Modeling American Kestrel Decline Using Spatiotemporal Subsampling to Improve eBird Data Reliability

Archith Sharma[1] and Junhyeon Kwon*[2]

[1]Texas Academy of Mathematics and Science, University of North Texas
[2]Department of Mathematics, University of North Texas

October 31, 2024

Abstract

American Kestrel (Falco sparverius) numbers have plummeted over 40% since 1980, despite many other raptor species, including the closely related Merlin (Falco columbarius) rebounding in the same period, post the ban on the use of DDT. One plausible reason for this decline is the increase in Cooper's Hawks (Accipiter cooperii), a predatory species. To model the impact of Cooper's Hawk on American Kestrels, one can use citizen science data such as eBird. However, eBird contains a heavy sampling bias towards recent data in large cities as it has gained popularity over time. Spatiotemporal subsampling is performed to make the data more reliable and to enjoy a finer time resolution compared to the Breeding Bird Survey standard. The creation and analysis of synthetic data are also tested to verify that our framework reflects the ground truth of the bird abundance trend. The ecological relationship between two species is then modeled by regression with ARMA (Auto-Regressive Moving Average) errors. The optimal model based on AIC (Akaike Information Criterion) showed immediate forecasts of a stable population at the current pace, but with potential for a recovery to past levels.

*Corresponding author. Email: junhyeon.kwon@unt.edu

1 Introduction

The American Kestrel (*Falco sparverius*), seen in Figure 1 is North America's smallest and most abundant falcon. While Kestrels are about the same size as a mourning dove, they are known to be vicious predators. Kestrels are regular breeding inhabitants in most of North America. They prefer open, grassy habitat in search of their prey, mostly insects.

Figure 1: Female American Kestrel

In total, North America has lost an estimated 2 million Kestrels since 1970, despite the ban on DDT, according to the Peregrine Fund, a leading organization in raptor conservation. [1] Conservancy agencies and past research point to the decline being attributed to various potential factors, including predation by Cooper's Hawks (*Accipiter cooperii*), habitat change, and pesticide use. [2]

In this paper, we outline a methodology for considering covariates in a time-dependent model by considering the Cooper's Hawk theory - that predation is what has driven the numbers of Kestrels down. With vast amounts of data available, the ecological relationships between Kestrels and Cooper's Hawk can be explained through statistical relationships. Changes in the abundance can be better modeled if we consider the temporal dependence in the data.

Past research in bird abundance relationships has focused on the Northern Goshawk,

The time-series method of analysis used in this paper is regression with AutoRegressive Moving Average (ARMA) errors [3]. Regression with ARMA errors combines the potential explanatory variables with the time-dependent nature of Kestrel abundance, as the abundance of Kestrels this year can be influenced by the abundance last year. Standard linear regression models fail to account for autocorrelated errors, and adding ARMA to these errors can enhance model forecast accuracy and make it more interpretable.

We hypothesize that time-series methods used with eBird and BBS data can be used to find evidence of how the quick rebound of predatory species after the ban of DDT has continued to drive Kestrel numbers down. Since Cooper's Hawks do not eat primary consumers, the pesticides affecting Kestrels' food sources have forced them into foraging larger areas where they are vulnerable to predation.

1.1 Data Description

The data that is used in this research comes from eBird [2], a community science project where users can document bird sightings and the Breeding Bird Survey, an annual survey of birds during the breeding season (June). eBird is used as the main dataset in this project, along with Breeding Bird Survey (BBS) data to verify the quality of eBird data.

eBird eBird data is a public repository of bird observations maintained by the Cornell Lab of Ornithology. The dataset contains records of bird sightings reported by the general public. Birders enter where they birded (location), for how long, and what birds they saw. This data is uploaded to the eBird database, where automated systems and regional reviewers maintain data quality. eBird has grown tremendously over time, with over 1 billion observations recorded as of 2024.

For this research, the spatial domain of eBird data was restricted to the Atlantic Waterfowl Flyway in the Continental US, and the data for the time period between 1976-2022 in June is used. [4] The Atlantic Waterfowl Flyway consists of the Eastern United States, which is the region where Kestrels have seen the most noticeable decline in the last 40 years.[5] The time period 1976-2022 is used as the most accurate comparison to existing BBS data. Even though eBird's recordkeeping began in 2002, data prior to that exists in the eBird archives since many users have uploaded prior checklists before eBird was created. This extra data creates a historical context for current population trends. The full plot of eBird data (at different time intervals) is present in Figure 2.

