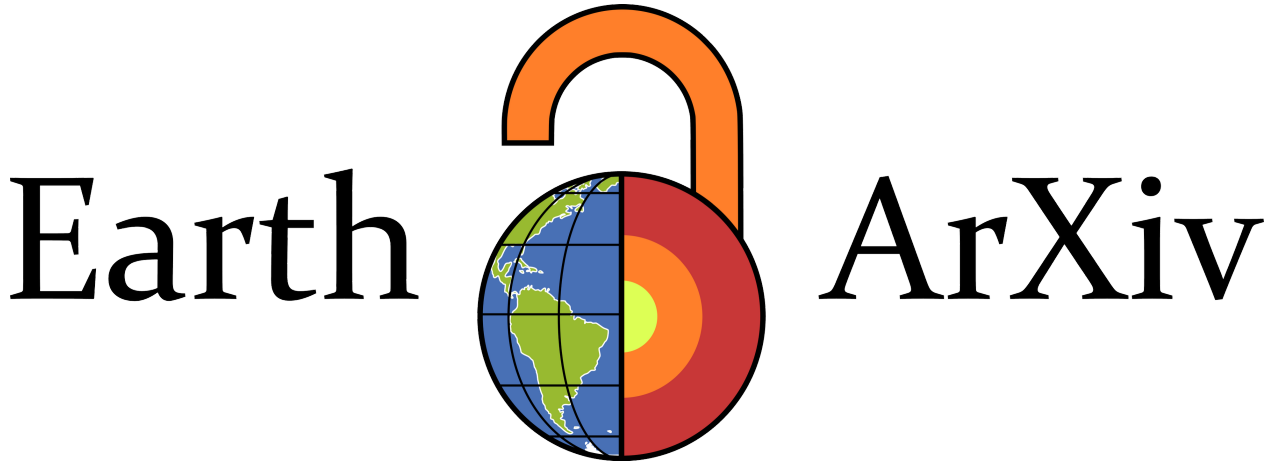

This manuscript has been submitted for publication in *Urban Climate*. Please note the manuscript has yet to be formally accepted for publication. Subsequent versions of this manuscript may have slightly different content. If accepted, the final version of this manuscript will be available via the 'Peer-reviewed 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)

Scalable and robust Gaussian processes for reanalysis of urban air temperature with crowdsourced meteorological data

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, which limits 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 produce interpolated air temperature fields with calibrated uncertainty estimates. To test our approach, we apply our model to six years of hourly data across Durham County, North Carolina. This area includes a medium-sized city, so we expect our model to generalize to similarly sized regions with sufficient sensor coverage. Compared to ERA5-Land, it improves estimates of canopy temperature (MAE=0.57°C versus ERA5-Land MAE=3.20°C on held-out locations) 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 these data,

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

we can better quantify urban heat exposure and its impact on health and energy.

Keywords: urban heat island, citizen-science, uncertainty quantification, Gaussian process, reanalysis, 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 assimilating 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 processes (GPs) excel 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 GPs at a fraction of the computational cost. The machine learning community has further improved our ability to apply SVGPs through the develop-

ment 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 which 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 incorporating noisy citizen-science into 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, non-linear 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 capture temperature variability.

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 Processes

2.3.1. Gaussian processes

A Gaussian process, also known as kriging, is a type of spatial interpolation technique in which the correlation structure between data points is learned as a function of distance, which provides a probabilistic estimate of the most likely values between data points. This

technique has long been used in spatial and spatiotemporal prediction, as these data often exhibit structured spatial and/or temporal covariances. A Gaussian process is formalized by a mean function and covariance function:

$$\begin{aligned} m(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})], \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x})(f(\mathbf{x}') - m(\mathbf{x}')] \end{aligned} \tag{1}$$

And given these definitions, we denote a Gaussian process as:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \tag{2}$$

We refer the reader to Rasmussen and Williams (2006) for a more complete introduction to the topic. For the purposes of this manuscript, we highlight the importance of defining the mean and covariances functions. The mean function is often simple, either a constant or a linear function. In the covariance function, we seek to define a parametric formula to describe how the residual error of the mean function correlates with respect to the covariates. We define this function in terms of the distance between points. In this manuscript, we deal with the Matérn class of covariance functions, which is defined by a smoothness parameter ν and length-scale parameter ℓ. A larger ν value connotes a smoother decay in covariance over distance, while a larger ℓ indicates a stronger correlation between distant points. In practice, we set ν beforehand and learn ℓ from the data.

Given a mean and covariance function definition, we then train our model by maximizing the marginal log likelihood with respect to the mean and covariance function hyperparameters $\boldsymbol{\theta}$. In the case of a linear mean function, we learn the linear parameters $\boldsymbol{\beta}$, such that $m(\mathbf{x}; \boldsymbol{\theta}) = \boldsymbol{\beta}\mathbf{x}$. For an example Matérn covariance function with $\nu = 1/2$, we typically learn an additional scale parameter σ^2 and the length-scale parameter ℓ, such that $k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}) = \sigma^2 \exp(-||\mathbf{x} - \mathbf{x}'||_2/\ell)$, with $||\cdot||_2$ referring to the Euclidean distance. Finally, because we assume some noise in our data, we add an additional hyperparameter ε^2. This corresponds to evaluating the Gaussian process at training points X with observations $\mathbf{y}$ as having value $\mathbf{f}$, then assuming $p(\mathbf{y} \mid \mathbf{f}) = \mathcal{N}(\mathbf{f}, \varepsilon^2)$. This completes the set of hyperparameters $\boldsymbol{\theta} = \{\boldsymbol{\beta}, \sigma^2, \ell, \varepsilon^2\}$. Given a set of N observations $\mathbf{y}$ at points X, we then define the marginal log likelihood with $\boldsymbol{\mu}_\theta = m(X; \theta)$ and $K_\theta = k(X, X'; \theta)$:

$$\log p(\mathbf{y} \mid X, \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_{\boldsymbol{\theta}})^\top (K_{\boldsymbol{\theta}} + \varepsilon^2 I)^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\boldsymbol{\theta}}) - \frac{1}{2}\log|K_{\boldsymbol{\theta}} + \varepsilon^2 I| - \frac{N}{2}\log 2\pi \tag{3}$$

Note that I indicates the identity matrix. We maximize this objective by iteratively differentiating with respect to the partial derivatives $\partial\boldsymbol{\theta}$ and performing gradient ascent. In practice, many packages exist to implement this algorithm. However, this objective demonstrates the principal computational challenge of Gaussian processes. The inversion of the $N \times N$ covariance matrix $K_{\boldsymbol{\theta}}$ has a computational complexity of $O(N^3)$, which must be computed for each iteration. This means that exact Gaussian processes with $N > 10,000$ often become infeasible on currently available hardware. In our case, one month of WU data is $N > 200$ sensors $\times$ 744 hours $> 100,000$ over a relatively small domain, so exact Gaussian processes are intractable.

