Gaussian process forecasting of sparse ecological time series

Parul V. Patil, *[1], Robert B. Gramacy[1], and Leah R. Johnson[1]

[1]Department of Statistics, Virginia Tech, Blacksburg, Virginia 24061, USA

Abstract

Ecological time series are often unevenly sampled in time. That is, because the sampling processes used are resource intensive, data may be collected infrequently, or with adaptive frequencies triggered by presence of a target variable. When the data are irregularly spaced, standard time series methods may not be directly applicable. Instead, approaches that take inspiration from linear regression (LR) may be appropriate. In this paper, we explore flexible, nonparametric Gaussian Process (GP) models as tools for producing forecasts of unevenly sampled observations. Our example is data on abundances of nymphal *Amblyomma americanum* from nine locations spread across the eastern United States, collected by NEON. The data exhibit highly variable sampling regimes and abundance levels across sites and time. We implement two versions of GPs to forecast tick abundance and benchmark them models against LR approaches. Both GPs are able to capture population patterns without the need for forecasting additional drivers, such as temperature, or specifying a specific relationship between the response and predictors. Further, our Heteroskedastic Gaussian Process (HetGP) model allows for flexible bounds and yields improved uncertainty quantification. We find that GP models provide an effective method to forecast irregularly sampled populations at short to intermediate time scales.

Keywords: *Amblyomma americanum*; Ecological forecasting; Gaussian process; Lone-star ticks; NEON forecasting challenge; Tick populations.

*Corresponding author: parulvijay@vt.edu

1 Introduction

Accurate forecasts of ecological time series are important for a variety of applications. For example, forecasts of algal blooms can ensure proper resource management and safety of drinking water quality (Carey et al., 2022) or pest forecasting can be useful for farmers so they can spray crops with pesticides in a timely manner improving efficacy while reducing the risk of adverse health effects (Liu et al., 2022; Ponomarev et al., 2019). Forecasting of species dynamics over medium to long-term horizons can also be used to inform policy, for example, forecasting the effects of climate change on *Aedes*-borne viruses to inform public health measures Ryan et al. (2019) or predicting the growth of chytrid fungus in frogs to help prevent rapid decline of amphibians (Gajewski et al., 2021). Although forecasting populations of species is useful in many domains, it can be challenging as monitoring and sampling populations can be difficult. For instance, observing endangered frogs affected by chytrid in the Sierra Nevada mountains of California requires researchers haul sampling equipment on foot for hours to collect a single time point. As a result, populations may be irregularly and infrequently sampled which presents challenges for forecasting efforts.

One example of irregularly sampled populations for which short to medium term forecasts would be useful for making decisions are ticks that spread infectious diseases. In the United States and around the world, ticks adversely affect humans and animals, primarily by spreading tick-borne pathogens (Mitchell et al., 2020; Levin et al., 2017; Paddock et al., 2016; Piesman and Eisen, 2008). In the United States, diseases caused by tick-born pathogens include Lyme disease and Babesiosis, both transmitted by the blacklegged tick, *Ixodes scapularis*; and human Ehrlichiosis (Gaines et al., 2014; Raghavan et al., 2014; Childs and Paddock, 2003), Heartland virus (Savage et al., 2013), and Tularemia (Goddard and Varela-Stokes, 2009) (among others) transmitted by the lone-star tick (*Amblyomma americanum*). Apart from spreading pathogens, bites from ticks can also cause long term issues such as alpha-gal syndrome which can result in a life long meat allergy (Steinke et al., 2015; van Nunen, 2015; Commins et al., 2011). Predicting tick density, as a proxy for bite risk, can help assess likelihood of pathogen transmission (Léger et al., 2013; Pfäffle et al., 2013). This can be beneficial for resource management, public safety, and preventative health measures (Estrada-Peña, 2001a; Randolph and Rogers, 1997)

In this study, we focus on modeling the seasonal density of the lone star tick, *Amblyomma americanum*, in forested plots across geographic locations within the United States. Lone star ticks are native to the eastern and southeastern United States and parts of Mexico (Molaei et al., 2019; Monzón et al., 2016; Springer et al., 2014). Their range has been expanding due to climate change (Rochlin et al., 2022; Raghavan et al., 2019; Cortinas and Spomer, 2013; Masters et al., 2008). Notably, there have been reports of these ticks in northern regions of the U.S. that were historically too cold for their survival (Paddock and Yabsley, 2007). This expansion is concerning, as lone star ticks are responsible for a significant proportion of tick-borne diseases in North America (Goddard and Varela-Stokes, 2009).

To forecast tick density, scientists have examined seasonal trends that affect population dynamics (Kollars Jr. et al., 2000). Classification models such as Generalized linear models (GLM), Maximum Entropy models (MaxEnt; Raghavan et al., 2016), and Classification and Regression Trees (CART) which are usually used in the field of Machine Learning (Zannou et al., 2021) can be used to study factors influencing presence of ticks. These models can incorporate weather and climate data such as temperature, moisture, precipitation (Williams et al., 2015; Springer et al., 2015; Estrada-Peña, 2001b; Haile and Mount, 1987). Additionally, they can incorporate location-based features such as elevation and forest cover (Slatculescu et al., 2020; Hahn et al., 2016) which may influence tick abundance (Paddock and Yabsley, 2007). These models may also include density of hosts such as deer density in addition to climate and vegetation (Wimberly et al., 2008). Although classification models cannot be used to forecast abundance,

they provide insights into the factors affecting tick distribution.

Tick density would conventionally be modeled using a time series analysis (Brugger et al., 2018; Meltzer and Norval, 1992), wherein we study the data and capture a trend using xxan auto-regressive model with fine tuning. This approach usually works fairly well when a time series is evenly sampled in time. However, most tick data sets are not evenly sampled in time. Ticks can be collected in a variety of ways which include tick drags, CO_2 traps, and baited traps. The trap methods are easier as they require less time and effort, but are not as effective as tick drags (Falco and Fish, 1992). The drag method involves an experimenter dragging a cloth across the field to collect ticks which are then identified and reported based on their taxonomy. This method is time and effort intensive, and so adaptive sampling schedules are often used to minimize the change that a drag is conducted when ticks are not present. This can lead to irregular sampling resulting a small dataset irregular gaps between sampling times. In such cases, time series methods may be inapplicable without substantial data pre-processing (such ass aggregating over large time intervals), which could dampen the signal or bias trends (Little and Rubin, 2019).

In this paper, we explore the use of Gaussian Process (GP) modeling for time dependent species data. One advantage of GP models is that they do not require regularly sampled time series. Additionally, we can build models that do not rely on direct environmental/weather factors (e.g., temperature, humidity) that themselves need to be forecast, nor must we assume specific functional forms between the response and predoctors. Rather, we use GPs to learn the historic patterns across locations and forecast tick density for the future years. Below, we begin by introducing the problem setup with details of the data source followed by visualization and pre-processing steps (Section 2). Next we describe a Linear Regression (LR) modeling technique that is a more classic approach this problem that we use as a comparitor (Section 3). We then introduce the GP model and its implementation (Section 4). In Section 5 we describe an extension of the GP method and our proposed approach for this specific forecasting problem - the Heteroskedastic Gaussian Process (HetGP) Model. Finally, we compare all the modeling approaches across multiple metrics (Section 6) and conclude with a brief discussion of the methods and the future scope (Section 7).

2 Data Description and Preparation

2.1 Data Overview

2.1.1 *Amblyomma americanum*

This study focuses specifically on *Amblyomma americanum*, commonly known as the lone star tick, also referred to as the northeastern water tick or the turkey tick. These ticks inhabit wooded areas and primarily feed on white-tailed deer (Allan et al., 2010; Paddock and Yabsley, 2007), though wild turkeys are common hosts for immature stages, especially in the Midwest (Mock et al., 2001). Lone star ticks aggressively look for hosts judging carbon-dioxide and movement patterns (Childs and Paddock, 2003). These ticks make up a large majority of human biting ticks (Mixson et al., 2006; Goddard, 2002; Felz et al., 1996) responsible for a wide range of tick borne diseases in the Unites States (Mitchell et al., 2020; Levin et al., 2017).

Lone star ticks follow a three-host life cycle that spans approximately 8 months to 2 years, depending on environmental conditions such as temperature (Troughton and Levin, 2007). The life cycle includes larval, nymphal, and adult stages. At each stage, the tick requires a blood meal from a host, typically mammals, birds, rodents, or humans to develop further (Kollars Jr. et al., 2000). Feeding usually lasts one to three weeks per stage (Troughton and Levin, 2007). Notably, lone star ticks are capable of feeding

on humans during all developmental stages (Fisher et al., 2006), thereby maintaining a consistent risk of human exposure to tick-borne pathogens throughout their life cycle.

2.1.2 NEON Forecasting Challenge

Our forecasting effort was based on a challenge set up by Ecological Forecasting Initiative Research Coordination Network (EFI-RCN) using data from National Ecological and Observatory Network (NEON; Thomas et al., 2023). We focus on Tick Populations forecasting using data from the NEON challenge. These data were collect across nine NEON locations located in eastern and southern parts of the United States shown in Fig. 1.

