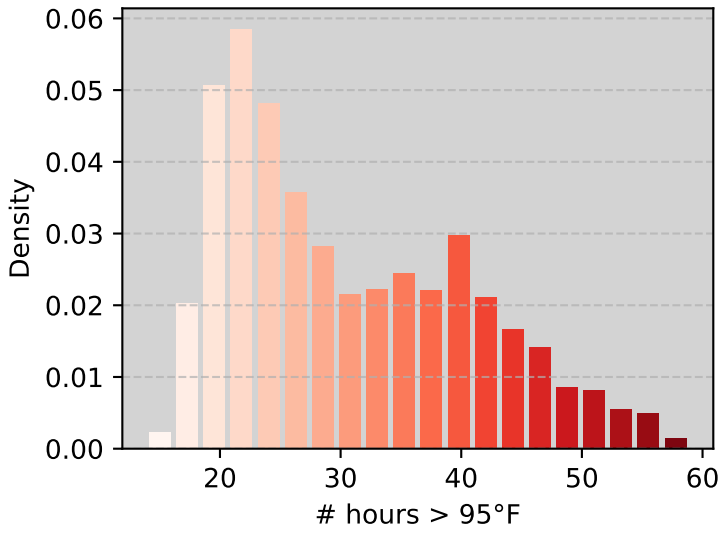

(c) Distribution

Figure 10: The estimated number of hours exceeding 35°C (95°F).

performance, too (Hensman et al., 2013). However, NGD is not always stable with non-Gaussian likelihoods.

We could also improve the covariance definition and feature set. While the SVGP approach allows our model to only focus on learning spatially coherent variations from the ERA5-Land, it is restricted to rather large length-scales. In locations with denser networks, it is likely that additional additive kernel terms could be used to capture smaller length-scale variability. We also restricted our parameterization of urban features to a simple principal component. This approach assumes that similar component values should co-vary, but we could incorporate further terms into our parameterization of urban areas. For example, if we had more information about the shading conditions of the stations, we would assume that unshaded/shaded stations will correlate, too.

Beyond the covariance definition and feature set, we made additional design choices that could be revisited (e.g., using linearly interpolated ERA5-Land, fitting data over months, optimization hyperparameters, etc). While we provided reasons for justifying many of these choices, good performance can be achieved through alternative choices, too. We argue here that a good interpolation method will reduce the impact of preferential sampling, explicitly model heteroscedastic noise, and avoid overfitting. Moreover, we used an optimization procedure that worked well on our data, but other strategies could work, too. Generally, we advise reasonable parameter initialization and sufficient iterations over the training data to ensure good model fit. The GPyTorch implementation of SVGP regression offers flexibility; we hope that our work encourages the urban climate community to consider this approach when interpolating noisy data.

Limitations remain in the interpretation of the predictive variance. This is because the predictive variance is learned from sensors with an unknown noise distribution, so the predictive variance at any point reflects a sum of the sensor noise and the uncertainty due to gaps in observations. In other words, the model tells us the range of noisy values expected if we placed a sensor in a new location. If the sensor noise has mean zero, then our model gives an unbiased estimate of the actual air temperature. However, we do not know if this is the case. Sensors in citizen's backyards may have a slight negative bias and misrepresent the conditions faced by pedestrians, whereas sensors fully in the sun may still have a positive

bias due to radiative heating, even after the quality control steps taken here. Despite this limitation, our model is precise: it can accurately estimate the temperature of held-out sensors. This suggests that the observations come from a stable distribution, a necessary but not sufficient condition to guarantee accurate estimation of true temperature; more ground-truth observations are needed to validate that the quality-controlled observations are unbiased. Nonetheless, this problem reinforces the need for a probabilistic approach. Without quantifying the potential noise in crowdsourced data, it is challenging to confidently answer scientific questions from this source.

Regardless, a significant effect on local temperature is due to siting conditions, which are generally unknown. Sensors are exposed to effects at the micro-, neighborhood-, city-, and meso-scale (Muller et al., 2013). A common critique of urban heat studies is on the lack of adequate characterization of what scale is being captured by the network in question (Stewart, 2011). Our model provides a principled solution to this problem through the covariance function, which captures mean estimates at the learned length-scales, and variance estimates to explain variability below this length-scale (i.e., capturing micro-scale effects). In this respect, we can determine when the city-scale urban heat island dominates (low spatial variance), versus local-scale effects (high spatial variance).

While the S-PLL model represents an improvement over homoscedastic interpolation techniques for quantifying uncertainty, we still find that the predictive uncertainty is not perfect when looking at quantile coverage error for specific times of day (see Appendix B). This suggests that the noise process is more complex than the approach taken here, and could be further improved by modeling noise as a function of other meteorological conditions. For example, the noise itself could be modeled as a separate Gaussian process, entirely separate from the random field, trained on both time of day and shortwave radiation. Such an approach would increase computational cost, but may improve predictive uncertainty still. Nonetheless, our results suggest that modeling heteroscedasticity should be prioritized when interpolating crowdsourced data, and that the S-PLL approach offers a computationally efficient means to do so.

Lastly, the estimates produced from this model present several new questions worth exploring. First, to our knowledge, the asymmetry in CDD/HDD due to urban heating has

not been shown in prior works. Future work should replicate our work on additional cities to determine if this is a more generalizable pattern. If so, we would be able to highlight the additional energy burden for communities on the urban periphery, where increased cooling demand is not offset by reduced heating demand. Second, this method is readily extended to other crowdsourced meteorological variables. Weather Underground provides information on dew point, precipitation, solar radiation, wind speed, and pressure. While these variables may be noisier, we expect that our modeling approach is robust enough to provide estimates of how they vary in an urban setting. We demonstrate the application to dew point in Appendix C, which highlights the seasonal variability in the dry island effect. Precipitation and solar radiation could be treated through the same model, but likely require a log-transform for use in a Gaussian process. An additional application of our uncertainty-aware model may be in the field of data assimilation with high-resolution urban canopy models. In data assimilation, we often contend with both observation uncertainty and model uncertainty. A barrier to the use of crowdsourced data in this application is on the quantification of uncertainty due to siting. Our method provides an effective means to explain this variance, which allows for crowdsourced data to be fed into data assimilation pipelines.