2.3.2. Sparse variational approximations

Over the past two decades, research in the fields of computer science and machine learning has produced novel techniques for overcoming this problem. Many of these techniques make approximations using a set of learned $M \ll N$ inducing points which best describe the data (Williams and Seeger, 2000; Smola and Bartlett, 2000). These inducing points are parameterized with locations Z and values $\mathbf{u}$. This allows us to learn an approximation of $\mathbf{f}$ at training points X by conditioning on $\mathbf{u}$. That is, we have a joint distribution:

$$p(\mathbf{f}, \mathbf{u}) = \mathcal{N}\left(\begin{bmatrix} \mathbf{f} \\ \mathbf{u} \end{bmatrix}, \begin{bmatrix} m(X) \\ m(Z) \end{bmatrix} \begin{bmatrix} k(X,X) & k(X,Z) \\ k(Z,X) & k(Z,Z) \end{bmatrix} \right) \tag{4}$$

To obtain the conditional distribution, we use the shorthand $\mu_\cdot = m(\cdot)$ and $K_{\cdot\cdot} = k(\cdot,\cdot)$:

$$p(\mathbf{f} \mid \mathbf{u}) = \mathcal{N}\Big(\mathbf{f} \mid \mu_z + K_{XZ}K_{ZZ}^{-1}(\mathbf{u} - \mu_z), K_{XX} - K_{XZ}K_{ZZ}^{-1}K_{ZX}\Big) \tag{5}$$

The key benefit of this form is that we can evaluate the efficacy of our mean and covariance function with smaller computational complexity $O(NM^2 + M^3)$. However, to ensure this model provides a good approximation, we must now learn the locations Z and values $\mathbf{u}$. Furthermore, to calculate the marginal log likelihood, we need to integrate out $p(\mathbf{f}, \mathbf{u}) = p(\mathbf{f} \mid \mathbf{u})p(\mathbf{u})$. Without parameterizing $\mathbf{u}$ somehow, it is unclear how we learn the best inducing point locations and values.

The variational approach approaches this problem by approximating $p(\mathbf{u})$ using the parameterized distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ (Titsias, 2009). By ensuring this approximation is

close to the distribution of our model $p(\mathbf{u})$, we then have a set of global variables that we can optimize to maximize the marginal log likelihood. We refer the reader to Hensman et al. (2015) for complete derivations. To summarize their result, we balance the dual objectives of maintaining a close approximation while maximizing the marginal log likelihood by instead maximizing the evidence lower bound (ELBO):

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

With KL referring to the Kullback-Leibler divergence. Because this objective is summed over the data, we may achieve further computational improvements by dividing our data into mini-batches, which produces an unbiased estimate of the ELBO. By randomly sampling our data into batches of size B, we therefore reduce the per-iteration computational complexity of the dominant calculations from $O(NM^2)$ to $O(BM^2)$. So long as $B \ll N$, we improve convergence time significantly; in our experiments, we set $B = 512$. With both B and M fixed, this model scales efficiently to large sensor networks (although M should be set higher when fitting over larger areas or larger time periods).

2.3.3. Non-conjugate likelihoods

An additional benefit of the above objective is that we are no longer confined to parameterizing $p(\mathbf{y}|\mathbf{f})$ as a Gaussian distribution. If we recall Equation 3, conjugacy is required with exact Gaussian processes because it allows for closed form calculations of the marginal log likelihood and its partial derivatives. This precludes the use of non-conjugate likelihoods. Therefore, we cannot fit exact Gaussian processes to classification problems (e.g., with a binomial likelihood), and we cannot consider distributions with heavier tails, such as the Student's-t distribution. In the case of citizen science data, we may indeed have data that is noisier than possible to model with a Gaussian distribution.

2.3.4. An alternative loss function

An alternative interpretation of the sparse variational Gaussian process is suggested in Jankowiak et al. (2020). As the variational approximation introduces new parameters 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:

$$\log p(\mathbf{y} \mid X) \geq \sum_{i=1}^{N} \log \mathbb{E}_{q(f_i)}[p(y_i \mid f_i)] - \mathrm{KL}[q(\mathbf{u})||p(\mathbf{u})] \tag{7}$$

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 SVGP-ELBO, and with this alternative loss, the SVGP-PLL (predictive log likelihood).

2.3.5. Computational considerations

Many packages have since emerged to fit Gaussian processes and their variational approximations. We use the Python package GPyTorch, as it includes implementations of both standard and sparse variational Gaussian processes, support for non-conjugate likelihoods, as well as both loss functions discussed here (Gardner et al., 2018). Furthermore, because this package is built on PyTorch, GPU compatibility is supported, allowing for further computational gains.

3. Methods

3.1. Data Collection

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 station locations found in Weather Underground inform the linear interpolation of ERA5-Land data, as well as the processing of the land-use data to construct the features $z(\mathbf{s}, t)$, representing a land-use feature at location $\mathbf{s}$ at time t, 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.

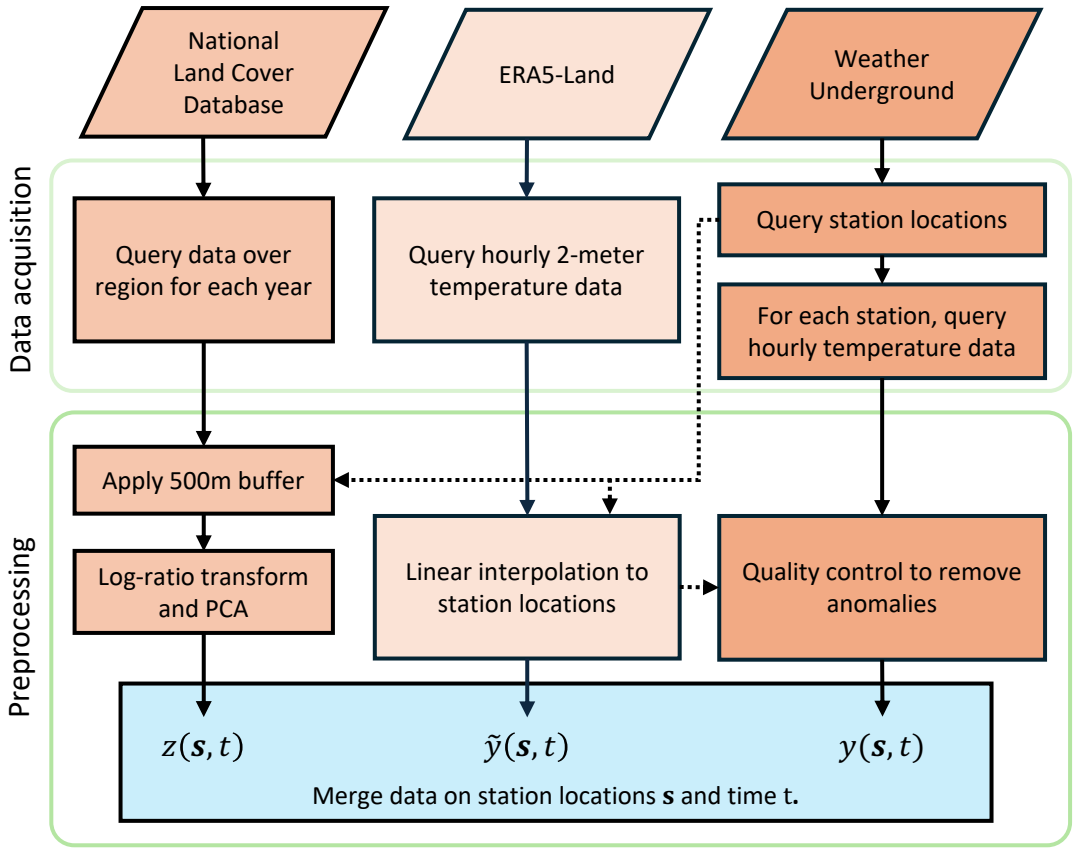

Figure 1: Summary of data acquisition and pre-processing pipeline.