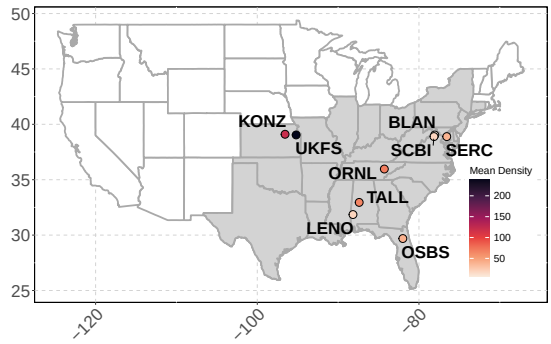

Site ID	Name	State
BLAN	Blandy Experimental Farm	VA
KONZ	Konza Prairie Biological Station	KS
LENO	Lenoir Landing	AL
ORNL	Oak Ridge National Lab	TN
OSBS	Ordway-Swisher Biological Station	FL
SCBI	Smithsonian Conservation Biology Institute	VA
SERC	Smithsonian Environmental Research Center	MD
TALL	Talladega National Forest	AL
UKFS	University of Kansas Field Station	KS

Figure 1

The target data includes the date of collection, the location, the target species, observed density of ticks per plot (which is the response) and the iso-week of collection. We have visualized the average tick density observed in the past at each location in Fig. 1. As part of the NEON challenge, we also have a subsidiary dataset (site dataset; Site Metadata, 2025), containing details about each of the nine locations such as the coordinates, altitude (mean, minimum and maximum). Additionally, the site dataset also contains information about certain ecological features that could be of interest (e.g., precipitation, temperature, vegetation, foliage).

Ticks are sampled using the "tick drag" method by dragging a $1m^2$ cloth across the plots in transects of 160m (National Ecological Observatory Network (NEON), 2024). This occurs once every few weeks at a given location during the sampling season which is determined by the foliage at the location. The samples are then taken to the lab for taxonomic identification where the life stage and the taxonomic rank of the species is identified. The species of interest for this study is *Amblyomma americanum*, and the focus is specifically on ticks that are currently in their nymphal phase.

2.1.3 Data Visualization

In this section, we highlight key observations across locations that inform our modeling decisions. We illustrate the location-level variation in Fig. 2 by showing tick density recorded at three representative sites: Lenoir Landing (LENO), Smithsonian Conservation Biology Institute (SCBI), and Talladega National Forest (TALL) out of the nine locations. The number of samples collected varies across locations, as tick dragging did not begin simultaneously and was not conducted on the same days across locations. For example, some locations such as TALL and SCBI have data dating back to 2014, while others like LENO only have records starting in 2016.

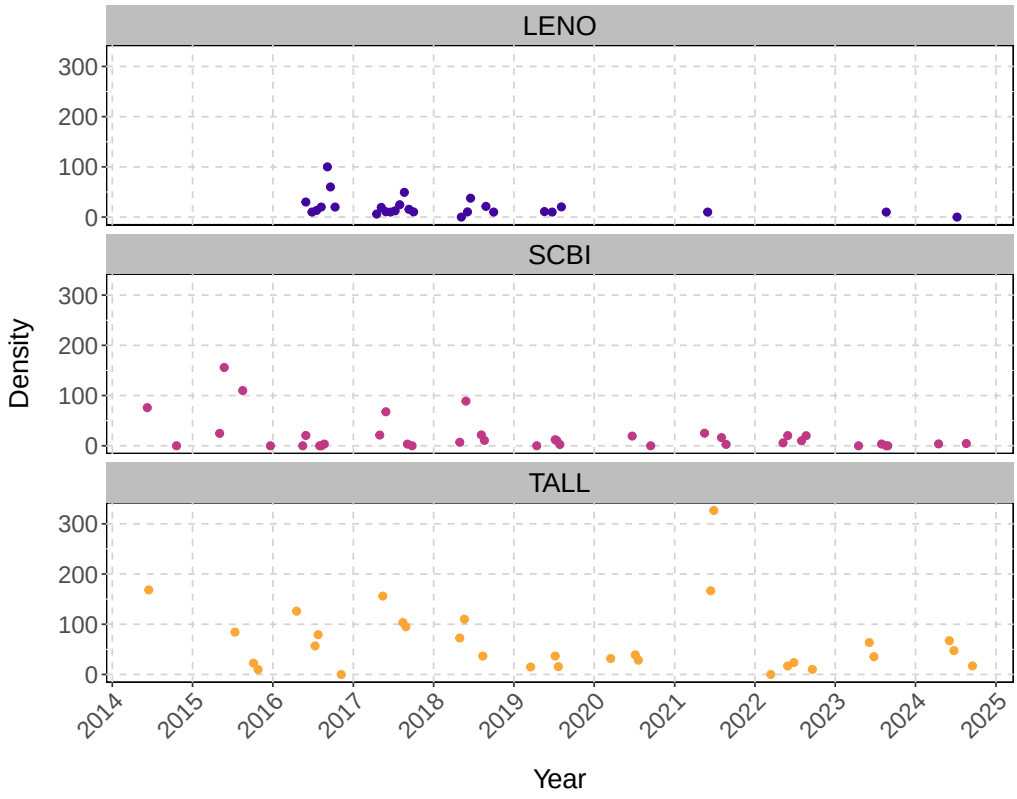

Figure 2

A broader review of all the locations suggests that data collection was influenced by the seasonality which was determined by environmental factors, such as the likely presence of snow. Typically, drags begin around March and conclude by October. This is because tick dragging during the snow is not possible and tick abundance and activity is very low in cold temperatures. On rare occasions, drags did occur in peak winter at warmer locations (e.g. SCBI, BLAN) where ticks remain active year-round.

The amplitude of tick density varies substantially across locations. For instance, peak densities at TALL reach around 300 ticks per plot, while LENO peaks at approximately 100 ticks per plot. Some other sites show extreme peaks with over 600 ticks per plot during peak season. Observe in Fig. 2 tick density has a seasonal pattern. It is is lower in colder months and higher during warmer months at each location. However, the magnitude of seasonal peaks can vary significantly from year to year, even at the same site.

2.2 Preprocessing Steps

2.2.1 Data Transformation

To forecast tick density, we explore some simple modeling methods such as Linear Regression (LR) and the Gaussian Process (GP) approach. A common assumption of both types of models is that the response must be normally distributed around the mean. However, tick density is non-negative which makes transformation necessary. We must ensure that any chosen transformation will allow for a sensible and stable back-transformation for predicted means and uncertainty estimates. Standard transformations like the square root or logarithm are commonly used in such cases. However, both have drawbacks. The

log transform is undefined at zero, over expands smaller values in the forward transform. Further, it exponentiates the magnitude when back-transformed, making uncertainty estimates unstable. A square-root transform can result in negative mean predicted densities and bounds which when squared, will result in insensible UQ (upper bound < lower bound). We propose a hybrid transform to the shifted tick density as,

$$Y' = \begin{cases} \sqrt{Y+1}, & \text{if } (Y+1) > 1. \\ \log(Y+1), & \text{otherwise;} \end{cases} \tag{1}$$

where, Y' is the transformed response and Y is on the original scale, i.e., the observed tick density. We apply this transform to shifted density so we eliminate discontinuity at zero. The forward transform is always square root so we avoid over-expansion of smaller values. While back-transforming, we avoid squaring negative values; rather exponentiate them $\forall\, Y \leq 0$. We apply this transformation to the response before modeling to ensure consistency and enable direct comparison of results.

2.2.2 Train-Test Split

To ensure good forecasts and evaluate our method, we need a training set, as well as a holdout (testing) set. For time series data, it is standard practice to use earlier data for training and more recent data for testing. We selected a cutoff date of 2022-12-31. All data before this date were used for training, and all data after it formed the testing set. This provided a sufficient number of points to assess model performance. Because predictions made only at sparse time points do not capture underlying periodic trends well, we constructed a uniform weekly grid from 2014 to the latest date in the dataset. Each of the trained models was used to predict tick density at every week in this grid. We then picked out the predictions corresponding to our holdout date and evaluated performance based on these points.

We benchmarked our method using root mean square error (RMSE; Appendix: Section S1), continuous rank probability score (CRPS) (Gneiting and Raftery, 2007), and coverage. RMSE measures prediction accuracy but ignores uncertainty, while CRPS accounts for both accuracy and uncertainty and can be computed in closed form for Gaussian distributions. Coverage is the percentage of data points within the 90% predictive intervals (PIs), which ideally should match the nominal level. We calculated RMSE and CRPS on the transformed scale for all models. We also assess the quality of predictions over time, by visualizing each model on the original scale. This enables us to evaluate model performance, particularly during data-sparse months such as November through February where we expect near-zero densities. We present selected model fits that best illustrate specific characteristics of each modeling approach. Full model fits for all locations are provided in Appendix: Section S3.