Appendix A. The evidence lower bound

Here, we offer a brief derivation of the evidence lower bound in the context of sparse variational Gaussian processes, with the goal of providing an accessible introduction to this topic assuming minimal background in statistics or machine learning. This derivation combines components from other work before it (e.g., Hensman et al. (2015); Blei et al. (2017)). However, many of the existing derivations take the result from variational inference, first, then show how it applies to sparse variational Gaussian processes. Here, we take the opposite approach: starting from the marginal log-likelihood, we demonstrate how the evidence lower bound is achieved through implementing a sparse approximation. We have not seen such a derivation taken in prior work.

First, recall that our goal is to obtain an approximation of the marginal log likelihood:

$$\log p(\mathbf{y} \mid X) = \log \int p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid X) d\mathbf{f} \tag{A.1}$$

Given the variational approximation $p(\mathbf{f}) \approx q(\mathbf{f})$, we start by introducing $q(\mathbf{f})$ into the objective (Equation A.2a), in which we place $q(\mathbf{f})$ in both the numerator and denominator. Then, by definition, we write this equation as an expectation (Equation A.2b).

$$\log \int p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X}) d\mathbf{f} = \log \int \frac{q(\mathbf{f}) p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X})}{q(\mathbf{f})} d\mathbf{f} \tag{A.2a}$$

$$= \log \mathbb{E}_{q(\mathbf{f})} \left[\frac{p(\mathbf{y} \mid \mathbf{f}) p(\mathbf{f} \mid \mathbf{X})}{q(\mathbf{f})} \right] \tag{A.2b}$$

By Jensen's inequality, which states that for a concave function g, $g(\mathbb{E}[\cdot]) \geq \mathbb{E}[g(\cdot)]$, we obtain Equation A.3a. This step results in the "lower bound" component, causing the resulting objective to be an approximation of the likelihood. When we re-arrange the equation to put similar terms together, we see that this step decouples the model estimates from the random field.

$$\log p(\mathbf{y} \mid X) \geq \mathbb{E}_{q(\mathbf{f})} \left[\log p(\mathbf{y} \mid \mathbf{f}) + \log p(\mathbf{f} \mid X) - \log q(\mathbf{f}) \right] \tag{A.3a}$$

$$= \mathbb{E}_{q(\mathbf{f})} \left[\log p(\mathbf{y} \mid \mathbf{f}) \right] - \mathbb{E}_{q(\mathbf{f})} \left[\log \frac{q(\mathbf{f})}{p(\mathbf{f})} \right] \tag{A.3b}$$

By definition, the second term is the Kullback-Leibler (KL) divergence, a metric used to calculate the difference between two distributions. This is typically written as:

$$\mathbb{E}_{q(\mathbf{f})} \left[\log \frac{q(\mathbf{f})}{p(\mathbf{f})} \right] = \mathrm{KL}(q(\mathbf{f}) || p(\mathbf{f})) \tag{A.4}$$

It's worth pointing out that the above definition is known as the *reverse* KL. In contrast, the *forward* KL swaps the positions of the true distribution and the approximating distribution (i.e., $\mathrm{KL}(p(\mathbf{f}) || q(\mathbf{f}))$.

The problem with this objective is that we would prefer to work with $p(\mathbf{u})$ and $q(\mathbf{u})$, because this is ultimately our variational approximation, and is a simpler calculation. In

fact, we can show that these two are equivalent algebraically:

$$\mathbb{E}_{q}(\mathbf{f})\left[\log \frac{q(\mathbf{f})}{p(\mathbf{f})}\right] = \int q(\mathbf{f}) \log \frac{q(\mathbf{f})}{p(\mathbf{f})} d\mathbf{f} \tag{A.5a}$$

$$= \int \int q(\mathbf{f}, \mathbf{u}) d\mathbf{u} \log \frac{\int p(\mathbf{f} \mid \mathbf{u}) q(\mathbf{u}) d\mathbf{u}}{\int p(\mathbf{f} \mid \mathbf{u}) p(\mathbf{u}) d\mathbf{u}} d\mathbf{f} \tag{A.5b}$$

$$= \int \int q(\mathbf{f}, \mathbf{u}) d\mathbf{f} \log \frac{q(\mathbf{u})}{p(\mathbf{u})} d\mathbf{u} \tag{A.5c}$$

$$= \int q(\mathbf{u}) \log \frac{q(\mathbf{u})}{p(\mathbf{u})} d\mathbf{u} \tag{A.5d}$$

$$= \mathbb{E}_{q(\mathbf{u})}\left[\log \frac{q(\mathbf{u})}{p(\mathbf{u})}\right] = \mathrm{KL}(q(\mathbf{u}) || p(\mathbf{u})) \tag{A.5e}$$

Intuitively, this makes sense, given that the only difference between $p(\mathbf{f})$ and $q(\mathbf{f})$ is $q(\mathbf{u})$, which we see in Equation A.5b. Now, if we insert this re-parameterized term into the complete equation, we have:

$$\mathcal{L}_{ELBO} = \mathbb{E}_{q(\mathbf{f})}\left[\log p(\mathbf{y} \mid \mathbf{f})\right] - \mathrm{KL}(q(\mathbf{u}) || p(\mathbf{u})) \tag{A.6}$$

Finally, by definition, $\log p(\mathbf{y} \mid \mathbf{f}) = \sum_{i=1}^{N} \log p(y_i \mid f_i)$; $\mathbf{y}$ consists of conditionally independent observations, given $\mathbf{f}$. Because of this, we can re-write the evidence lower bound as:

$$\mathcal{L}_{ELBO} = \sum_{i=1}^{N} \mathbb{E}_{q(f_i)}\left[\log p(y_i \mid f_i)\right] - \mathrm{KL}(q(\mathbf{u}) || p(\mathbf{u})) \tag{A.7}$$

As mentioned in the main text, this property allows for "mini-batching," or estimating the loss function over smaller subsets of the data. In contrast with the marginal log-likelihood for exact Gaussian processes, in which we perform gradient descent, "mini-batching" allows for *stochastic* gradient descent, which is computationally cheaper per iteration, allowing for quicker convergence to performance within an acceptable error tolerance of gradient descent solution (Bottou, 2010).