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. 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 sub-

tract 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 real spatial variability.

ERA5-Land. ERA5-Land is a reanalysis dataset developed from the European Centre for Medium-Range Weather Forecasts (ECMWF) that assimilates 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 is 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 an assimilation technique to incorporate the noisy personal weather station data into the weather model.

NLCD. Lastly, we use the National Land Cover Database for an annually updated map of land-cover 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, to impute 2023's data for 2024. The primary objective of using this dataset is to develop a maximally informative covariate to feed into our model. To do this, we first calculate

the average percentage of each class in a 500m radius around each station location. This radius was chosen following guidance from Schatz and Kucharik and Oke (2006). We then perform a centered-log ratio transform to extract the first principal component from the transformed percentages, using the Principal Component Analysis implementation in sci-kit learn (Pedregosa et al., 2011). This provides us with a continuous feature for use in our model, in which higher values indicate more urbanization, and less values more rural. To produce final maps, we used the principal components to generate a gridded dataset of NLCD aligned with the interpolated ERA5-Land dataset at a 0.005° resolution. The transformed NLCD feature can be seen in Figure 2.

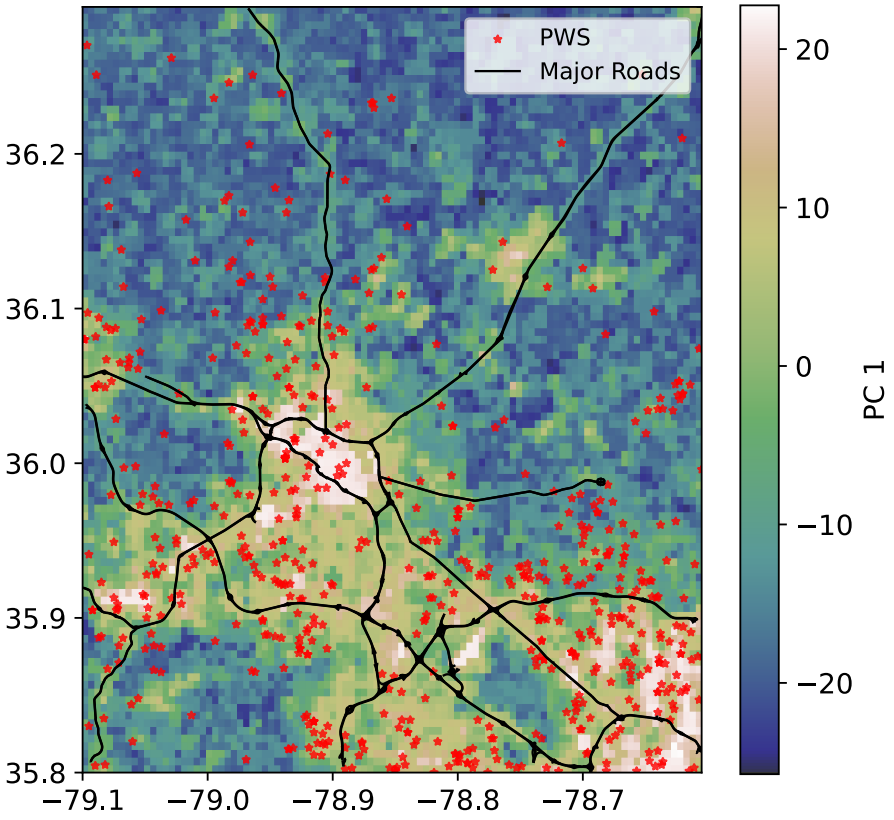

Figure 2: Principal component 1. Sensor locations denoted PWS, with major roads provided for reference. Pixels shown are 0.005° by 0.005°(approximately 500 by 500 meters).

3.2. Modeling approach

To develop our model, we start with a Sparse Variational Gaussian Process (SVGP), then define the mean and covariance functions, and finish the model definition by setting a likelihood and loss function for model training. We provide a graphical summary of these components in Figure 3.

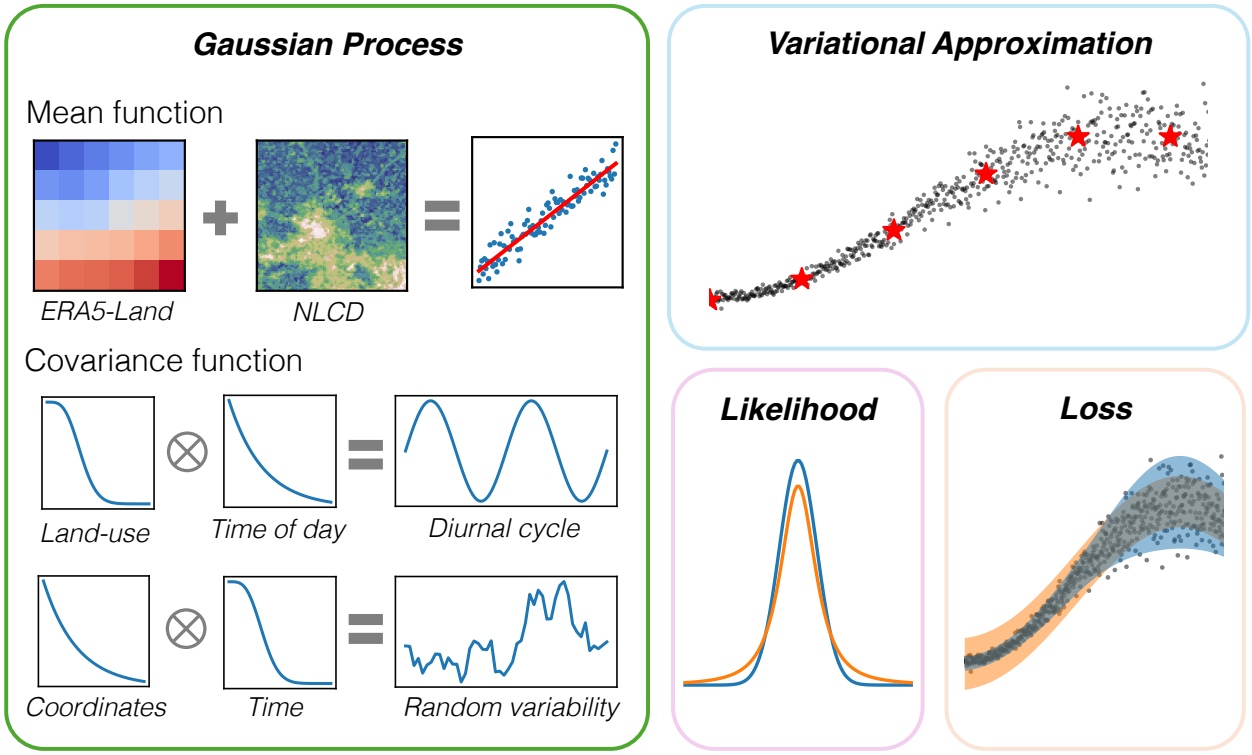

Figure 3: Model summary. The Gaussian process is defined with a mean and covariance function. The mean function centers model estimates on ERA5-Land, with an adjustment for the average effect of land-use on temperature. The covariance function captures the diurnal effect of land-use, along with random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

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.