3 Linear Regression Modeling

We start with describing a relatively simple approach to model sparse time series data using Linear Regression (LR) Models. The models can be expressed as $Y' = X\beta + \epsilon$ where $\epsilon_i \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \sigma^2\right)$. Here, Y' is the vector of responses (transformed tick density) of size n. The matrix X has dimensions $n \times (p+1)$, where n is the number of samples and $(p+1)$ represents p predictors with one column (of ones) corresponding to the model intercept. The parameter vector $\beta = (\beta_0, \beta_1, \ldots \beta_p)$ includes the model intercept and p slopes. The random errors, ϵ_i are assumed to be independent and identically distributed (i.i.d.) following a normal distribution with mean zero and constant variance σ^2.

Linear Regression models offer a simple alternative to time series modeling when the data is irregularly sampled. These models can incorporate predictors such as time, temperature, and relative humidity along with periodicity. In this study, we explore two variants: *LR-Time* and *LR-Temp*. Each model is fit separately at each location using only the data observed at that location. Thus, for each variant, we fit a total of nine models, one for each unique location. We fit the linear models using `stats` package in R (R Core Team, 2022).

3.1 Linear Regression with Time predictor (LR-Time)

We first begin with a simple time dependent model. We start with using iso-week as one of the predictors in the LR-Time model. We incorporate periodicity using sine and cosine functions to capture the seasonal trend. We can mathematically write the model as

$$Y' = \beta_0 + \beta_1 X_1^{\mathrm{LR}} + \beta_2 X_2^{\mathrm{LR}} + \beta_3 X_3^{\mathrm{LR}} + \epsilon \tag{2}$$

$$X_1^{\mathrm{LR}} = 1, 2, \ldots, 53 \text{ i.e. the iso-week number}; \quad {X_2}^{\mathrm{LR}} = \frac{\sin(2\pi {X_1}^{\mathrm{LR}})}{53}; \quad {X_3}^{\mathrm{LR}} = \frac{\cos(2\pi {X_1}^{\mathrm{LR}})}{53} \tag{3}$$

We visualize one of the periodicity predictor i.e. the sinusoid (${X_2}^{\mathrm{LR}}$) in Fig. 3. The period we chose corresponds to the number of iso-weeks in a year i.e., 53. We observe in Fig. 3a that the sine wave completes two cycles but we suspect that it is not in phase with tick density. We use both sine and cosine predictors so the linear regression model can adjust the phasing.

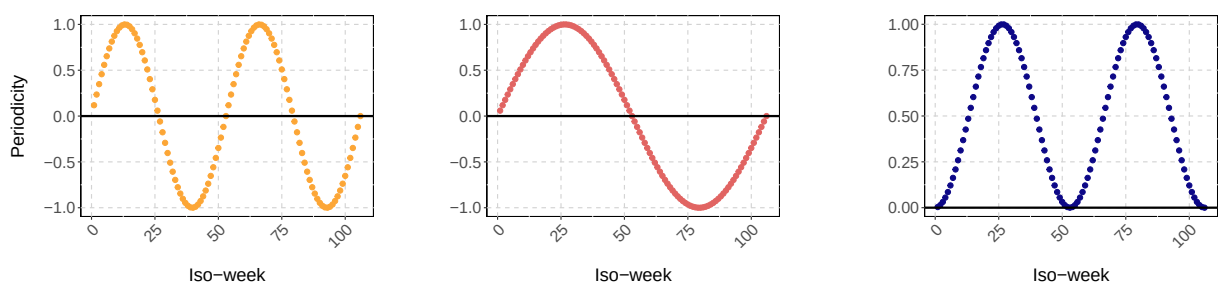

(a) sine term, 53-day period. **(b)** sine term, 106-day period. **(c)** Squared sine term, 106-day period.

Figure 3

The LR-Time model does not have predictors that vary across years and therefore it yields the same predictions year after year. Thus, we can look at the mean prediction as a function of iso-week. We visualize this model at two locations i.e., UKFS and OSBS in Fig. 4 to demonstrate the differences based on the model fits at each location. At UKFS (18 training points), the model captures the expected seasonal pattern with tick density increasing in summer and decreasing in winter but shows discontinuity across years and overly conservative upper bounds (Appendix: Section S3).

At OSBS (21 training points), the model fails to capture the expected seasonal pattern, and similarly exhibits discontinuity and overly conservative bounds. We observe that the mode of the tick density does not necessarily coincide with mid-year. We hypothesize that the peak tick density may be changing from one season to another. This makes it impossible to identify the real trend across seasons using only iso-week and periodicity as predictors.

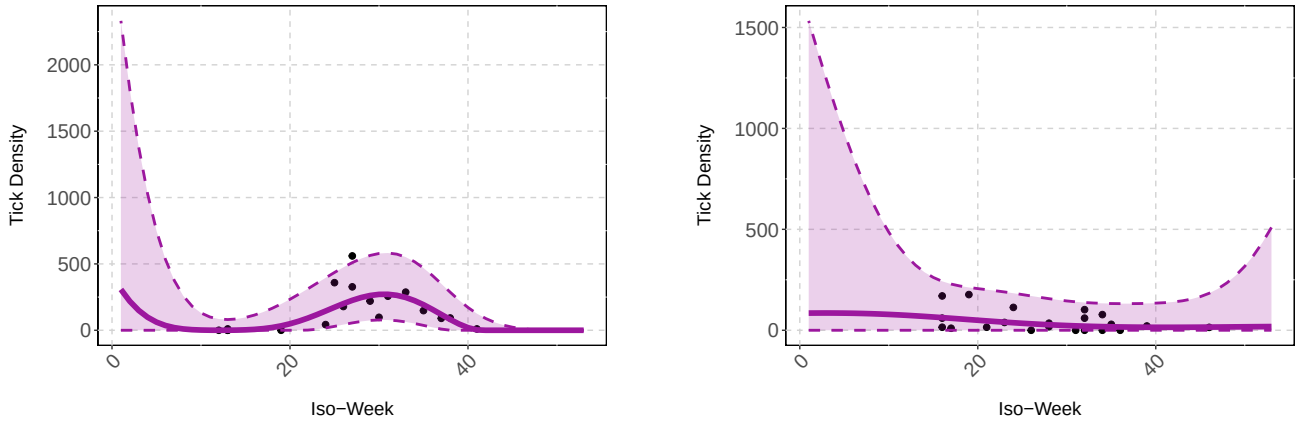

(a) LR-Time fit at UKFS

(b) LR-Time fit at OSBS

Figure 4

3.2 Linear Regression with Temperature Predictor (LR-Temp)

Empirical studies show that tick density depends on temperature and humidity (Vail and Smith, 1998; Sagurova et al., 2019). Including environmental factors like temperature may help capture tick density patterns (Raghavan et al., 2019). The NOAA Global Ensemble Forecasting System (Hamill et al., 2022)[1] provides 35-day forecasts for seven meteorological variables, including temperature, humidity, and precipitation—all known to influence tick density (Rochlin et al., 2023; Springer et al., 2015). Although humidity and precipitation are important (Clarke-Crespo et al., 2020; Raghavan et al., 2016), they are difficult to measure and are highly correlated with temperature. Thus, using temperature-driven model based on NOAA forecasts may be sufficient. However, as weather is forecasted into the future, the uncertainty in the forecast increases and will affect model performance.

Instead, we attempt to build a model using climatology rather than weather forecasts from NOAA. Climatology is the study of long-term weather patterns and averages whereas weather forecasts predict short-term atmospheric conditions. Since climatology averages over weather patterns, it smooths out the noise and yields a more stable forecast which can be used for long-term modelling. We considered minimum temperature (tmin) from climatology data as the predictor for this study and use a LR Model. The model is the same as in Eq. (2) replacing X_1^{LR} to $X_1^{\text{LR-Temp}}$ which represents tmin observed in the iso-week corresponding to when the response was recorded. We use periodicity as described in Eq. (3).

To incorporate tmin as a predictor, we make use of the Historic monthly weather data obtained from CRU-TS 4.09 (Harris et al., 2020) downscaled with WorldClim 2.1 (Fick and Hijmans, 2017). This dataset provides historic averages for minimum temperature, maximum temperature and precipitation from 1950 upto the last complete year, i.e. 2024. We used data from the years 2014-2024 at 5 mins resolution to obtain tmin (in Celsius). These data were recorded at a monthly resolution while the ticks samples were collected and reported by iso-week. We thus converted tmin from monthly to weekly resolution by repeating the same temperature for all weeks in the same month.

We show an example fit of the LR-Temp model obtained at location ORNL in Fig. 5. This model performs well on the training data. For example, it is able to capture uncertainty and predict zero tick density during colder months. However, for true forecasts into the future, temperature is unknown and

[1]See `https://www.ncei.noaa.gov/products/weather-climate-models/global-ensemble-forecast`.