BBS BBS data is a long-term dataset collected through annual surveys of bird populations across North America maintained by the USGS and the Canadian Wildlife Service. The survey consists of thousands of fixed routes over uniformly distributed regions across the United States and Canada as shown in Figure 3. Similar to the eBird dataset, BBS data was also restricted to the same spatiotemporal domain. [6] [7]

The paper is organized as follows. Section 2 introduces our methods for debiasing eBird data and estimating and forecasting abundance trends for the focal bird species. Section 3 describes the simulation study designed to verify that our framework accurately captures trend dynamics in the presence of spatiotemporal sampling bias in synthetic data. In Section 4, we apply this framework to real-world data collected in the Atlantic Waterfowl Flyway. Finally, Section 5 concludes with an interpretation of our model's forecast results, discussing the implications, contributions, and limitations of our approach, and potential directions for future research.

2 Data Debiasing and Trend Estimation Framework

The methods used in this section are denoted in Figure 4. eBird data is sampled according to different spatial buffers, and be tested against synthetic data and BBS data. If similar, regression with ARMA errors will be run and then based on the ability of the model to fit along with evaluation with rolling forecast origin, forecasting will be performed.

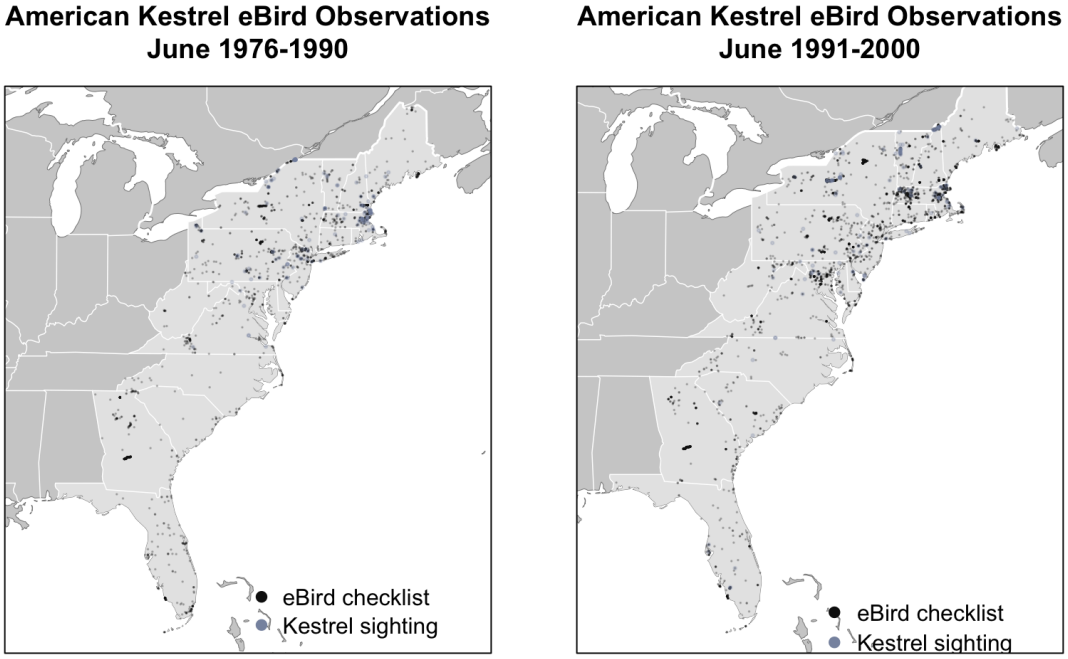

(a) 1976-1990

(b) 1991-200

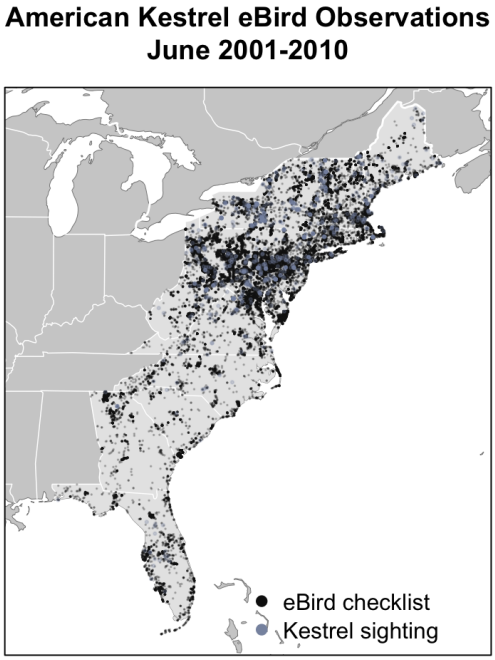

(c) 2001-2009

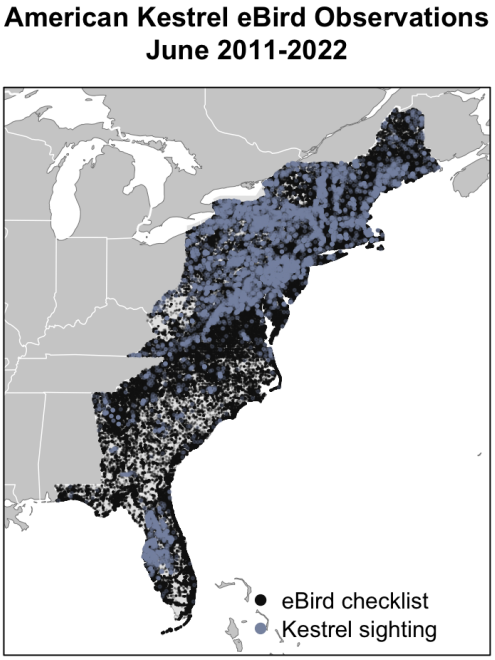

(d) 2011-2022

Figure 2: eBird exponential growth over time

2.1 Spatiotemporal Subsampling

Despite the measures in place to ensure eBird data quality, biases are present in eBird data, most notably spatial and temporal bias. Users sample near their homes or in accessible