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 kriging 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 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.

To define this function, we introduce variables z to refer to the principal component, $\boldsymbol{s}$ to refer to the spatial coordinates, and t to refer to time. We additionally generate a feature for capturing diurnal variability, using a sine and cosine transform:

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

Together we refer to these as $\Phi = \{\phi_1, \phi_2\}$. The complete covariance function is then defined as:

$$K = K_1(\Phi, \Phi', \mathbf{z}, \mathbf{z}') + K_2(t, t', \mathbf{s}, \mathbf{s}')$$

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:

$$\begin{aligned}
K_1 &= \sigma_1^2 K_a(\Phi, \Phi') K_b(\mathbf{z}, \mathbf{z}') \\
K_a &= \alpha_1 \exp\left(-\frac{||\Phi - \Phi'||_2}{\ell_a}\right) + \alpha_2 \\
K_b &= \left(1 + \frac{\sqrt{3}||\mathbf{z} - \mathbf{z}'||_2}{\ell_b}\right) \exp\left(-\frac{\sqrt{3}||\mathbf{z} - \mathbf{z}'||_2}{\ell_b}\right)
\end{aligned}$$

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 an added constant α, and K_b 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:

$$
\begin{aligned}
K_2 &= \sigma_2^2 K_c(t, t') K_d(\boldsymbol{s}, \boldsymbol{s'}) \\
K_c &= \left(1 + \frac{\sqrt{3}||t - t'||_2}{\ell_c}\right) \exp\left(-\frac{\sqrt{3}||t - t'||_2}{\ell_c}\right) \\
K_d &= \exp\left(-\frac{||\boldsymbol{s} - \boldsymbol{s'}||_2}{\ell_d}\right) + \alpha_3
\end{aligned}
$$

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. 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.

3.2.1. 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 $\boldsymbol{f}$, we typically assume that our observations $\boldsymbol{y}$ are normally distributed around $\boldsymbol{f}$, so $\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{f}, \sigma^2 I)$. Instead, we may decide to improve our model's robustness to outliers by considering the noise as following a Student's t-distribution. In this case, we must contend with a non-conjugate likelihood. Fortunately, the variational approximation naturally allows us to account for this case.

Predictive Log Likelihood. We also test the use of the aforementioned the predictive log likelihood (PLL) (Jankowiak et al., 2020). In the case of using the Student's t-likelihood, we replace the Normal distribution ($\mathcal{N}$) with the Student's t-distribution (t_ν with ν degrees of freedom). This gives us the loss function:

$$\mathcal{L}_{\text{pll}} = \sum_{i=1}^{N} \log t_\nu(y_i \mid \mu_{\mathbf{f}}(\mathbf{x}_i), \sigma_{obs}^2 + \sigma_{\mathbf{f}}(\mathbf{x}_i)^2) - \beta_{reg}\text{KL}(q(\mathbf{u})||p(\mathbf{u})) \tag{8}$$

3.3. Fitting procedure

To fit each model (i.e., each month's worth of data), we initialize inducing points by randomly sampling M points from the training set. Using the Adam optimizer, a learning rate of 0.01, and a Cosine Annealing learning rate scheduler, we train the model for 10 epochs using a batch size of 512. During model fitting, we highlight that the covariance hyperparameters are learned in conjunction with the inducing point locations and values. The variational distribution is defined as a Cholesky variational distribution, which uses the lower triangle of the covariance matrix for computational stability, while still maintaining covariance terms. This is in contrast with alternative variational distribution definitions, such as the mean field distribution, which assumes no covariance between induced values. The mean field distribution is computationally cheaper but typically performs worse.

3.4. Experimental set-up

The goals of our experiments are: (1) to determine the optimal number of inducing points; (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. By holding out sensors, we test spatial generalization across all locations. Performance is evaluated on these held-out stations using the metrics described below.

3.4.1. 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 X, we calculate this metric with respect to our held-out data points X_* with true values $\mathbf{y}_*$:

$$\text{NLPD}(\mathbf{y}_* \mid \mathbf{f}, X, X_*) = -\log p(\mathbf{y}_* \mid \mathbf{f}, X, X_*) \tag{9}$$

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 $\mathbf{x}_* \in X_*$, yielding the set of estimates $\mathbf{f}_*$. This procedure is given below:

$$\text{NLPD}(\mathbf{y}_* \mid \mathbf{f}, X, X_*) = -\log \int \text{Student}(\mathbf{y}_* \mid \mathbf{f}_*)\mathcal{N}(\mathbf{f}_* \mid \mathbf{f}, X_*)d\mathbf{f}_* \tag{10}$$

$$\approx -\log\left(\frac{1}{S}\sum_{i=1}^{S}\text{Student}(y_* \mid \mathbf{f}_{*,i})\right) \tag{11}$$

$$= -\log\left(\sum_{i=1}^{S}\text{Student}(y_* \mid \mathbf{f}_{*,i})\right) + \log(S) \tag{12}$$

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 look at model performance in Table 1 and the 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.

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.

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 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 provide 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.