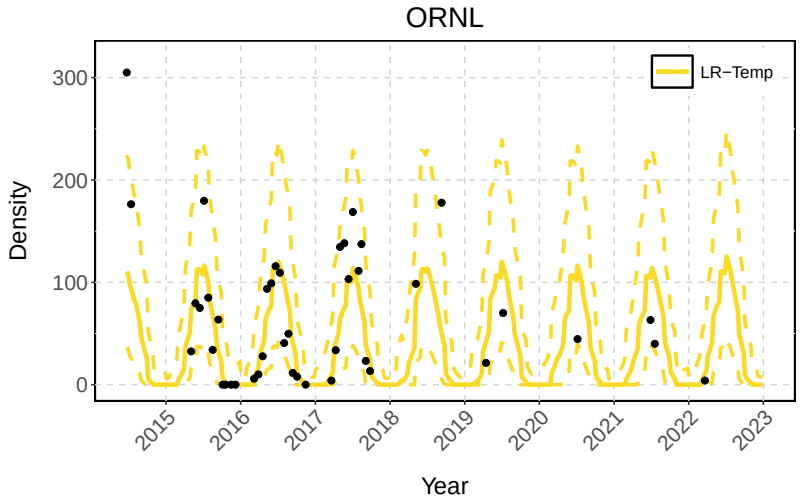

Figure 5

must be forecasted. Forecasting climate variables such as monthly average temperatures even one year in advance is inherently challenging due to short-term variability and uncertainty. We can potentially strengthen short-term forecasts, by eliminating environmental factors which require forecasting. Further, we can train one common model across all locations to capture broader geographical patterns, improve forecasting ability and enhance its generalizability.

4 Gaussian Processes

Homoskedastic Gaussian Processes (GPs; Gramacy, 2020) are highly flexible and non-parametric and can be used to learn patterns within data. GP regression models rely primarily on the distance between input values rather than their magnitude, which has made them widely used in spatial statistics (where they are known as kriging; Cressie, 1991). GPs have also gained popularity in machine learning due to their ability to provide fast predictions at new locations along with reliable uncertainty quantification (Rasmussen and Williams, 2005). We apply GPs in our ecological context because they do not require equally spaced time series, are flexible and particularly effective at distinguishing signal from noise (Johnson et al., 2018).

A GP model assumes that the response Y has a Multivariate Normal Distribution (MVN) centered at mean μ_n, with covariance structure Σ_n. The covariance between two responses depends on distances between the corresponding inputs. Simply stated, if training points x and x' are closer together in the input space, the GP assumes responses $y(x)$ and $y(x')$ are more highly correlated. A GP model can be written as

$$Y_n \sim \mathcal{N}\left(\mu_n, \Sigma_n\right); \quad \Sigma_n = \tau^2 \left(\sum_{k=1}^{d} C_\theta(X_n) + g\mathbb{I}_n \right) \tag{4}$$

where we specify $\mu_n = 0$ i.e., zero mean (for simplicity) and Σ_n is the covariance matrix for this process corresponding to input locations X_n. The covariance matrix is typically built using a kernel $C_\theta(X_n)$ with three hyperparameters. The scale, τ^2, captures the amplitude of the response; the length-scale, θ, captures the rate of decay of correlation; and the nugget, g, captures the noise observed in the response. For our

analysis, we used the distance based squared exponential kernel given by,

$$C_\theta(X_n)^{ij} = C(x_i, x_j) = \exp\left(-\sum_{k=1}^{d} \frac{(x_{ik} - x_{jk})^2}{\theta_k}\right) \quad \text{for } i, j = 1, 2, \ldots, n \tag{5}$$

This specification is used to describe an anisotropic GP where we have different length-scale in each input dimension so $\boldsymbol{\theta} = \{\theta_1, \theta_2, \ldots, \theta_d\}$, where d is the dimensionality of the input space. We infer hyper-parameters $\{\tau^2, \boldsymbol{\theta}, g\}$ through Maximum Likelihood Estimation (MLE).

We can then obtain the predictive distribution of $\mathcal{Y}$ at new locations $\mathcal{X}$ by calculating the conditional distribution of $\mathcal{Y}$, given (X_n, Y_n). To do so, we first stack the distribution of $\mathcal{Y}$ and Y_n and jointly write it as,

$$\begin{bmatrix} \mathcal{Y} \\ Y_n \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma(\mathcal{X}, \mathcal{X}) & \Sigma(\mathcal{X}, X_n) \\ \Sigma(X_n, \mathcal{X}) & \Sigma_n \end{bmatrix} \right) \tag{6}$$

where $\Sigma(\mathcal{X}, \mathcal{X})$ is the covariance matrix for $\mathcal{Y}$ and $\Sigma(X_n, \mathcal{X})$ is the covariance matrix between $\mathcal{Y}$ and Y_n. Then, using MVN conditioning properties, we have $\mathcal{Y}|Y_n, X_n \sim N\left(\mu(\mathcal{X}), \sigma^2(\mathcal{X})\right)$, and

$$\mu(\mathcal{X}) = \Sigma(\mathcal{X}, X_n)\Sigma_n^{-1} Y_n \quad \text{and} \quad \sigma^2(\mathcal{X}) = \Sigma(\mathcal{X}, \mathcal{X}) - \Sigma(\mathcal{X}, X_n)\Sigma_n^{-1}\Sigma(X_n, \mathcal{X}) \tag{7}$$

where $\mu(\mathcal{X})$ is the mean and $\sigma^2(\mathcal{X})$ is the variance at the predictive locations.

In order to use a Gaussian Process (GP) model effectively, we want to build predictors that capture "closeness" in our system. Further, we want predictors that do not require their own forecasts. Additionally, we aim to train one model on the all the training data at all locations (unified dataset) and design predictors which can capture both, temporal and geographical patterns. Temporal patterns can be represented using variables such as week number and periodic components, which we call *common predictors* and elaborate on in the following paragraph. Rather than using categorical predictors to capture geographical patterns, we introduce proxy variables that serve as locations identifiers. These are numeric inputs that distinguish locations while still providing useful information about tick density. We will refer to as *location-specific predictors*.

The *common predictors* include week number (X_1^{GP}) and periodicity (X_2^{GP}), which do not depend on location. The week number, (X_1^{GP}) is a sequence, $\{1, 2, \ldots, n, n+1, \ldots, n_p\}$, which represents grid of time points from the start date (week 1) up to the end date (week n_p) in the entire dataset, including the testing period. Week number n corresponds to the last week in the training data set. By using week number as a predictor we create a relationship between the previous week and the current week since the distance between consecutive weeks will be the smallest. We incorporate information that weeks across years should be similar through periodicity (X_2^{GP}). This predictor is based on the sine predictor, X_2^{LR} with some modifications. We show X_2^{GP} in Fig. 3c which is created by increasing the frequency from 53 to 106 weeks (Fig. 3b) and squaring the sinusoid so we have $X_2^{\text{GP}} = \sin^2(2\pi X_1/106)$. With this adjustment, the curve completes half a cycle in one year and flattens out towards the end of the year allowing for smoother transition between years.

In contrast to our common predictors, *location specific predictors* are designed to distinguish one location from another. The intuition behind these predictors is that if observations y and y' from our unified dataset are collected at the similar locations, the distances would be small resulting in higher correlation. We obtained our first location-specific predictor directly from the information provided by NEON (National Ecological Observatory Network (NEON), 2025). We incorporated mean elevation (X_3^{GP})

of location which is correlated with maximum tick density (Fig. 6a).

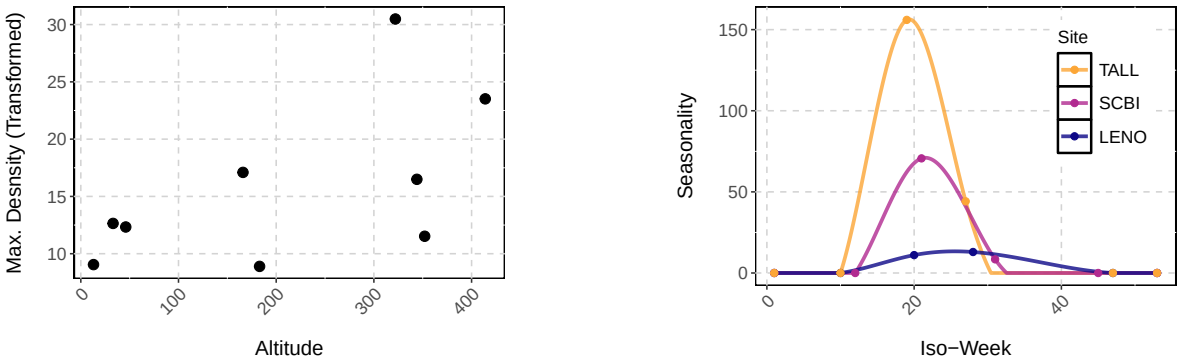

(a) Maximum tick density on transformed scale observed by mean elevation.

(b) Location-specific composite seasonality metric for three locations based on tick densities and seasonality.

Figure 6

Our last predictor is the location-specific seasonality predictor (X_4^{GP}) which captures seasonal patterns for each site using foliage data (Site Metadata, 2025), including the days-of-year (DOYs) when greenness increases, peaks, decreases, and reaches its minimum. These DOYs are converted to iso-weeks, X_0^{GP}, and the corresponding tick densities Y, observed at those weeks are paired with X_0^{GP}. To ensure continuity between years, we set tick density to zero for weeks 1 and 53, reflecting low winter activity. This gives us $(X_{0_i}^{\text{GP}}, Y_i)_{i=1}^{6}$ pairs, each representing an iso-week and its corresponding tick density. We then fit a cubic spline through these points as shown in Fig. 6b for three locations (see Appendix: Section S2 for one detailed location). To use this as a predictor, we evaluate the spline for all iso-weeks (1 to 53) and map the predictions to the corresponding location and week, creating X_4^{GP}.