Appendix B. Comparing S-PLL and G-ELBO calibration

We compare and contrast the S-PLL calibration with the G-ELBO calibration in Figure B.11 using a probability integral transform (PIT) for the held-out July 2024 observations. This figure more generally demonstrates the phenomenon seen in Figure 5a in the main text, wherein the G-ELBO model has predictive intervals that are two wide at night, and too narrow during the day. If a model is perfectly-calibrated, the cumulative distribution function (CDF) of observations with respect to the model will have a uniform distribution. For the G-ELBO model, the hump in the center corresponds to the night, when the model is underconfident, and the higher counts at the tails correspond to the afternoon, when the model is over-confident. In contrast, the S-PLL CDF is approximately uniform, but has slightly less mass at 0, suggesting that the model tends to underpredict at the held-out stations shown here. In other words, the held-out observations include fewer instances of lower extreme values than the model expects.

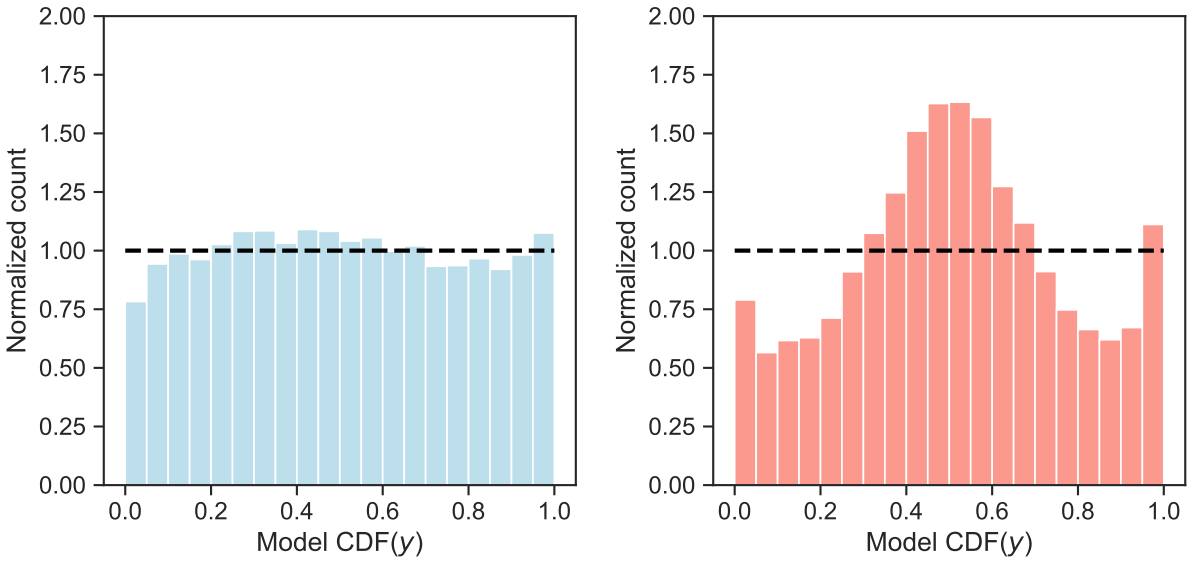

(a) S-PLL

(b) G-ELBO

Figure B.11: Probability integral transform (PIT) diagrams for the S-PLL and G-ELBO models for the July 2024 observations. The dashed black line represents the values for a perfectly calibrated model.

We analyze the calibration performance in the morning and late afternoon in Figure B.12. The hump for morning observations suggests underconfidence, or that the estimated variance

is too large. We see that the S-PLL model is slightly underconfident for morning observations, indicating that the morning variance may be too conservative. However, the estimated variance is significantly better than the G-ELBO model. We see the opposite problem with the 6pm estimates: the model is overconfident, but again, the S-PLL model less so than the G-ELBO model. Additionally, the slight right skewness suggests that among the held-out observations, the model tends to slightly underpredict in the afternoon. These results correspond to the 95% predictive interval for the S-PLL model covering the held-out values 97% and 90% of the time at 6am and 6pm, respectively, in comparison with the G-ELBO model covering these values 99% and 84% of the time.

Overall, the PIT diagrams suggest that modeling heteroscedastic noise improves calibration when interpolating noisy crowdsourced meteorological data, and probabilistic models that model this heteroscedasticity should be favored over standard probabilistic interpolation techniques (e.g., ordinary kriging, the most similar method to the G-ELBO comparison used here).