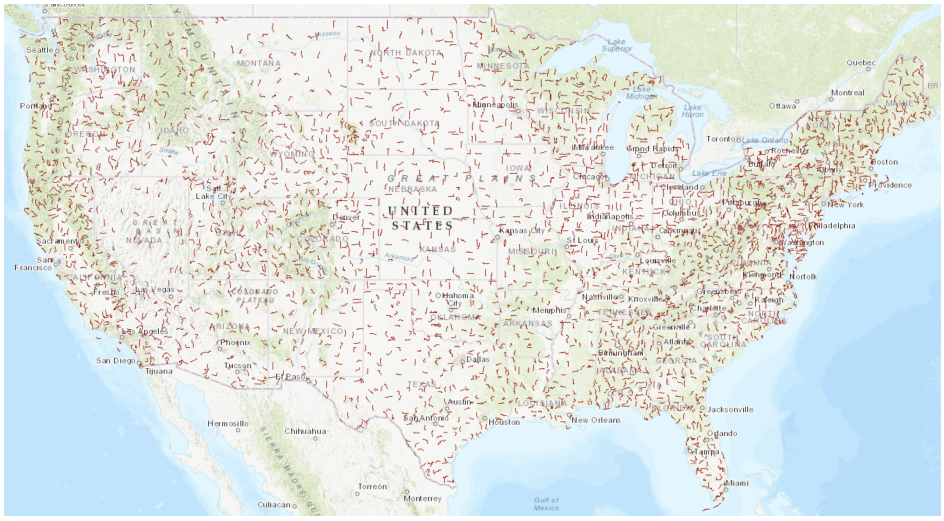

Figure 3: BBS Annual Survey Routes

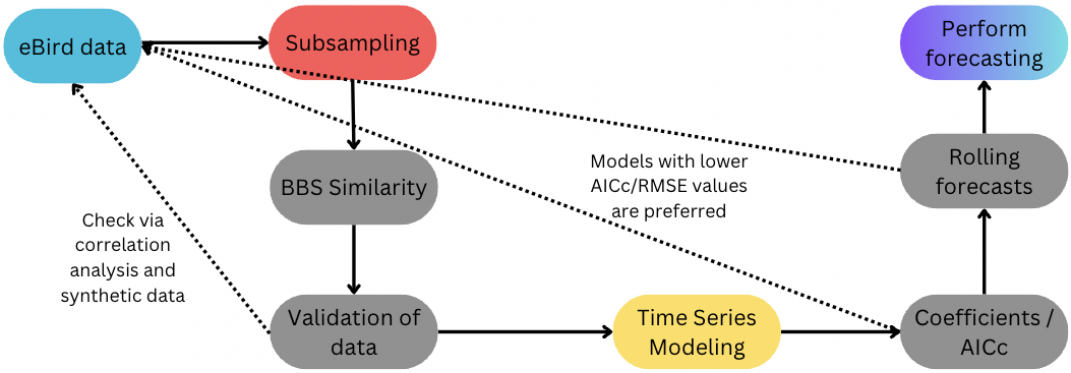

Figure 4: Methods for eBird time series

areas. [8] eBird has also had exponential growth in the number of users since its launch, leading to a strong bias towards recent data.

Both of these problems have fixes that can be implemented in R. Creating equal area grids over the study region and sampling a given amount per box for each year addresses both of these biases, as shown in Figure 5. This spatiotemporal subsampling is done at three different resolutions - 3 km, 10 km, and 30 km - to see what sample is most representative of the Kestrel population (i.e. similarity to the longer-standing BBS standard). [9]

2.2 Trend Estimation and Forecast

Regression with ARMA errors is fit to each of the sampled series. The model is