This predictor is periodic and follows the same trend as foliage, i.e., it increases in the summer and decreases in the winter. Each kernel of the seasonality metric has a different height which allows us to incorporate some information about the usual tick density of one location relative to another, so if the height is less for LENO than for TALL, the tick density should be less on average at LENO relative to TALL. Another specialty of this predictor (unlike the sinusoid) is that the peak and the rate of change of the seasonality kernel across iso-weeks is allowed to vary from one location to another.

Using our four predictors explained above, we can fit one unified GP model to all the observed data while making location-specific predictions. In order to fit the GP, the input space is scaled so all predictors lie within the interval [0,1]. We used an anisotropic GP (Eqs. 4–5) and fit the models using `hetGP` (Binois and Gramacy, 2021).

We show the fit of the GP model at BLAN for the years 2015 upto 2022 (included) in Fig.7. The model produces a smooth fit with good uncertainty bounds and continuous predictions across years. It captures the expected seasonal pattern, i.e., higher tick abundance during summer and little to no activity in winter. As noted earlier, we train one GP model for the combined training dataset and this model assumes constant noise across the entire input space (all weeks and locations). As a result, the model produces wider prediction intervals at some locations, such as BLAN, where it effectively provides 100% coverage.

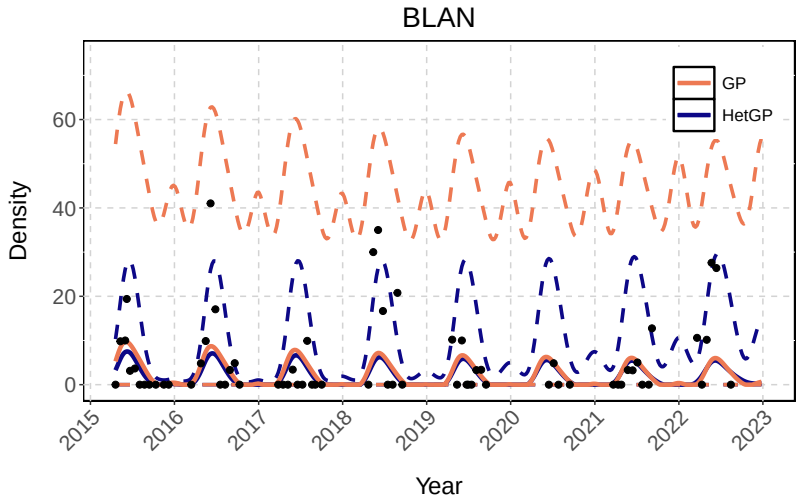

Figure 7

5 Heteroskedastic Gaussian Processes

We suspect that response variability changes substantially across locations and seasons (Fig. 2). Modeling heteroskedasticity will allow for tight bounds in regions of high signal and low noise, and broad bounds in regions with high noise. Thus, we can use a Heteroskedastic Gaussian Processes (HetGP; Binois et al., 2018) to improve UQ. A HetGP models the mean process similar to a GP but additionally infer a latent variance process for the noise. Specifically, the model assumes that the noise is input dependent,and changes across the input space. This is done by incorporating a vector of nuggets, $\boldsymbol{\lambda}_n$ (rather than the scalar nugget, g, in the GP setting). Mathematically, we can write the hetGP as,

$$Y_n \sim \mathcal{N}\left(0, \tau^2 K_n\right); \quad K_n = C_{\theta_Y}(X_n) + \Lambda_n \text{ and } \Lambda_n = \text{diag.}(\boldsymbol{\lambda}_n) \tag{8}$$

The response Y_n follows a MVN distribution similar to before but with a different covariance specification, $\tau^2 K_n$. Λ_n replaces the $g\mathbb{I}$ term in Eq. (4) and $C_{\theta_Y}(X_n)$ is the same as for a GP, following Eq. (5). The noise process, $\boldsymbol{\lambda}_n$, also follows a GP given by,

$$\log \boldsymbol{\lambda}_n \sim \mathcal{N}\left(0, \tau_\lambda^2 \Sigma_n\right); \quad \Sigma_n = C_{\theta_\lambda}(X_n) + g_\lambda \mathbb{I} \tag{9}$$

with covariance structure Σ_n with the same kernel structure as in Eq. (4) but a different set of hyper-parameters, θ_λ, τ_λ^2 and g_λ. All the hyper-parameters for mean process as well as the noise process are inferred jointly via MLE. The predictions follow Eqs. (6–7), with the new K_n matrix replacing the Σ_n matrix used previously. We use the same predictors as the GP and scale the inputs between [0, 1]. We fit the model using the `hetGP` library.

We start with observing the full time-series fit visualized in Fig. 7 to evaluate the performance of our HetGP model. Observe that the mean fits are similar for the GP and the HetGP model. However, we notice that the HetGP provides much tighter predictive bounds. The GP appears to be overly conservative at this location resulting in 100 % coverage whereas the HetGP provides ≈ 90% coverage. A full comparison of percentage coverage along with interval widths is provided in Section 6.1 which shows that the HetGP

outperforms all models with appropriate bounds.

To understand the difference in coverage between the GP and the hetGP models, we examine the prediction intervals, specifically the predictive variance obtained for a new location. The predictive variance is a combination of the process uncertainty and noise. Process uncertainty reflects variability in the response. The noise is unexplained variation which we suspect varies from one location to another and across seasons. A regular GP models the noise level using one parameter g which results in averaging the noise across all locations. The model thus underestimates noise at high-noise locations whilst over estimating noise at low noise locations such as BLAN (Fig. 7). In contrast, the HetGP models a separate noise process as described in Eqs. (8–9) and fits a different noise level across every week at every location. This results in adaptive prediction intervals that effectively capture input dependent noise and improve uncertainty estimation.

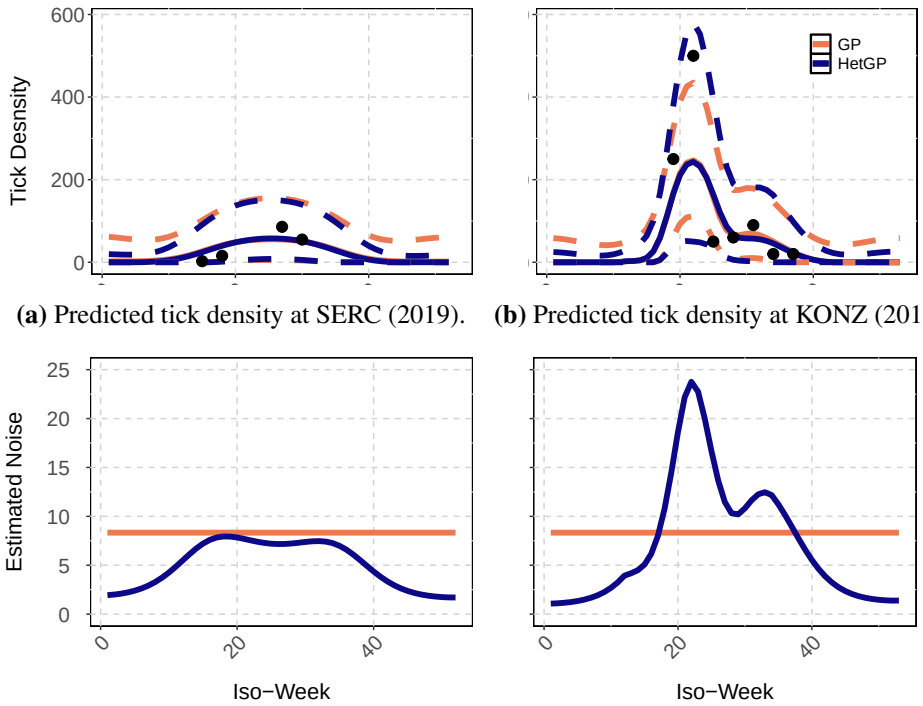

(a) Predicted tick density at SERC (2019). **(b)** Predicted tick density at KONZ (2018).

(c) Estimated noise level for SERC (2019). **(d)** Estimated noise level for KONZ (2018).

Figure 8

We demonstrate the flexibility of the HetGP model's UQ in comparison to the GP in Fig. 8. At location SERC, tick density ranges from 0 to 100 ticks/plot depending on the time of year. Both models produce similar mean predictions and achieve 100% coverage for 2019 (Fig. 8a). However, the GP's prediction intervals are slightly wider at the start and end of the year due to its constant noise assumption ($g \approx 10$), which averages noise across all training data (Fig. 8b). In contrast, HetGP estimates input-dependent noise, allowing it to capture changes in noise levels more effectively. The noise levels for the hetGP model $\boldsymbol{\lambda}_{n'}$ (for iso-weeks $n' = 1, \ldots, 53$) are consistently lower than g, resulting in tighter intervals.