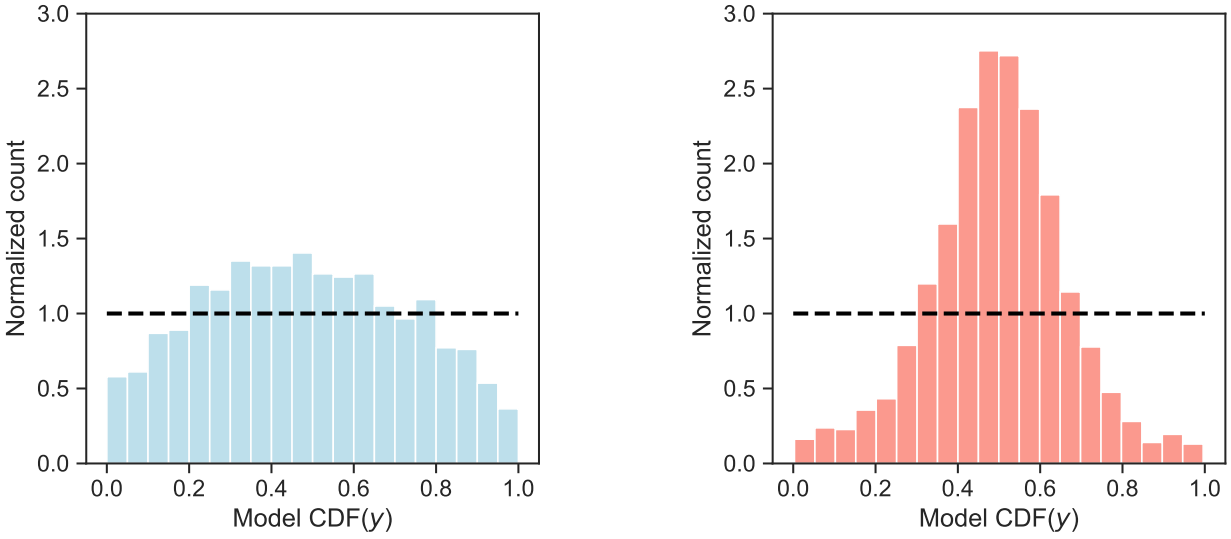

(a) S-PLL, 6am

(b) G-ELBO, 6am

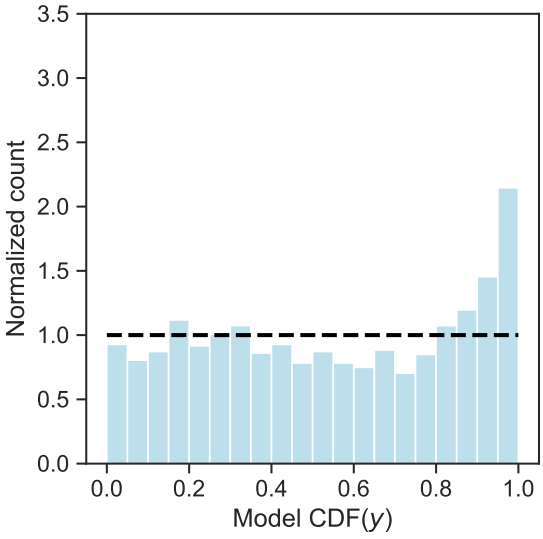

(c) S-PLL, 6pm

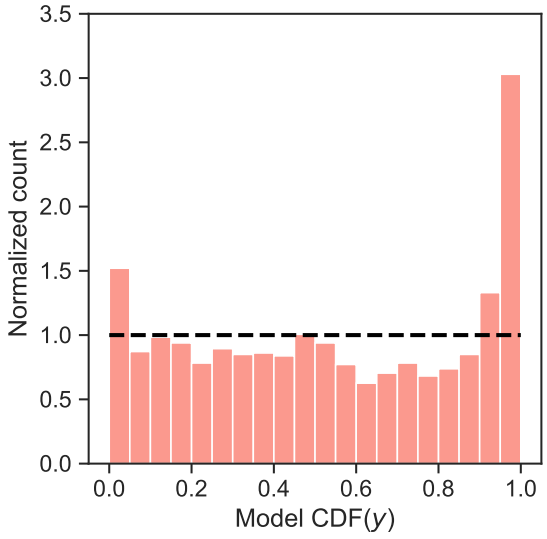

(d) G-ELBO,6pm

Figure B.12: PIT diagrams for 6am and 6pm (local time) residuals. Dashed black line represents values expected for a perfectly calibrated model.

Appendix C. Dew Point

Additional experiments were performed to understand our model's ability to calculate the spatiotemporal variability of dew point in Durham County. We preprocess the data according to the same process as with temperature, with two additions. We removed observations in which dew point exceeded the station's temperature, and we increased the lower α value used for removing outliers to 0.05. The dew point data appeared to contain significantly more outliers than the temperature data, so we added this additional filter to make our model more conservative. To our knowledge, there is no attempt in previous literature to validate crowdsourced dew point observations, and developing a specific filter for dew point is beyond the scope of our current work.

The only change to the model was then to update the covariates to ingest the ERA5-Land 2-meter dew point estimate rather than the 2-meter air temperature, and to initialize the linear term in the model corresponding to the principal component to -0.1, as we expect more urbanized regions to have a lower dew point (i.e., the urban dry island effect). We then tested our the *S-PLL* model on this data, using again 20% of the stations as a validation set. Our results, in Figure C.13, show that there is not the same diurnal pattern in performance as shown with temperature. Instead, performance generally improves over time, which we attribute to having more sensors.

We additionally show average evening dew points by season in Figure C.14. We see that the urban dry island is most salient in July, almost non-existent in January, and moderately visible in April and October.

The motivation behind modeling the dew point is that we may further improve our estimates of heat stress by factoring in moisture into key urban heat stress metrics (e.g., the heat index or wet-bulb globe temperature (WBGT).