$$y_t = \beta_0 + \beta_1 x_t + \eta_t \tag{1}$$

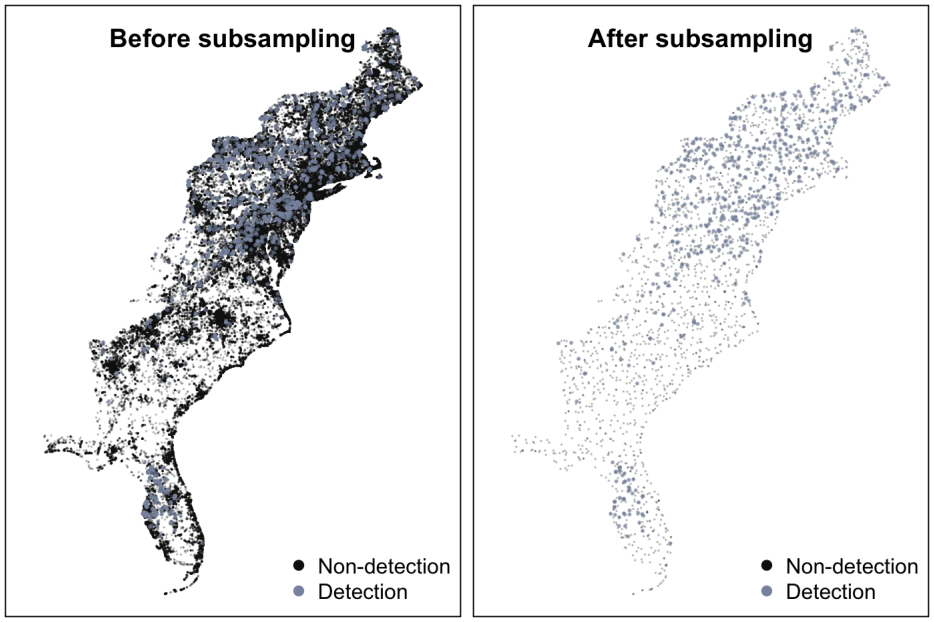

Figure 5: 30 km subsampled grid

where x_t is a covariate at time t and β_1 is its coefficient. β_0 is the intercept of the data, and η_t is the ARMA error, noted as

$$\eta_t = \phi_1 n_{t-1} + \cdots + \phi_p n_{t-p} - \theta_1 z_{t-1} - \cdots - \theta_q z_{t-q} + z_t \tag{2}$$

with order (p, q) where z_t is a zero-mean white noise process. The most accurate model is selected based on the corrected Akaike Information Criterion (AICc) based on the maximum likelihood estimation. Forecasts can then generated based on this model. AIC is defined as

$$\text{AIC} = 2k - 2\ln(\hat{L}) \tag{3}$$

where k is is the number of estimated parameters and $\hat{L}$ is the maximum likelihood of the model.