In contrast, at KONZ, tick density ranges from 0 to 500 ticks/plot (Fig. 8c). The GP underestimates this range, with upper bounds around 400, as it assumes constant noise level at all locations. Its intervals are too narrow, often missing extreme values. HetGP, however, captures both seasonal and site-specific variation. During high-variability summer weeks ($n' \in 18, \ldots, 37$), it estimates higher noise ($\boldsymbol{\lambda} n' > g$),

leading to wider intervals. In winter, noise is lower ($\lambda'_n < g$ for $n' \notin \{18, 17, \ldots, 37\}$), resulting in tighter bounds (Fig. 8d). This adaptability enables HetGP to provide better coverage and more reliable forecasts.

6 Comparative Prediction Performance

We have gauged how all models perform visually and captured some pitfalls in the LR models. We improved upon those using the GP models by fitting a unified model across data from all locations capturing spatial variation. However, we are yet to quantitatively assess model performance and, benchmark our GP models against the LR models. In this section, we will look at metrics that can determine how well each model has performed on the entire training set at every single location (*within sample*) as well as how well it predicts at our hold out dates across different locations (*forecasts*).

6.1 Within sample

We evaluate within-sample fits using data before 2023 by analyzing Q-Q plots of standardized residuals (Fig. 9). Residuals are computed using training data and model predictions, then standardized by dividing by the predicted standard error. The Q-Q plot compares predicted vs. observed tick densities, and closeness to the Q-Q line indicates better fit and uncertainty quantification. Both linear models perform poorly, with points far from the line, indicating overly conservative bounds and high standard errors. The GP and HetGP models show slight skewness but align more closely with the Q-Q line, especially at lower quantiles. On the left, both models underestimate low tick densities, with GP doing so more severely. Overall, HetGP provides better fit and uncertainty estimation than GP.

We assess sample coverage to evaluate whether prediction intervals provide adequate UQ. However, coverage does not provide a complete picture by itself. It is necessary to asses the width of the bounds as well. We calculate these metrics based on the transformed predictive bounds to capture the full width of the intervals (back transforming results in lower PIs bounded at zero). We show the coverage of all models along with the mean and median width of the interval across all locations, for all observations in Tab. 1a. The coverage of all the models is between 93 - 94% for a 90% prediction interval, with LR-Temp being the most conservative where as the GP appears to be the closest to the actual level of confidence. However, HetGP has the smallest average width and median width for the interval while providing higher coverage than the GP model which is desirable while making reliable forecasts.

(a) Within Sample

Model	Coverage	Mean Width	Median Width
HetGP	93.3	8.37	7.54
GP	93.0	10.0	9.92
LR-Time	93.6	10.4	9.54
LR-Temp	93.9	10.5	9.84

(b) Out of sample

Model	RMSE	CRPS
HetGP	3.4765	2.1663
GP	3.4657	2.2113
LR-Time	3.6789	2.6967
LR-Temp	3.7437	2.7008

Table 1

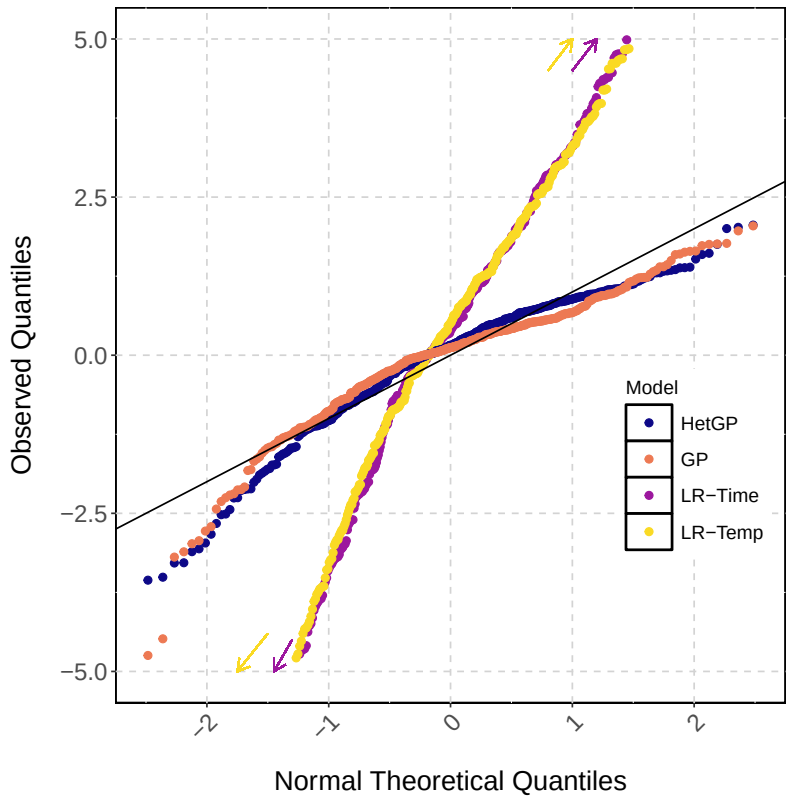

Figure 9

6.2 Out of sample

Our testing set consisted of all the data observed from 2023 onwards, which we use to evaluate the predictions. The metrics we used to compare the models included RMSE and CRPS as explained in Section 2.2.2. We show the calculated RMSE and CRPS for the testing data set in Tab. 1b. The RMSEs for both the GP models are very similar as they provide similar mean predictions. However, CRPS for the HetGP has the smallest values indicating superior UQ than the regular GP. The linear models have much larger CRPS and RMSE values indicating overall poor mean prediction and UQ.

7 Discussion

In this paper, we explored the use of Gaussian Processes to forecast tick populations for one future season (here one year). We show that the GP and HetGP models overall out perform LR models using several metrics such as coverage, RMSE and CRPS. We had three requirements in mind while building both GP models. First, we ensured that the predictors used in the models did not themselves require forecasting. To achieve this, we identified alternative predictors that did not rely on forecasting, replacing common weather-based variables such as temperature and humidity. By selecting predictors that could be constructed directly, we mitigated the issue of uncertainty propagation, and improved model UQ. Second, we aimed to fit a single GP model across all training data, incorporating both seasonal patterns and location-based variations. This ensured that we had a sufficiently large dataset to capture the trends as well as estimate

the underlying variability in each location. Lastly, we leveraged a core property of GP models: reliance on relative distances between predictor/input values rather than their absolute magnitudes.

Despite these advantages, GP models retain some of the challenges typical of more common time-series approaches. For example, GP models are best suited for short-term forecasting (e.g., one season/year). Similarly to most other time series approaches, we must assume that the process is stationary, which may not be the case over longer time horizons. Traditional time series models also assume weak stationarity. However, time series models can incorporate autoregressive (AR) and moving average (MA) terms or explicit temporal trends which can allow the model to extrapolate trends over longer periods.The trade off is that these time series methods require regularly sampled data, which is uncommon in ecological applications. Although it is possible to manipulate datasets in order to apply these methods, for example by aggregating data to more coarse time steps or by interpolating, for cases such as ours where whole seasons may be skipped, these options would reduce signal substantially. Consequently, our GP model makes what we think is a reasonable a trade-off: we are able to use sparse and irregularly sampled data, but we must focus on short to medium term forecasts.

Although the flexible, distance-based approach taken by by GPs is part of what enables it to capture patterns efficiently, this approach means that factors included in the model are not interpretable in the same way as they would be in LR. GPs also do not include mechanistic components that describe the dynamics of the population being modeled. This constrains how they may be used. For example, differential equation-based models can describe population changes over time in terms of demographic changed (births and deaths) as a function of factors such as temperature and humidity (Raghavan et al., 2016). This allows these, or other mechanistic models, to explore a range of scenarios about how long term changes in climate may impact in populations (Lou and Wu, 2017; Norman et al., 2016; Ostfeld and Brunner, 2015). In contrast, because the GP models do not posit either mechanisms or causality, they cannot be used to explore possible changes in populations across time due to climactic shifts.

We note that the current HetGP implementation lacks flexibility in specifying how noise varies across input dimensions. In ecological datasets, noise levels often fluctuate due to seasonal and geographical factors. In this paper, the HetGP model fits an individual noise level at every point in time across all locations. While this worked well in our case, it may lead to overfitting in scenarios where noise variation is driven by some inputs/factors that vary less across time, such as if noise only varies as we move from one location to another but does not depend on time/seasons. Future work could explore alternative formulations and implementations where the noise process only changes in some dimensions of the input space.

The primary goal of this study was to develop a short-term forecasting model for a nymphal lone-star tick density across multiple locations in the United States. Although this seems like a highly specific choice, our GP approach could be applied more widely, for example to other life stages or species of ticks, especially since the constructed predictors are fairly general. Studying short-term patterns of tick density enables health departments and policy makers to take action and ensure disease prevention (Kessler et al., 2019; Goddard and Varela-Stokes, 2009). Outdoor workers, hikers, etc can avoid locations where higher ticks are forecasted. Ticks also pose risks to animals. We can protect wildlife and livestock by taking precautionary measures when high tick density is expected (Springer et al., 2015). Similarly pet infections due to tick-borne diseases can be reduced by better understanding tick patterns (Sykes, 2023). Thus, GP models such as those presented here that are relatively simple to setup and can provide reliable short-term forecasts, could support a range of policy decisions.