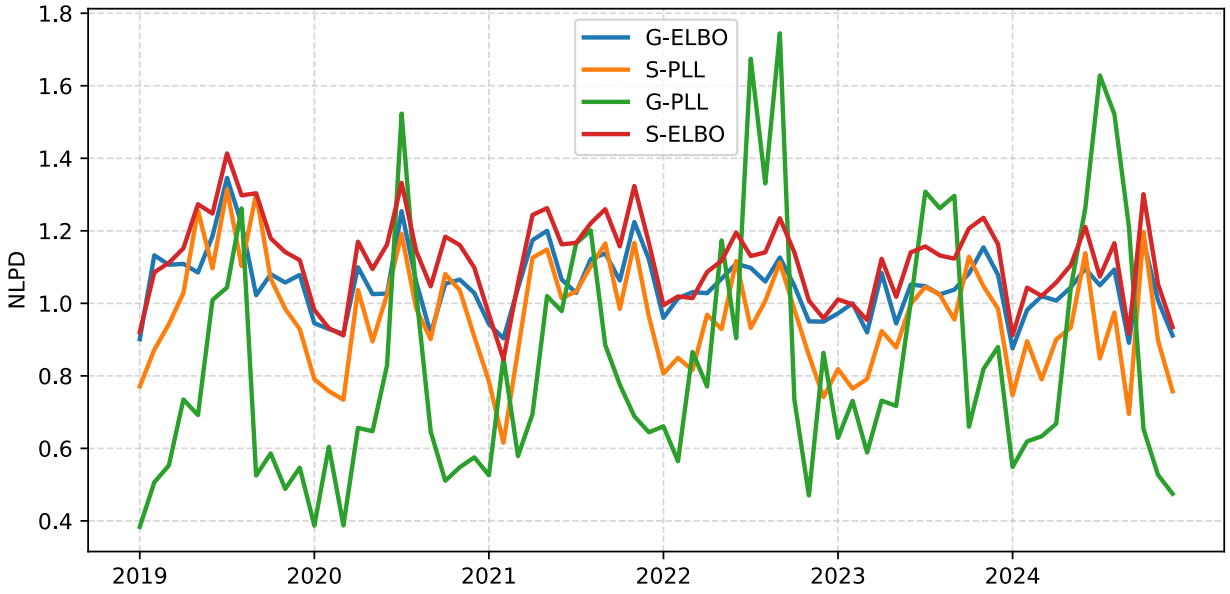

Figure 4: Negative Log Predictive Density by month on the held-out data.

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.

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

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.

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.

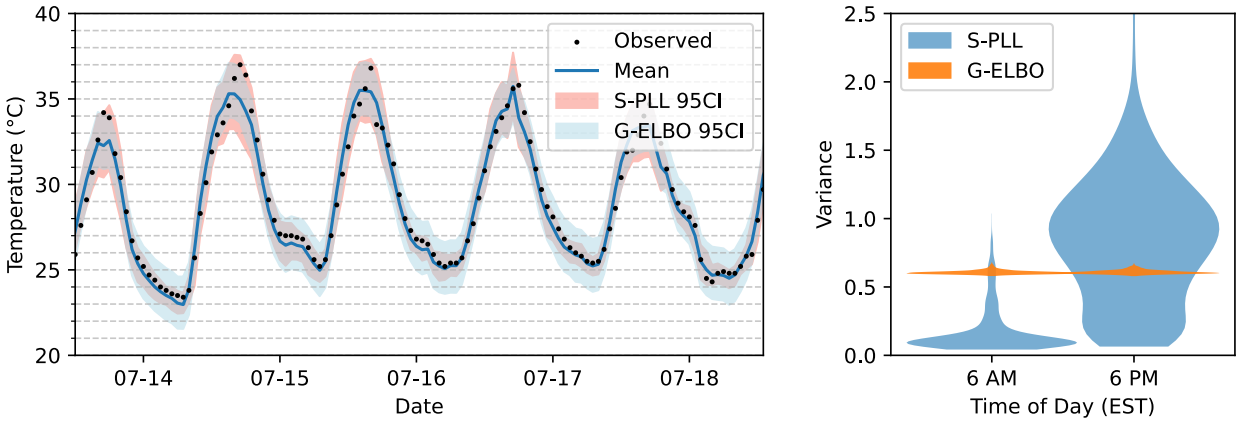

(a) (b)

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.

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 6, 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.

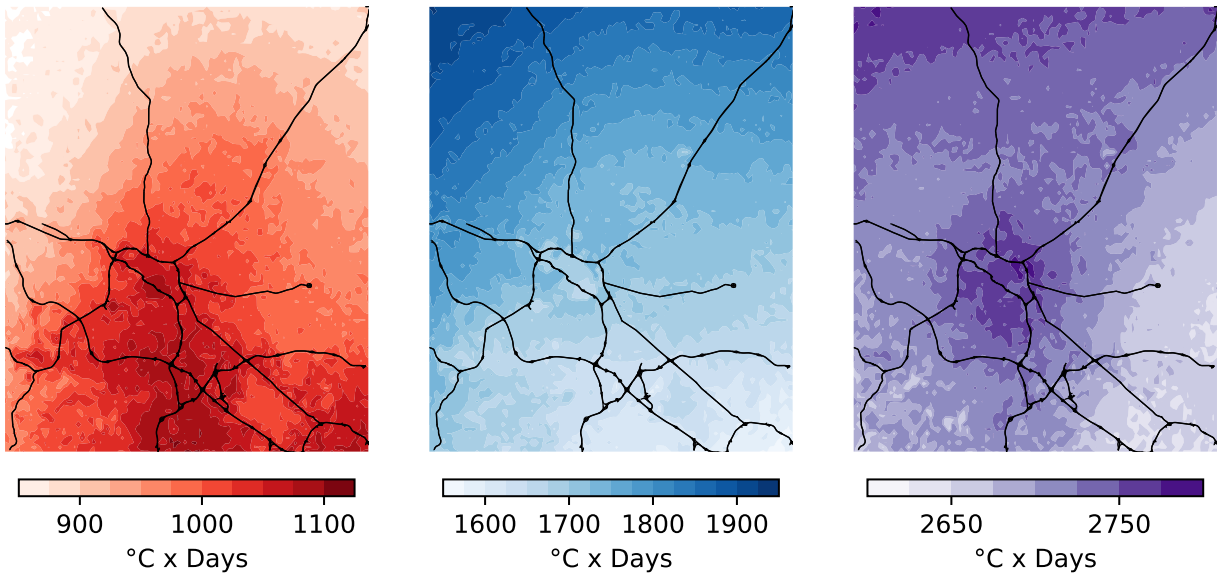

(a) Cooling Degree Days (b) Heating Degree Days (c) CDD + HDD

Figure 6: 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 7. 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.

Perhaps the most dramatic difference is seen when comparing clusters 4 and 5 in Figure 8, 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 days, whereas cluster 5 represents about 17% of the days.