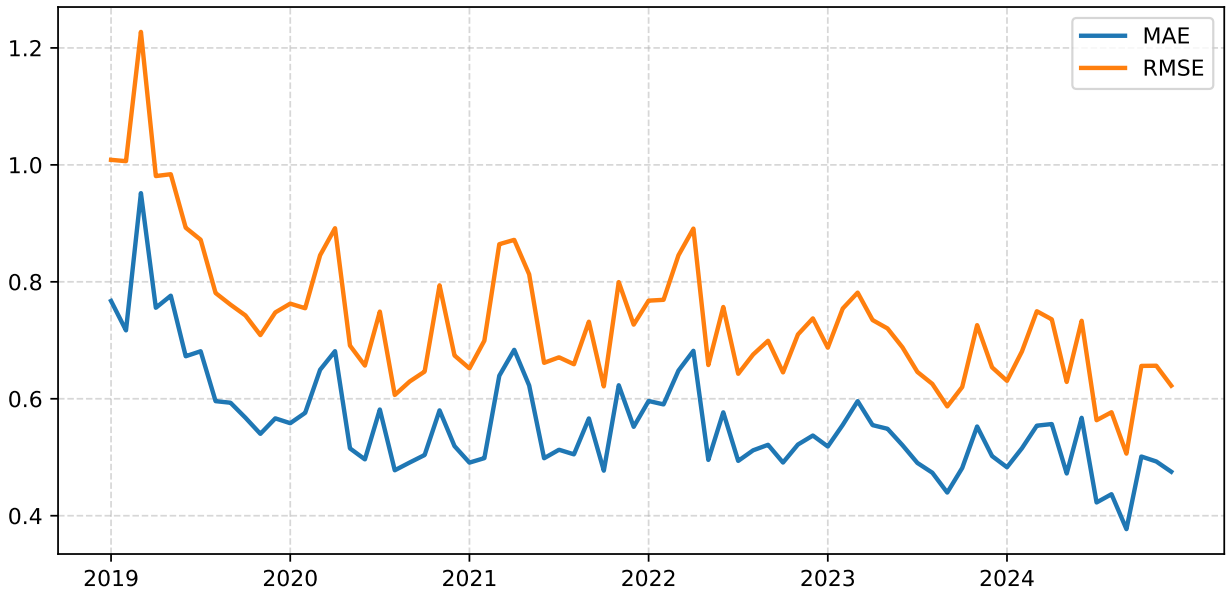

Figure C.13: Dew point mean absolute error and root mean squared error, monthly.

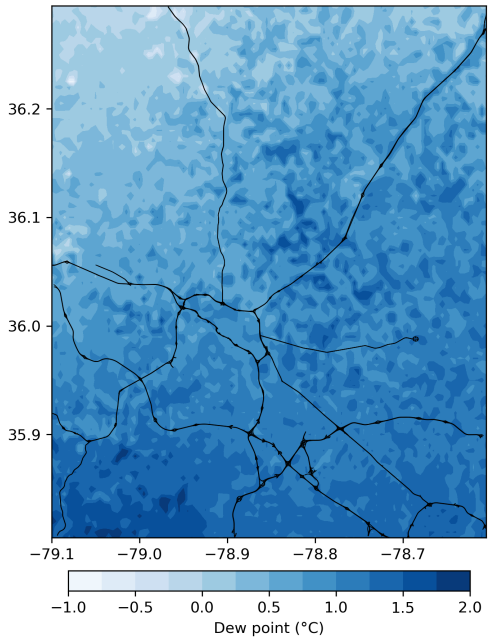

(a) January

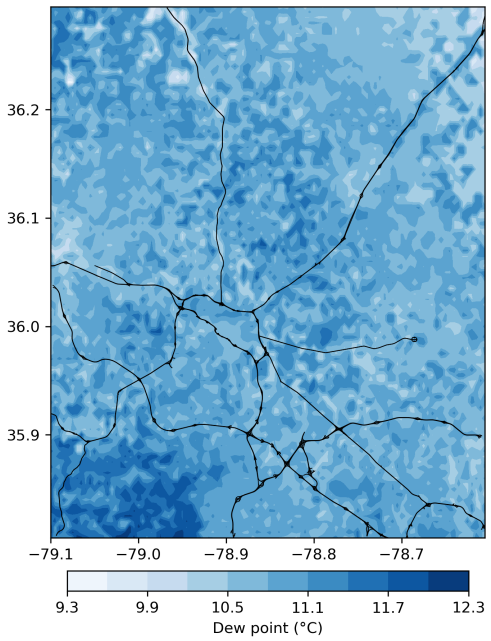

(b) April

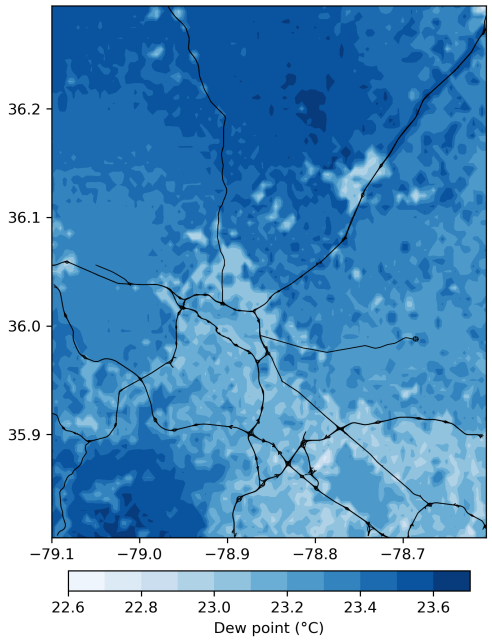

(c) July

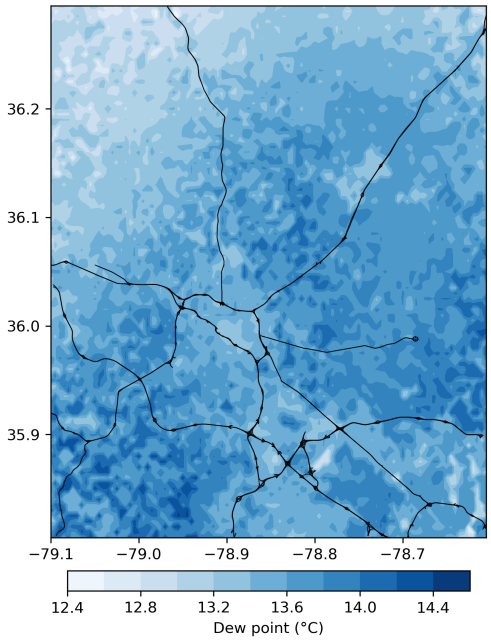

(d) October

Figure C.14: Average 8pm (local time) Dew Point. Evening is used, as this is when the urban dry island effect is most active.