Beyond tick forecasting, our GP-based methodology has broader applications in short-term ecological modeling. Many ecological forecasting problems require robust short-term predictions but lack dense,

regularly sampled datasets. Our approach avoids weather-dependent predictors and instead relies on carefully chosen predictors based on tick knowledge (Springer et al., 2015). This same recipe can be extended to other ecological applications, such as mosquito population forecasting or endangered species modeling. In cases where weather-based models struggle due to uncertainty propagation, GPs provide a viable alternative with reliable uncertainty quantification. Additionally, GP models are useful when dealing with large datasets that require short-term forecasts with minimal assumptions about long-term trends. Overall, our study highlights the effectiveness of GP methods in capturing ecological trends while minimizing forecasting uncertainty.

References

Allan, B. F., Goessling, L. S., Storch, G. A., and Thach, R. E. (2010). "Blood meal analysis to identify reservoir hosts for *Amblyomma americanum* ticks." *Emerging Infectious Diseases*, 16, 3, 433.

Barry, R. G. (2008). *Mountain weather and climate*. Cambridge University Press.

Binois, M. and Gramacy, R. B. (2021). "`hetgp`: Heteroskedastic Gaussian process modeling and sequential design in R." *Journal of Statistical Software*, 98, 1–44.

Binois, M., Gramacy, R. B., and Ludkovski, M. (2018). "Practical heteroscedastic gaussian process modeling for large simulation experiments." *Journal of Computational and Graphical Statistics*, 27, 4, 808–821.

Brugger, K., Walter, M., Chitimia-Dobler, L., Dobler, G., and Rubel, F. (2018). "Forecasting next season's *Ixodes ricinus* nymphal density: the example of southern Germany 2018." *Experimental and Applied Acarology*, 75, 281–288.

Carey, C. C., Woelmer, W. M., Lofton, M. E., Figueiredo, R. J., Bookout, B. J., Corrigan, R. S., Daneshmand, V., Hounshell, A. G., Howard, D. W., Lewis, A. S., et al. (2022). "Advancing lake and reservoir water quality management with near-term, iterative ecological forecasting." *Inland Waters*, 12, 1, 107–120.

Childs, J. E. and Paddock, C. D. (2003). "The ascendancy of *Amblyomma americanum* as a vector of pathogens affecting humans in the United States." *Annual Review of Entomology*, 48, 1, 307–337.

Clarke-Crespo, E., Moreno-Arzate, C. N., and López-González, C. A. (2020). "Ecological niche models of four hard tick genera (Ixodidae) in Mexico." *Animals*, 10, 4, 649.

Commins, S. P., James, H. R., Kelly, L. A., Pochan, S. L., Workman, L. J., Perzanowski, M. S., Kocan, K. M., Fahy, J. V., Nganga, L. W., Ronmark, E., Cooper, P. J., and Platts-Mills, T. A. (2011). "The relevance of tick bites to the production of IgE antibodies to the mammalian oligosaccharide galactose." *Journal of Allergy and Clinical Immunology*, 127, 5, 1286–1293.e6.

Cortinas, R. and Spomer, S. (2013). "Lone Star Tick (Acari: Ixodidae) Occurrence in Nebraska: Historical and Current Perspectives." *Journal of Medical Entomology*, 50, 2, 244–251.

Cressie, N. (1991). *Statistics for Spatial Data*. A Wiley-Interscience Publication. J. Wiley.

Estrada-Peña, A. (2001a). "Forecasting habitat suitability for ticks and prevention of tick-borne diseases." *Veterinary Parasitology*, 98, 1-3, 111–132.

— (2001b). "Forecasting habitat suitability for ticks and prevention of tick-borne diseases." *Veterinary Parasitology*, 98, 1-3, 111–132.

Falco, R. C. and Fish, D. (1992). "A comparison of methods for sampling the deer tick, *Ixodes dammini*, in a Lyme disease endemic area." *Experimental & Applied Acarology*, 14, 165–173.

Felz, M. W., Durden, L. A., and Oliver Jr, J. H. (1996). "Ticks parasitizing humans in Georgia and South Carolina." *The Journal of Parasitology*, 505–508.

Fick, S. E. and Hijmans, R. J. (2017). "WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas." *International Journal of Climatology*, 37, 12, 4302–4315.

Fisher, E. J., Mo, J., and Lucky, A. W. (2006). "Multiple pruritic papules from lone star tick larvae bites." *Archives of Dermatology*, 142, 4, 491–494.

Gaines, D. N., Operario, D. J., Stroup, S., Stromdahl, E., Wright, C., Gaff, H., Broyhill, J., Smith, J., Norris, D. E., Henning, T., et al. (2014). "Ehrlichia and spotted fever group Rickettsiae surveillance in *Amblyomma americanum* in Virginia through use of a novel six-plex real-time PCR assay." *Vector-Borne and Zoonotic Diseases*, 14, 5, 307–316.

Gajewski, Z., Stevenson, L. A., Pike, D. A., Roznik, E. A., Alford, R. A., and Johnson, L. R. (2021). "Predicting the growth of the amphibian chytrid fungus in varying temperature environments." *Ecology and Evolution*, 11, 24, 17920–17931.

Gneiting, T. and Raftery, A. E. (2007). "Strictly proper scoring rules, prediction, and estimation." *Journal of the American Statistical Association*, 102, 477, 359–378.

Goddard, J. (2002). "A ten-year study of tick biting in Mississippi: implications for human disease transmission." *Journal of Agromedicine*, 8, 2, 25–32.

Goddard, J. and Varela-Stokes, A. S. (2009). "Role of the lone star tick, *Amblyomma americanum* (L.), in human and animal diseases." *Veterinary Parasitology*, 160, 1, 1–12.

Gramacy, R. B. (2020). *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC.

Hahn, M. B., Jarnevich, C. S., Monaghan, A. J., and Eisen, R. J. (2016). "Modeling the geographic distribution of *Ixodes scapularis* and *Ixodes pacificus* (Acari: Ixodidae) in the contiguous United States." *Journal of Medical Entomology*, 53, 5, 1176–1191.

Haile, D. and Mount, G. (1987). "Computer simulation of population dynamics of the lone star tick, *Amblyomma americanum* (Acari: Ixodidae)." *Journal of Medical Entomology*, 24, 3, 356–369.

Hamill, T. M., Whitaker, J. S., Shlyaeva, A., Bates, G., Fredrick, S., Pegion, P., Sinsky, E., Zhu, Y., Tallapragada, V., Guan, H., et al. (2022). "The reanalysis for the global ensemble forecast system, version 12." *Monthly Weather Review*, 150, 1, 59–79.

Harris, I., Osborn, T. J., Jones, P., and Lister, D. (2020). "Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset." *Scientific Data*, 7, 1, 109.

Johnson, L. R., Gramacy, R. B., Cohen, J., Mordecai, E., Murdock, C., Rohr, J., Ryan, S. J., Stewart-Ibarra, A. M., and Weikel, D. (2018). "Phenomenological forecasting of disease incidence using heteroskedastic Gaussian processes: A dengue case study." *The Annals of Applied Statistics*, 12, 1, 27–66.

Kessler, W. H., Ganser, C., and Glass, G. E. (2019). "Modeling the distribution of medically important tick species in Florida." *Insects*, 10, 7, 190.

Kollars Jr., T. M., Oliver Jr., J. H., Durden, L. A., and Kollars, P. G. (2000). "Host Associations and Seasonal Activity of *Amblyomma americanum* (Acari: Ixodidae) in Missouri." *Journal of Parasitology*, 86, 5, 1156 – 1159.

Léger, E., Vourc'h, G., Vial, L., Chevillon, C., and McCoy, K. D. (2013). "Changing distributions of ticks: causes and consequences." *Experimental and Applied Acarology*, 59, 1, 219–244.

Levin, M. L., Zemtsova, G. E., Killmaster, L. F., Snellgrove, A., and Schumacher, L. B. (2017). "Vector competence of *Amblyomma americanum* (Acari: Ixodidae) for Rickettsia rickettsii." *Ticks and Tick-borne Diseases*, 8, 4, 615–622.

Little, R. J. and Rubin, D. B. (2019). *Statistical analysis with missing data*. John Wiley & Sons.

Liu, C., Zhai, Z., Zhang, R., Bai, J., and Zhang, M. (2022). "Field pest monitoring and forecasting system for pest control." *Frontiers in Plant Science*, 13, 990965.

Lou, Y. and Wu, J. (2017). "Modeling Lyme disease transmission." *Infectious Disease Modelling*, 2, 2, 229–243.

Masters, E. J., Grigery, C. N., and Masters, R. W. (2008). "STARI, or Masters Disease: Lone Star Tick–Vectored Lyme-like Illness." *Infectious Disease Clinics of North America*, 22, 2, 361–376. Tick-borne Diseases, Part I: Lyme Disease.

Meltzer, M. and Norval, R. (1992). "The use of time-series analysis to forecast bont tick (*Amblyomma hebraeum*) infestations in Zimbabwe." *Experimental & Applied Acarology*, 13, 4, 261–279.

Mitchell, C. L., Lin, F.-C., Vaughn, M., Apperson, C. S., Meshnick, S. R., and Commins, S. P. (2020). "Association between lone star tick bites and increased alpha-gal sensitization: evidence from a prospective cohort of outdoor workers." *Parasites & vectors*, 13, 1–4.