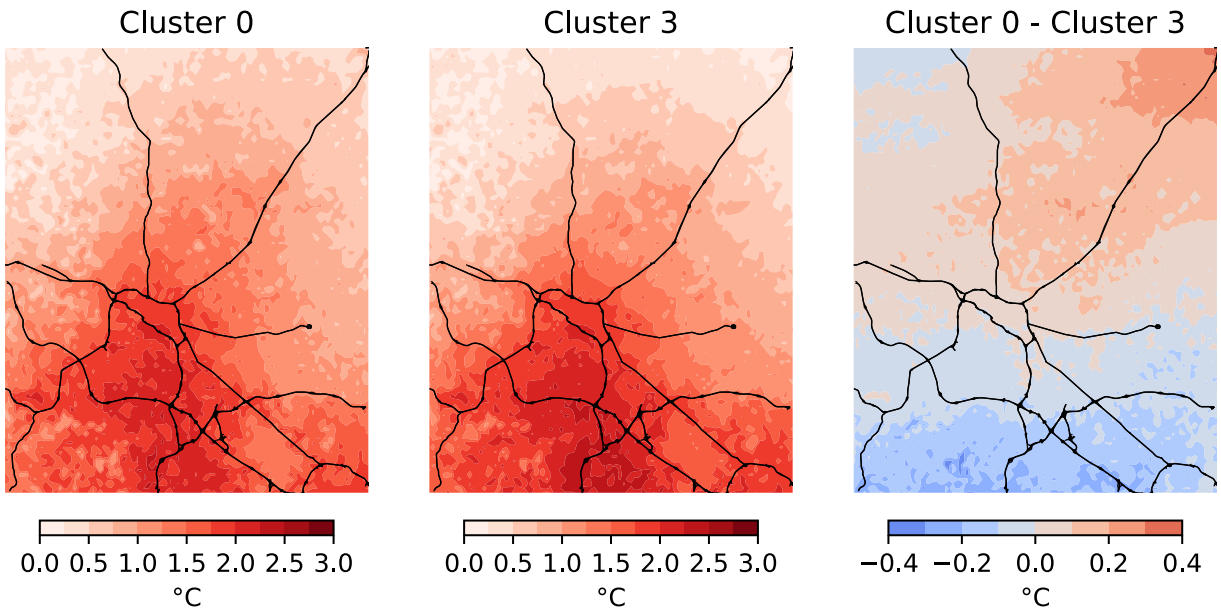

(a) Southerly (b) Northerly (c) Difference

Figure 7: The effect of southerly versus northerly winds

4.3. Hours above threshold

Lastly, we use our model to evaluate the number of hours above 35°C (95°F), shown in Figure 9. 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 (9a) and standard deviation (9b). We visualize the spatial density of these estimates in Figure 9c, 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, as well as the potential for localized cooling interventions. 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.

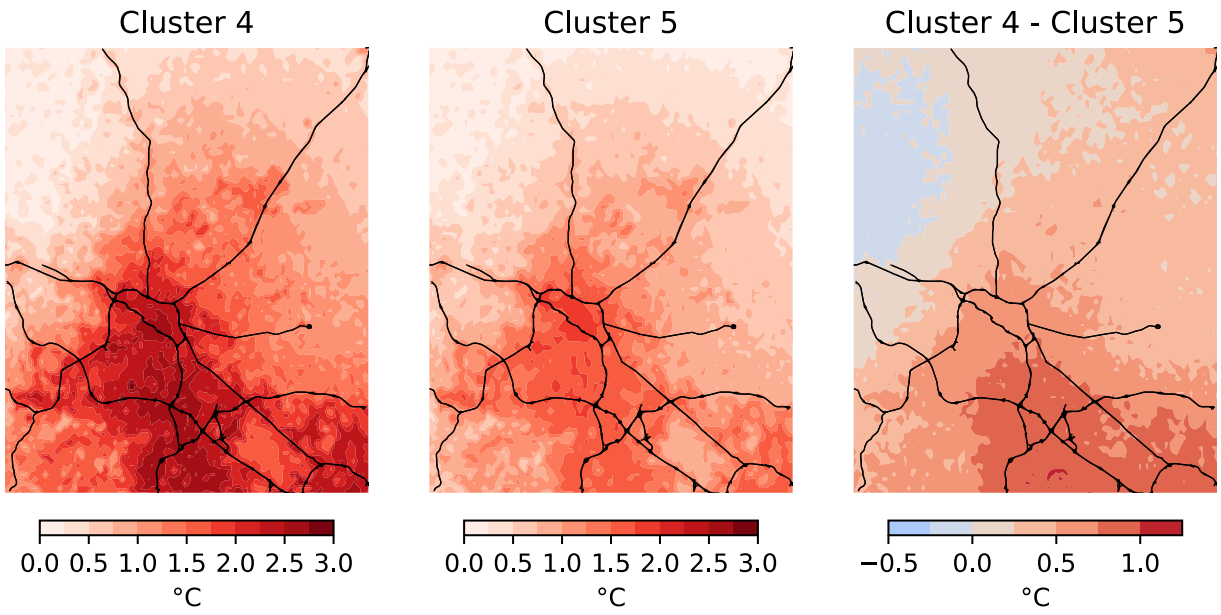

(a) Southeasterly (b) Southwesterly (c) Difference

Figure 8: The effect of southeasterly versus southwesterly winds.