Appendix D. Gaussian processes for prediction

Given the training dataset ($\mathbf{X}$, $\mathbf{y}$), the learned random field $\mathbf{f}$ (with mean function $m(\cdot)$) and covariance function $k(\cdot,\cdot)$), we produce estimates of the underlying random field $\mathbf{f}_*$ by conditioning on the multivariate normal distribution. To see this, we first consider the joint distribution of $\mathbf{y}$ and $\mathbf{f}_*$:

$$p\left(\begin{bmatrix}\mathbf{y}\\ \mathbf{f}_*\end{bmatrix} \,\middle|\, \begin{bmatrix}\mathbf{X}\\ \mathbf{X}_*\end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix}m(\mathbf{X})\\ m(\mathbf{X}_*)\end{bmatrix}, \begin{bmatrix}k(\mathbf{X},\mathbf{X}) + \varepsilon^2\mathbf{I} & k(\mathbf{X},\mathbf{X}_*)\\ k(\mathbf{X}_*,\mathbf{X}) & k(\mathbf{X}_*,\mathbf{X}_*)\end{bmatrix}\right) \tag{D.1}$$

The covariance matrix represents a block matrix defined with respect to the covariance matrices for the training data and the test data. From this joint distribution, we condition on the observations to produce estimates at new locations. A property of the multivariate normal distribution is that the conditional distribution has the following analytical result:

$$p(\mathbf{f}_* \mid \mathbf{y}, \mathbf{X}, \mathbf{X}_*) = \mathcal{N}(\boldsymbol{\mu}_*, \mathbf{K}_*) \tag{D.2}$$

$$\boldsymbol{\mu}_* = m(\mathbf{X}_*) + k(\mathbf{X}_*,\mathbf{X})\Big[k(\mathbf{X},\mathbf{X}) + \varepsilon^2\mathbf{I}\Big]^{-1}\Big(\mathbf{y} - m(\mathbf{X})\Big) \tag{D.3}$$

$$\mathbf{K}_* = k(\mathbf{X}_*,\mathbf{X}_*) - k(\mathbf{X}_*,\mathbf{X})\Big[k(\mathbf{X},\mathbf{X}) + \varepsilon^2\mathbf{I}\Big]^{-1}k(\mathbf{X},\mathbf{X}_*) \tag{D.4}$$

Gaussian processes produce probabilistic estimates through this property. We see that prediction requires inversion of a potentially large matrix, re-introducing a computational burden mentioned in the main text. The sparse approximation overcomes this computational limitation by conditioning on the inducing points.

For a longer discussion of this conditional distribution, we direct the reader to Rasmussen and Williams (2006), Chapter 2.

References

Agrawal, R., Lizana, J., Vargas, P.S.N., Núñez-Peiró, M., 2025. Spatio-Temporal Distribution and Data Quality of Citizen Weather Stations: A Comparative Study of Netatmo and Wunderground in Paris. doi:10.5194/icuc12-59.

de Baar, J.H., van Der Schrier, G., 2025. Come as you are: Reconsidering the need for complex quality control when gridding crowd-sourced weather data. Quarterly Journal of the Royal Meteorological Society 151, e4890. doi:10.1002/qj.4890.

Bassett, R., Cai, X., Chapman, L., Heaviside, C., Thornes, J.E., Muller, C.L., Young, D.T., Warren, E.L., 2016. Observations of urban heat island advection from a high-density monitoring network. Quarterly Journal of the Royal Meteorological Society 142, 2434–2441. doi:10.1002/qj.2836.

van Beekvelt, D., Garcia-Marti, I., de Baar, J., 2024. Towards high-resolution gridded climatology stemming from the combination of official and crowdsourced weather observations using multi-fidelity methods. PLOS Climate 3, e0000216. doi:10.1371/journal.pclm.0000216.

Belcher, S.E., Coceal, O., Goulart, E.V., Rudd, A.C., Robins, A.G., 2015. Processes controlling atmospheric dispersion through city centres. Journal of Fluid Mechanics 763, 51–81. doi:10.1017/jfm.2014.661.

Blei, D.M., Kucukelbir, A., McAuliffe, J.D., 2017. Variational inference: A review for statisticians. Journal of the American Statistical Association 112, 859–877. URL: http://dx.doi.org/10.1080/01621459.2017.1285773, doi:10.1080/01621459.2017.1285773.

Bottou, L., 2010. Large-scale machine learning with stochastic gradient descent, in: Proceedings of COMPSTAT'2010: 19th International Conference on Computational StatisticsParis France, August 22-27, 2010 Keynote, Invited and Contributed Papers, Springer. pp. 177–186.

Brousse, O., Simpson, C., Kenway, O., Martilli, A., Krayenhoff, E.S., Zonato, A., Heaviside, C., 2023. Spatially explicit correction of simulated urban air temperatures using crowdsourced data. Journal of applied meteorology and climatology 62, 1539–1572. doi:10.1175/JAMC-D-22-0142.1.

Brousse, O., Simpson, C., Walker, N., Fenner, D., Meier, F., Taylor, J., Heaviside, C., 2022. Evidence of horizontal urban heat advection in London using six years of data from a citizen weather station network. Environmental Research Letters 17, 044041. doi:10.1088/1748-9326/ac5c0f.

Brousse, O., Simpson, C.H., Poorthuis, A., Heaviside, C., 2024. Unequal distributions of crowdsourced weather data in England and Wales. Nature Communications 15, 4828. doi:10.1038/s41467-024-49276-z.

Burt, D.R., Rasmussen, C.E., van der Wilk, M., 2020. Convergence of sparse variational inference in gaussian processes regression. Journal of Machine Learning Research 21, 1–63. URL: http://jmlr.org/papers/v21/19-1015.html.

Båserud, L., Lussana, C., Nipen, T.N., Seierstad, I.A., Oram, L., Aspelien, T., 2020. TITAN automatic spatial quality control of meteorological in-situ observations, in: Advances in Science and Research, Copernicus GmbH. pp. 153–163. doi:10.5194/asr-17-153-2020.

Calhoun, Z.D., Black, M.S., Bergin, M., Carlson, D., 2024. Refining citizen climate science: Addressing preferential sampling for improved estimates of urban heat. Environmental Science & Technology Letters 11, 845–850. doi:10.1021/acs.estlett.4c00296.