AICc is generally the better criterion. AICc is calculated as

$$\text{AICc} = \text{AIC} + \frac{2k(k+1)}{n-k-1} \tag{4}$$

Notice that the correction for small sample sizes $\frac{2k(k+1)}{n-k-1}$ goes to zero as $n \to \infty$. Thus, as the sample size gets larger, AICc converges to AIC. AICc is especially better for models in which the sample size n is small compared to the amount of parameters k. n is small compared to k in ARMA since for an ARMA(p, q) with an external regressor and intercept,

$$k = p + q + 2 \tag{5}$$

If $\frac{n}{k} < 40$, AICc should be used. [10] Since most models will have at least one order of p or q, $\frac{n}{k} < 40$ as k is at least two and $n = 47$ years.

After selecting the models with the lowest AICc values, a diagnostic check will be performed by analyzing the residuals for stationarity and lack of autocorrelation. The Ljung-Box test p-values will be used to assess if the residuals are uncorrelated, and a Normal Q-Q plot will be reviewed to confirm if the model captured the trend. [11]

To verify the model's ability at generalizing to unseen data, the process of rolling forecast origin in is used to verify forecast accuracy at different steps ahead based on Root Mean Squared Error (RMSE) to evaluate whether the time-dependent structure has better regression performance. [3] Then, multiple forecasts will be generated with different changes to the Cooper's regressor to determine which way Kestrel populations will trend.

3 Simulation Study

In this section, we present the details of our simulation study, including the creation of synthetic data to mimic both Breeding Bird Survey (BBS) and eBird data, the number of repetitions (over 200), and the confidence intervals of estimated parameters based on these repetitions. We discuss the implications of the findings from this study.

3.1 Simulation Setting

JK: This subsection describes the setup of our simulation study. Key elements include the length of synthetic data, the dimensions of the spatial grid (e.g., a 10×10 grid), and the number of repetitions (200+) we performed. These settings are chosen to replicate typical conditions in citizen science datasets like BBS and eBird. The simulation setup is done by creating a 10x10 grid with a length of 47 years. To simulate real-world collection of eBird data, we increase the amount of 'checklists' that we are sampling over time. The ground truth model is constructed based on a reasonable estimation for the intercept and slope of Kestrels and Cooper's Hawks. These trends are then tested through repeated simulation. This increasing checklist abundance estimates the amount of Kestrels found on checklists in a spatial region each year. This process is repeated 200 times to estimate the α_1 and β_1 parameters with an empirical confidence interval. .

3.2 Ground Truth Model

JK: Here, we outline the ground truth model assumptions for the synthetic abundance data of American Kestrels and Cooper's Hawks. This model serves as the baseline against which we assess our framework's performance in estimating abundance trends, simulating true population dynamics under predefined conditions.

This spatial randomness can be shown through a Gaussian process in a Gaussian Random Field (GRF). The GRF assigns a random value to an individual grid cell, but still follows a spatial dependence structure. The value which changes the intercept and slope is randomly determined by independent GRFs, denoted as $f_a(s)$ and demonstrated in Figure 6. The intercept and slope estimates, α_0^* and α_1^* are given reasonable values based on estimated change over time and rarity on checklists. At a point on the grid s, the GRF values can add

spatial-dependent randomness to both of these parameters mathematically as

$$\alpha_0(s) = \alpha_0^* + f_\alpha(s) \tag{6}$$

and the same for α_1^*. Assuming the model is the same as described in Equation 1, we can describe each of the abundances in a spatial grid at box s at time t as

$$x(s,t) = \alpha_0(s) + \alpha_1(s) \times t + \eta_1(s,t) \tag{7}$$

where η follows an ARMA(1, 1) error. We can then model x as the series representing Cooper's Hawks, and assume that α_1* is positive to represent the increasing trend.

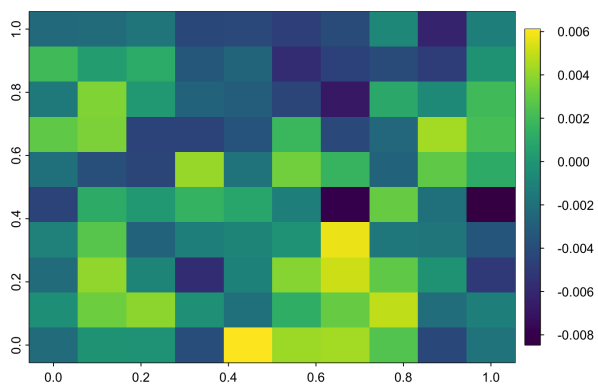

Figure 6: Gaussian Random Field on a 10x10 grid

β_0^* and β_1^* then represent the intercept and slope of the American Kestrel series. Following a similar process to the construction of the Cooper's Hawk series in Equation 7 with GRFs being added to the parameter estimation as in Equation 6. Instead of making time the independent factor, we make the Cooper's Hawk series the independent factor as in Equation 8.