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,

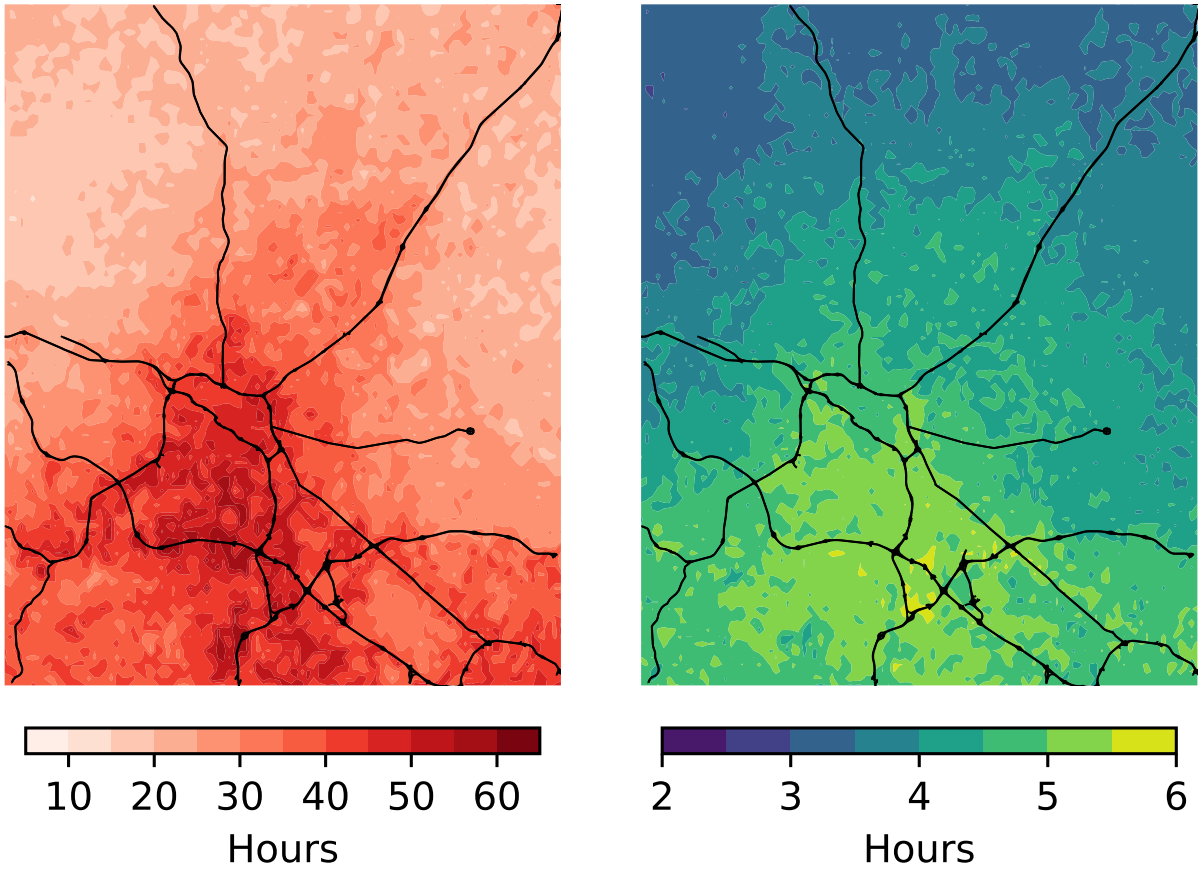

(a) Mean

(b) Standard Deviation

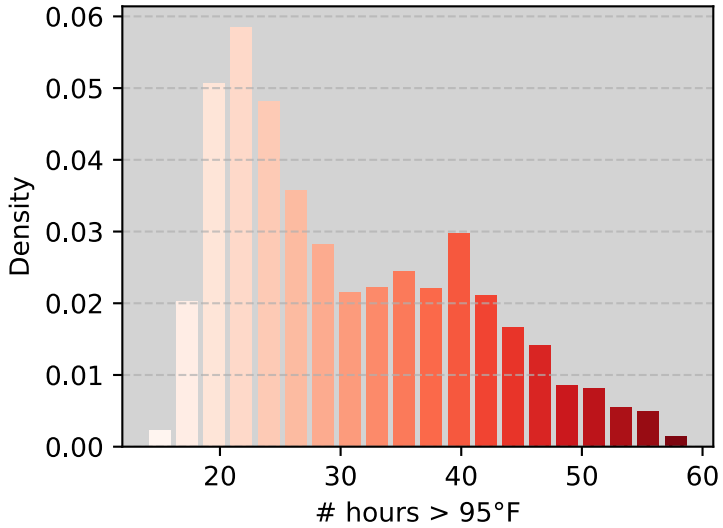

(c) Distribution

Figure 9: The estimated number of hours exceeding 35°C (95°F).

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

We could also improve the covariance definition and feature set. While the sparse GP 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.

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, versus local-scale effects.

In addition to improving the parameterization of the urban surface, this model presents 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 the appendix, 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. 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 A.10, 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.

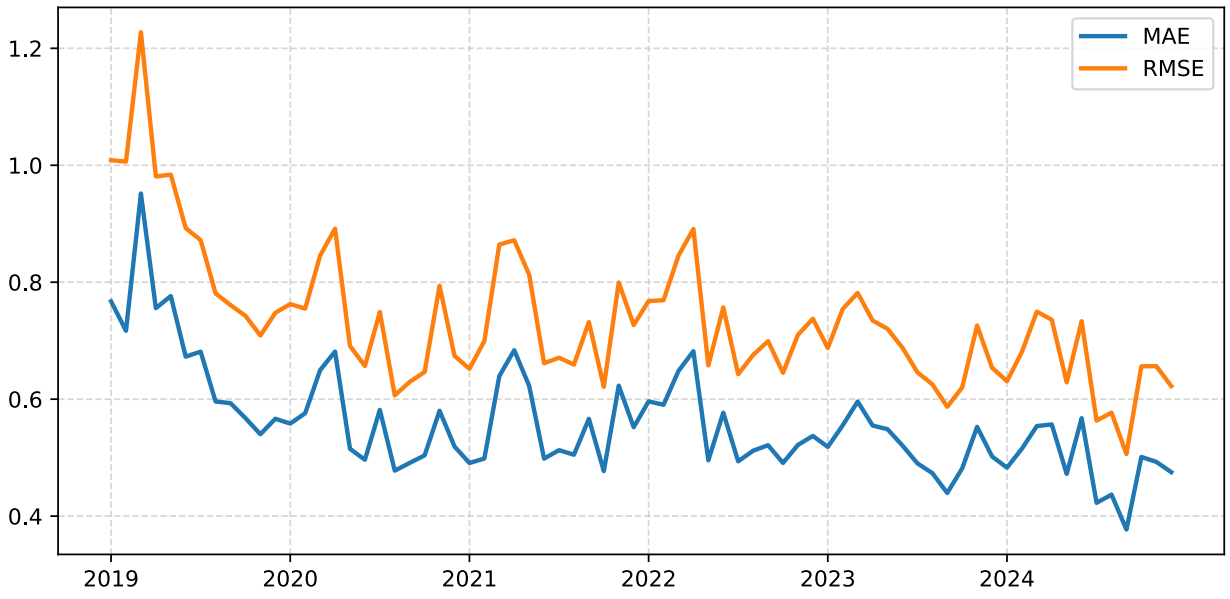

Figure A.10: Dew point mean absolute error and root mean squared error, monthly.

We additionally show average evening dew points by season in Figure A.11. 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).