Carrión, D., Arfer, K.B., Rush, J., Dorman, M., Rowland, S.T., Kioumourtzoglou, M.A., Kloog, I., Just, A.C., 2021. A 1-km hourly air-temperature model for 13 northeastern U.S. states using remotely sensed and ground-based measurements. Environmental Research 200, 111477. doi:10.1016/j.envres.2021.111477.

Castro Medina, D., family=Guerrero Delgado, given=Mcarmen, g.i., Sánchez Ramos, J., Palomo Amores, T., Romero Rodríguez, L., Álvarez Domínguez, S., 2024. Empowering urban climate resilience and adaptation: Crowdsourcing weather citizen stations-enhanced temperature prediction. Sustainable Cities and Society 101, 105208. doi:10.1016/j.scs.2024.105208.

Chen, A.B., Behl, M., Goodall, J.L., 2021. Assessing the trustworthiness of crowdsourced rainfall networks: A reputation system approach. Water Resources Research 57, e2021WR029721. doi:10.1029/2021WR029721.

Chen, G., Shi, Y., Wang, R., Ren, C., Ng, E., Fang, X., Ren, Z., 2022. Integrating weather observations and local-climate-zone-based landscape patterns for regional hourly air tem-

perature mapping using machine learning. Science of The Total Environment 841, 156737. doi:10.1016/j.scitotenv.2022.156737.

Diggle, P.J., Menezes, R., Su, T.l., 2010. Geostatistical inference under preferential sampling. Journal of the Royal Statistical Society Series C: Applied Statistics 59, 191–232. doi:10.1111/j.1467-9876.2009.00701.x.

Fenner, D., Bechtel, B., Demuzere, M., Kittner, J., Meier, F., 2021. CrowdQC+—A Quality-Control for Crowdsourced Air-Temperature Observations Enabling World-Wide Urban Climate Applications. Frontiers in Environmental Science 9. doi:10.3389/fenvs.2021.720747.

Flückiger, B., Kloog, I., Ragettli, M.S., Eeftens, M., Röösli, M., de Hoogh, K., 2022. Modelling daily air temperature at a fine spatial resolution dealing with challenging meteorological phenomena and topography in Switzerland. International Journal of Climatology 42, 6413–6428. doi:10.1002/joc.7597.

Gardner, J.R., Pleiss, G., Bindel, D., Weinberger, K.Q., Wilson, A.G., 2018. Gpytorch: blackbox matrix-matrix gaussian process inference with gpu acceleration, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA. p. 7587–7597.

Hamelijnck, O., Wilkinson, W., Loppi, N., Solin, A., Damoulas, T., 2021. Spatio-Temporal Variational Gaussian Processes, in: Advances in Neural Information Processing Systems, Curran Associates, Inc.. pp. 23621–23633. doi:10.5555/3540261.3542070.

Hensman, J., Fusi, N., Lawrence, N.D., 2013. Gaussian processes for big data, in: Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, AUAI Press, Arlington, Virginia, USA. p. 282–290.

Hensman, J., Matthews, A., Ghahramani, Z., 2015. Scalable Variational Gaussian Process Classification, in: Lebanon, G., Vishwanathan, S.V.N. (Eds.), Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, PMLR, San Diego,

California, USA. pp. 351–360. URL: https://proceedings.mlr.press/v38/hensman15.html.

Hough, I., Just, A.C., Zhou, B., Dorman, M., Lepeule, J., Kloog, I., 2020. A multi-resolution air temperature model for France from MODIS and Landsat thermal data. Environmental Research 183, 109244. doi:10.1016/j.envres.2020.109244.

Hsu, A., Sheriff, G., Chakraborty, T., Manya, D., 2021. Disproportionate exposure to urban heat island intensity across major US cities. Nature Communications 12, 2721. doi:10.1038/s41467-021-22799-5.

Jankowiak, M., Pleiss, G., Gardner, J.R., 2020. Parametric gaussian process regressors, in: Proceedings of the 37th International Conference on Machine Learning, JMLR.org. pp. 4702–4712. doi:10.5555/3524938.3525375.

Kloog, I., Nordio, F., Coull, B.A., Schwartz, J., 2014. Predicting spatiotemporal mean air temperature using MODIS satellite surface temperature measurements across the Northeastern USA. Remote Sensing of Environment 150, 132–139. doi:10.1016/j.rse.2014.04.024.

Küppers, F., Schneider, J., Haselhoff, A., 2023. Parametric and Multivariate Uncertainty Calibration for Regression and Object Detection, in: Karlinsky, L., Michaeli, T., Nishino, K. (Eds.), Computer Vision – ECCV 2022 Workshops, Springer Nature Switzerland. pp. 426–442. doi:10.1007/978-3-031-25072-9_30.

Laaidi, K., Zeghnoun, A., Dousset, B., Bretin, P., Vandentorren, S., Giraudet, E., Beaudeau, P., 2012. The Impact of Heat Islands on Mortality in Paris during the August 2003 Heat Wave. Environmental Health Perspectives 120, 254–259. doi:10.1289/ehp.1103532.

Lee, J., Dessler, A.E., 2024. Improved Surface Urban Heat Impact Assessment Using GOES Satellite Data: A Comparative Study With ERA-5. Geophysical Research Letters 51, e2023GL107364. doi:10.1029/2023GL107364.

Levi, D., Gispan, L., Giladi, N., Fetaya, E., 2022. Evaluating and calibrating uncertainty prediction in regression tasks. Sensors 22. doi:10.3390/s22155540.

Li, D., Wang, L., Liao, W., Sun, T., Katul, G., Bou-Zeid, E., Maronga, B., 2024. Persistent urban heat. Science Advances 10, eadj7398. doi:10.1126/sciadv.adj7398.

Malings, C., Pozzi, M., Klima, K., Bergés, M., Bou-Zeid, E., Ramamurthy, P., 2017. Surface heat assessment for developed environments: Probabilistic urban temperature modeling. Computers, Environment and Urban Systems 66, 53–64. doi:10.1016/j.compenvurbsys.2017.07.006.