Mixson, T. R., Campbell, S. R., Gill, J. S., Ginsberg, H. S., Reichard, M. V., Schulze, T. L., and Dasch, G. A. (2006). "Prevalence of Ehrlichia, Borrelia, and Rickettsial agents in *Amblyomma americanum* (Acari: Ixodidae) collected from nine states." *Journal of Medical Entomology*, 43, 6, 1261–1268.

Mock, D. E., Applegate, R. D., and Fox, L. B. (2001). "Preliminary survey of ticks (Acari: Ixodidae) parasitizing wild turkeys (Aves: Phasianidae) in eastern Kansas." *Journal of Medical Entomology*, 38, 1, 118–121.

Molaei, G., Little, E. A., Williams, S. C., and Stafford, K. C. (2019). "Bracing for the worst—range expansion of the lone star tick in the northeastern United States." *New England Journal of Medicine*, 381, 23, 2189–2192.

Monzón, J. D., Atkinson, E. G., Henn, B. M., and Benach, J. L. (2016). "Population and evolutionary genomics of *Amblyomma americanum*, an expanding arthropod disease vector." *Genome Biology and Evolution*, 8, 5, 1351–1360.

National Ecological Observatory Network (NEON) (2024). "Ticks sampled using drag cloths (DP1.10093.001)."

— (2025). "Site management and event reporting (DP1.10111.001)."

Norman, R. A., Worton, A. J., and Gilbert, L. (2016). "Past and future perspectives on mathematical models of tick-borne pathogens." *Parasitology*, 143, 7, 850–859.

Ostfeld, R. S. and Brunner, J. L. (2015). "Climate change and Ixodes tick-borne diseases of humans." *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370, 1665, 20140051.

Paddock, C. and Yabsley, M. (2007). "Ecological havoc, the rise of white-tailed deer, and the emergence of *Amblyomma americanum*-associated zoonoses in the United States." *Wildlife and Emerging Zoonotic Diseases: The Biology, Circumstances and Consequences of Cross-Species Transmission*, 289–324.

Paddock, C. D., Lane, R. S., Staples, J. E., and Labruna, M. B. (2016). "Changing paradigms for tick-borne diseases in the Americas." In *Global health impacts of vector-borne diseases: workshop summary*, 221–258. National Academies Press Washington, DC, USA.

Pfäffle, M., Littwin, N., Muders, S. V., and Petney, T. N. (2013). "The ecology of tick-borne diseases." *International Journal for Parasitology*, 43, 12-13, 1059–1077.

Piesman, J. and Eisen, L. (2008). "Prevention of Tick-Borne Diseases*." *Annual review of entomology*, 53, 323–43.

Ponomarev, V., Napalkova, V., and Klobukov, G. (2019). "Forecasting the phenology of gypsy moth development and the time frames for the pest control based on sum of effective temperatures according to latitudinal origination of population." In *AIP Conference Proceedings*, vol. 2063. AIP Publishing.

R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Raghavan, R. K., Goodin, D. G., Hanzlicek, G. A., Zolnerowich, G., Dryden, M. W., Anderson, G. A., and Ganta, R. R. (2016). "Maximum entropy-based ecological niche model and bio-climatic determinants of lone star tick (*Amblyomma americanum*) niche." *Vector-Borne and Zoonotic Diseases*, 16, 3, 205–211.

Raghavan, R. K., Neises, D., Goodin, D. G., Andresen, D. A., and Ganta, R. R. (2014). "Bayesian Spatio-Temporal Analysis and Geospatial Risk Factors of Human Monocytic Ehrlichiosis." *PLOS ONE*, 9, 7, 1–8.

Raghavan, R. K., Peterson, A. T., Cobos, M. E., Ganta, R. R., and Foley, D. H. (2019). "Current and Future Distribution of the Lone Star Tick, *Amblyomma americanum* (L.) (Acari: Ixodidae) in North America." *PLoS ONE*, 14.

Randolph, S. and Rogers, e. D. (1997). "A generic population model for the African tick *Rhipicephalus appendiculatus*." *Parasitology*, 115, 3, 265–279.

Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning*. The MIT Press.

Rochlin, I., Egizi, A., and Ginsberg, H. S. (2023). "Modeling of historical and current distributions of lone star tick, *Amblyomma americanum* (Acari: Ixodidae), is consistent with ancestral range recovery." *Experimental and Applied Acarology*, 89, 1, 85–103.

Rochlin, I., Egizi, A., and Lindström, A. (2022). "The Original Scientific Description of the Lone Star Tick (*Amblyomma americanum*, Acari: Ixodidae) and Implications for the Species' Past and Future Geographic Distributions." *Journal of Medical Entomology*, 59, 2, 412–420.

Ryan, S. J., Carlson, C. J., Mordecai, E. A., and Johnson, L. R. (2019). "Global expansion and redistribution of Aedes-borne virus transmission risk with climate change." *PLoS neglected tropical diseases*, 13, 3, e0007213.

Sagurova, I., Ludwig, A., Ogden, N. H., Pelcat, Y., Dueymes, G., and Gachon, P. (2019). "Predicted northward expansion of the geographic range of the tick vector *Amblyomma americanum* in North America under future climate conditions." *Environmental Health Perspectives*, 127, 10, 107014.

Savage, H., Godsey, M., Lambert, A., Panella, N., Burkhalter, K., Harmon, J., Lash, R., Ashley, D., and Nicholson, W. (2013). "First Detection of Heartland Virus (Bunyaviridae: Phlebovirus) from Field Collected Arthropods." *The American Journal of Tropical Medicine and Hygiene*, 89.

Site Metadata (2025). "NEON Sites Metadata: `https://raw.githubusercontent.com/eco4cast/neon4cast-ci/main/catalog/sites/10.1002/fee.2616`."

Slatculescu, A. M., Clow, K. M., McKay, R., Talbot, B., Logan, J. J., Thickstun, C. R., Jardine, C. M., Ogden, N. H., Knudby, A. J., and Kulkarni, M. A. (2020). "Species distribution models for the eastern blacklegged tick, *Ixodes scapularis*, and the Lyme disease pathogen, *Borrelia burgdorferi*, in Ontario, Canada." *PLoS One*, 15, 9, e0238126.

Springer, Y. P., Eisen, L., Beati, L., James, A. M., and Eisen, R. J. (2014). "Spatial distribution of counties in the continental United States with records of occurrence of *Amblyomma americanum* (Ixodida: Ixodidae)." *Journal of Medical Entomology*, 51, 2, 342–351.

Springer, Y. P., Jarnevich, C. S., Barnett, D. T., Monaghan, A. J., and Eisen, R. J. (2015). "Modeling the present and future geographic distribution of the lone star tick, *Amblyomma americanum* (Ixodida: Ixodidae), in the continental United States." *The American Journal of Tropical Medicine and Hygiene*, 93, 4, 875.

Steinke, J. W., Platts-Mills, T. A., and Commins, S. P. (2015). "The alpha-gal story: Lessons learned from connecting the dots." *Journal of Allergy and Clinical Immunology*, 135, 3, 589–596.

Sykes, J. E. (2023). "Tick-borne diseases." *Veterinary Clinics: Small Animal Practice*, 53, 1, 141–154.

Thomas, R. Q., Boettiger, C., Carey, C. C., Dietze, M. C., Johnson, L. R., Kenney, M. A., McLachlan, J. S., Peters, J. A., Sokol, E. R., Weltzin, J. F., Willson, A., Woelmer, W. M., and contributors, C. (2023). "The NEON Ecological Forecasting Challenge." *Frontiers in Ecology and the Environment*, 21, 3, 112–113.

Troughton, D. R. and Levin, M. L. (2007). "Life Cycles of Seven Ixodid Tick Species (Acari: Ixodidae) Under Standardized Laboratory Conditions." *Journal of Medical Entomology*, 44, 5, 732–740.

Vail, S. G. and Smith, G. (1998). "Air temperature and relative humidity effects on behavioral activity of blacklegged tick (Acari: Ixodidae) nymphs in New Jersey." *Journal of Medical Entomology*, 35, 6, 1025–1028.

van Nunen, S. (2015). "Tick-induced allergies: mammalian meat allergy, tick anaphylaxis and their significance." *Asia Pacific Allergy*, 5, 1, 3–16.

Williams, H. W., Cross, D. E., Crump, H. L., Drost, C. J., and Thomas, C. J. (2015). "Climate suitability for European ticks: assessing species distribution models against null models and projection under AR5 climate." *Parasites & vectors*, 8, 1–15.

Wimberly, M. C., Yabsley, M. J., Baer, A. D., Dugan, V. G., and Davidson, W. R. (2008). "Spatial heterogeneity of climate and land-cover constraints on distributions of tick-borne pathogens." *Global Ecology and Biogeography*, 17, 2, 189–202.

Zannou, O. M., Ouedraogo, A. S., Biguezoton, A. S., Abatih, E., Coral-Almeida, M., Farougou, S., Yao, K. P., Lempereur, L., and Saegerman, C. (2021). "Models for studying the distribution of ticks and tick-borne diseases in animals: A systematic review and a meta-Analysis with a Focus on Africa." *Pathogens*, 10, 7, 893.