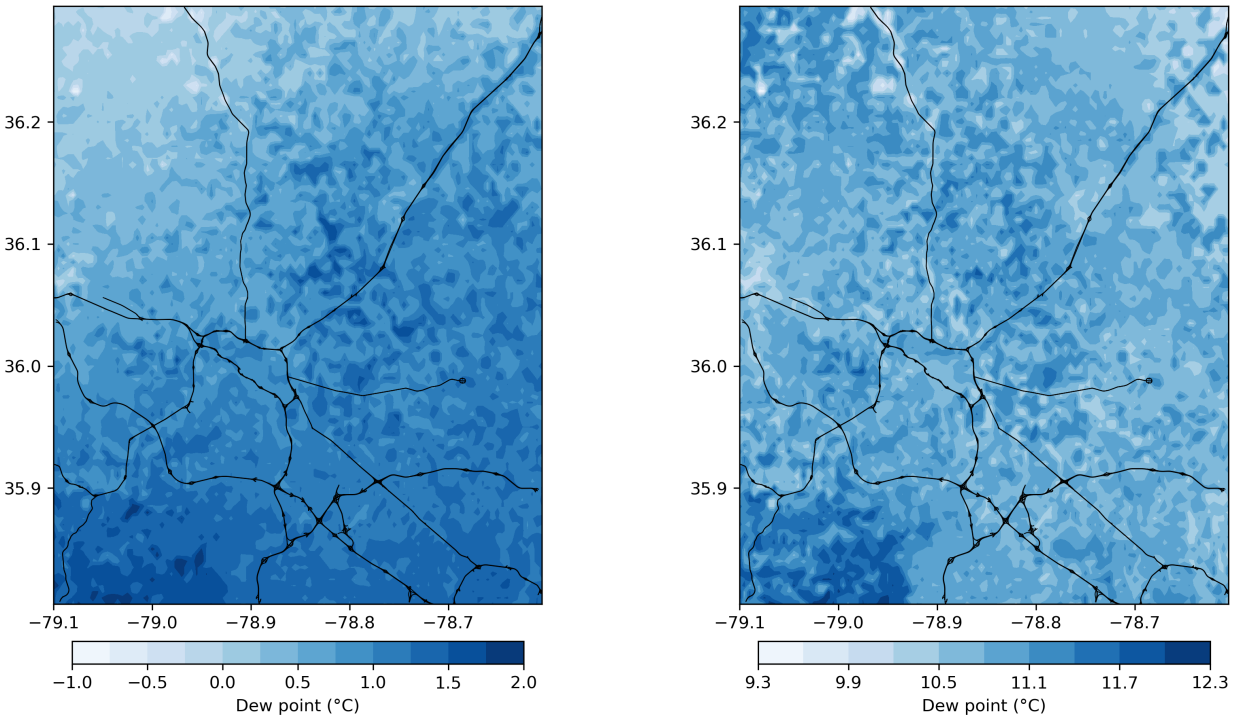

(a) January

(b) April

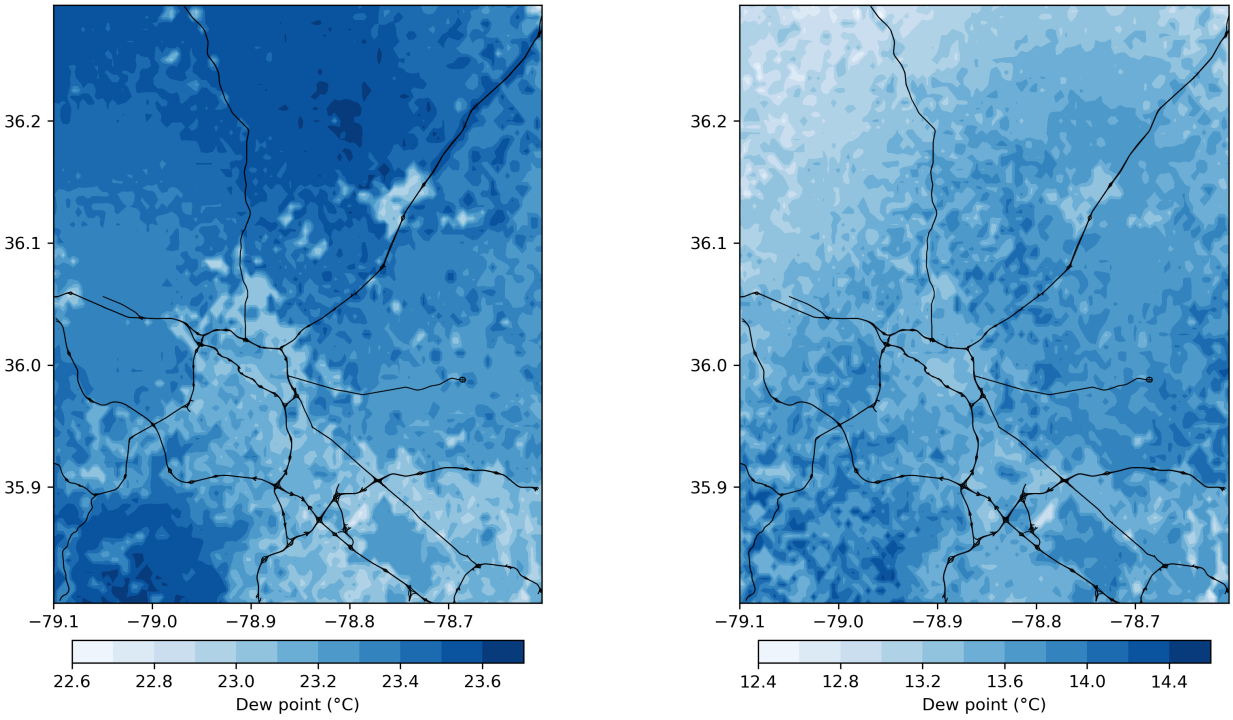

(c) July

(d) October

Figure A.11: Average 8pm (local time) Dew Point. Evening is used, as this is when the urban dry island effect is most active.

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, in: 12th International Conference on Urban Climate, Rotterdam, The Netherlands. URL: https://meetingorganizer.copernicus.org/ICUC12/ICUC12-59.html, doi: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: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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/qj.2836, doi: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 crowd-sourced weather observations using multi-fidelity methods. PLOS Climate 3, e0000216. URL: https://journals.plos.org/climate/article?id=10.1371/journal.pclm.0000216, doi: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. URL: https://www.cambridge.org/core/journals/journal-of-fluid-mechanics/article/processes-controlling-atmospheric-dispersion-through-city-centres/BADEBAA7CB3D552250BC345588555F1B, doi:doi:10.1017/jfm.2014.661.

Brousse, O., Simpson, C., Kenway, O., Martilli, A., Krayenhoff, E.S., Zonato, A., Heaviside, C., 2023. Spatially explicit correction of simulated urban air temperatures us-

ing crowdsourced data. Journal of applied meteorology and climatology 62, 1539–1572. doi: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. URL: https://iopscience.iop.org/article/10.1088/1748-9326/ac5c0f, doi: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. URL: https://www.nature.com/articles/s41467-024-49276-z, doi: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. URL: https://asr.copernicus.org/articles/17/153/2020/, doi: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: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: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. URL: https://www.sciencedirect.com/science/article/pii/S2210670724000374, doi: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: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 temperature mapping using machine learning. Science of The Total Environment 841, 156737. doi: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: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. URL: https://www.frontiersin.org/journals/environmental-science/articles/10.3389/fenvs.2021.720747/full, doi: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: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: Advances in Neural Information Processing Systems.

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. URL: https://proceedings.neurips.cc/paper/2021/hash/c6b8c8d762da15fa8dbbdfb6baf9e260-Abstract.html, doi: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: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. URL: https://www.nature.com/articles/s41467-021-22799-5, doi: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.

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: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: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. URL: https://ehp.niehs.nih.gov/doi/10.1289/ehp.1103532, doi: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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/2023GL107364, doi:doi:10.1029/2023GL107364.

Levi, D., Gispan, L., Giladi, N., Fetaya, E., 2022. Evaluating and calibrating uncertainty prediction in regression tasks. Sensors 22. URL: https://www.mdpi.com/1424-8220/22/15/5540, doi: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. URL: https://www.science.org/doi/full/10.1126/sciadv.adj7398, doi: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. URL: https://www.sciencedirect.com/science/article/pii/S0198971517300613, doi: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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/2025GH001451, doi: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. URL: https://www.sciencedirect.com/science/article/pii/S2212095524000191, doi: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:doi: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. URL: https://www.sciencedirect.com/science/article/pii/S2212095517300068, doi: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: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: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. URL: https://www.frontiersin.org/journals/earth-science/articles/10.3389/feart.2018.00118/full, doi: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:doi:10.1002/joc.4127.

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: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. URL: https://www.sciencedirect.com/science/article/pii/S2212095524002177, doi: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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/joc.4675, doi: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. URL: https://doi.org/10.1177/01436244251339727, doi:doi:10.1177/01436244251339727.

Smola, A., Bartlett, P., 2000. Sparse greedy gaussian process regression. Advances in neural information processing systems 13.

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: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. URL: https://journals.ametsoc.org/view/journals/bams/93/12/bams-d-11-00019.1.xml, doi: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. URL: https://proceedings.mlr.press/v5/titsias09a.html.

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

United States Geological Survey, 2024. Annual National Land Cover Database (NLCD) Collection 1 Products. URL: https://www.sciencebase.gov/catalog/item/655ceb8ad34ee4b6e05cc51a, doi: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. URL: https://www.science.org/doi/10.1126/sciadv.abb9569, doi: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: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.

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:doi:10.1016/j.scitotenv.2022.154459.