Marquès, E., Messier, K.P., 2025. Improved High Resolution Heat Exposure Assessment With Personal Weather Stations and Spatiotemporal Bayesian Models. GeoHealth 9, e2025GH001451. doi:10.1029/2025GH001451.

McBroom, B.D., Rahn, D.A., Brunsell, N.A., 2024. Urban fraction influence on local nocturnal cooling rates from low-cost sensors in Dallas-Fort Worth. Urban Climate 53, 101823. doi:10.1016/j.uclim.2024.101823.

van der Meer, W., Zantinge, F., Steeneveld, G., 2025. Urban fluxes for free: Estimating urban turbulent surface fluxes from crowdsourced meteorological canyon layer observations. City and Environment Interactions , 100201doi:10.1016/j.cacint.2025.100201.

Meier, F., Fenner, D., Grassmann, T., Otto, M., Scherer, D., 2017. Crowdsourcing air temperature from citizen weather stations for urban climate research. Urban Climate 19, 170–191. doi:10.1016/j.uclim.2017.01.006.

Muller, C., Chapman, L., Grimmond, C., Young, D., Cai, X., 2013. Sensors and the city: a review of urban meteorological networks. International Journal of Climatology 33, 1585–1600. doi:10.1002/joc.3678.

Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., Albergel, C., Arduini, G., Balsamo, G., Boussetta, S., Choulga, M., Harrigan, S., Hersbach, H., Martens, B., Miralles, D.G., Piles, M., Rodríguez-Fernández, N.J., Zsoter, E., Buontempo, C., Thépaut, J.N., 2021. ERA5-Land: A state-of-the-art global reanalysis dataset for land applications. Earth System Science Data 13, 4349–4383. URL: https://essd.copernicus.org/articles/13/4349/2021/, doi:10.5194/essd-13-4349-2021.

Napoly, A., Grassmann, T., Meier, F., Fenner, D., . Development and Application of a Statistically-Based Quality Control for Crowdsourced Air Temperature Data. Frontiers in Earth Science 6. doi:10.3389/feart.2018.00118.

Oke, T., 2006. Initial guidance to obtain representative meteorological observations at urban sites. IOM Report 81. World Meteorological Organization. Geneva.

Oyler, J.W., Ballantyne, A., Jencso, K., Sweet, M., Running, S.W., 2015. Creating a topoclimatic daily air temperature dataset for the conterminous United States using homogenized station data and remotely sensed land skin temperature. International Journal of Climatology 35, 2258–2279. doi:10.1002/joc.4127.

Pawlowsky-Glahn, V., Egozcue, J.J., 2006. Compositional data and their analysis: an introduction. Geological Society, London, Special Publications 264, 1–10. doi:10.1144/GSL.SP.2006.264.01.01.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research 12, 2825–2830.

Raharjo, M., Lukito, I., Sail, A., Soekirno, S., et al., 2025. A comparative analysis of air temperature data from automatic weather station and era5-land reanalysis in java, in: Journal of Physics: Conference Series, IOP Publishing. p. 012030. doi:10.1088/1742-6596/2989/1/012030.

Rasmussen, C.E., Williams, C.K.I., 2006. Gaussian Processes for Machine Learning. Adaptive Computation and Machine Learning, MIT Press.

Romero Rodríguez, L., family=Guerrero Delgado, given=Mcarmen, g.i., Castro Medina, D., Sánchez Ramos, J., Álvarez Domínguez, S., 2024. Forecasting urban temperatures through crowdsourced data from Citizen Weather Stations. Urban Climate 56, 102021. doi:10.1016/j.uclim.2024.102021.

Schatz, J., Kucharik, C.J., . Urban heat island effects on growing seasons and heating and cooling degree days in Madison, Wisconsin USA. International Journal of Climatology 36, 4873–4884. doi:10.1002/joc.4675.

Sharston, R., Singh, M., 2025. Urban morphology, urban heat island (UHI) and building energy consumption: A critical review of methods and relationships among influential parameters. Building Services Engineering Research & Technology 46, 561–584. doi:10.1177/01436244251339727.

Stewart, I.D., 2011. A systematic review and scientific critique of methodology in modern urban heat island literature. International Journal of Climatology 31, 200–217. doi:10.1002/joc.2141.

Stewart, I.D., Oke, T.R., 2012. Local Climate Zones for Urban Temperature Studies. Bulletin of the American Meteorological Society 93, 1879–1900. doi:10.1175/BAMS-D-11-00019.1.

Titsias, M., 2009. Variational Learning of Inducing Variables in Sparse Gaussian Processes, in: Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, PMLR. pp. 567–574.

Tong, S., Prior, J., McGregor, G., Shi, X., Kinney, P., 2021. Urban heat: An increasing threat to global health. BMJ 375, n2467. doi:10.1136/bmj.n2467, arXiv:34697023.

United States Geological Survey, 2024. Annual National Land Cover Database (NLCD) Collection 1 Products. doi:10.5066/P94UXNTS.

Venter, Z.S., Chakraborty, T., Lee, X., . Crowdsourced air temperatures contrast satellite measures of the urban heat island and its mechanisms. Science Advances 7, eabb9569. doi:10.1126/sciadv.abb9569.

Vicente-Serrano, S.M., Saz-Sánchez, M.A., Cuadrat, J.M., 2003. Comparative analysis of interpolation methods in the middle Ebro Valley (Spain): Application to annual precipitation and temperature. Climate Research 24, 161–180. doi:10.3354/cr024161.

Williams, C., Seeger, M., 2000. Using the nyström method to speed up kernel machines. Advances in neural information processing systems 13.

Wilson, A.G., Nickisch, H., 2015. Kernel interpolation for scalable structured gaussian processes (kiss-gp), in: Proceedings of the 32nd International Conference on Machine Learning - Volume 37, JMLR.org. p. 1775–1784.

Zou, J., Lu, N., Jiang, H., Qin, J., Yao, L., Xin, Y., Su, F., 2022. Performance of air temperature from era5-land reanalysis in coastal urban agglomeration of southeast china. Science of The Total Environment 828, 154459. doi:10.1016/j.scitotenv.2022.154459.