$$y(s,t) = \beta_0(s) + \beta_1(s) \times x(s,t) + \eta_2(s,t) \tag{8}$$

To take the region consensus, we take the average of the parameters at each spatial grid for n gridboxes using the operations in Equation 9, 10, and 11.

$$\overline{\alpha_0} = \frac{1}{n}\sum_{i=1}^{n} \alpha_0(s_i) = \alpha_0^* + \sum_{i=1}^{n} f_{\alpha_0}(s_i) \tag{9}$$

$$\overline{x}(t) = \frac{1}{n}\sum_{i=1}^{n} x(s,t) \tag{10}$$

$$\overline{\eta_1}(t) = \frac{1}{n}\sum_{i=1}^{n} \eta_1(s,t) \tag{11}$$

Thus, over the entire spatial grid,

$$\overline{x}(t) = \overline{\alpha_0} + \overline{\alpha_1} \times t + \overline{\eta_1}(t) \tag{12}$$

and

$$\overline{y}(t) = \overline{\beta_0} + \overline{\beta_1} \times \overline{x}(t) + \overline{\eta_2}(t) \tag{13}$$

3.2.1 Simulated Ground Truth coefficients

Based on the simulation process, the estimates for the four model parameters over the grid are denoted in Table 1. The simulation process yielded coefficients of:

Variable	**Value**
Cooper intercept (α_0)	0.02
Cooper slope (α_1)	0.005
Kestrel intercept (β_0)	0.09
Kestrel coefficient against Cooper's Hawk (β_1)	-0.17

Table 1: Simulation Ground Truth

3.3 Distribution of Checklists

JK: This subsection explains the methods used to introduce sampling bias into the synthetic data, designed to reflect the uneven geographical and temporal distribution of eBird observations. We describe how this sampling bias is embedded in the synthetic data to simulate real-world data collection tendencies.

eBird data contains uneven geographic and temporal distributions. Some geographic areas have more checklists than others, and the amount of checklists have increased over time. We then create a similar process in the synthetic dataset. Over the spatial grid, every 10 cells receives an increase in checklists to 'sample' from. The amount of checklists is also increasing over the sample period. The detection frequency is applied in a Bernoulli trial with the detection frequency being selected from the ground truth.

3.4 Performance Measurement

JK: In this subsection, we present the results of the simulation, showing how our estimators perform in relation to the ground truth values across 200+ simulation repetitions. We provide empirical 95% confidence intervals to demonstrate the accuracy and reliability of our estimators in capturing the true abundance trends.

The coefficients of the simulated data are presented below in Table 2:

Variable	**Ground Truth value**	**Simulated 95% confidence interval**
Cooper intercept (α_0)	0.02	-
Cooper slope (α_1)	0.005	0.004 ± 0.000416
Kestrel intercept (β_0)	0.09	-
Kestrel coefficient against Cooper's Hawk (β_1)	-0.17	-0.15 ± 0.0263

Table 2: Simulation Results (95% Confidence Interval

4 Application on Atlantic Waterfowl Flyway

In this section, we apply our method to real eBird data collected in the Atlantic Waterfowl Flyway region. We optimize subsampling of the eBird data to approximate the BBS data

during the breeding season. This optimized subsampling approach is then extended across the monthly dataset of eBird. Finally, we fit a time series model to the subsampled data and perform a forecast.

4.1 Determining Subsampling Resolution

JK: This subsection outlines the process of determining the optimal subsampling resolution for eBird data to align as closely as possible with BBS data during the breeding season. We discuss the criteria used for optimality, the methodology for selecting the appropriate subsampling resolution, and the resulting configuration applied to the eBird dataset.

4.1.1 Subsampled Series

The first section of the results explores the subsampling processes and how the estimated abundance of Kestrels changes over time depending on the sample. Figure 7 shows the abundance at each point in time for checklists with different samples, along with BBS data.

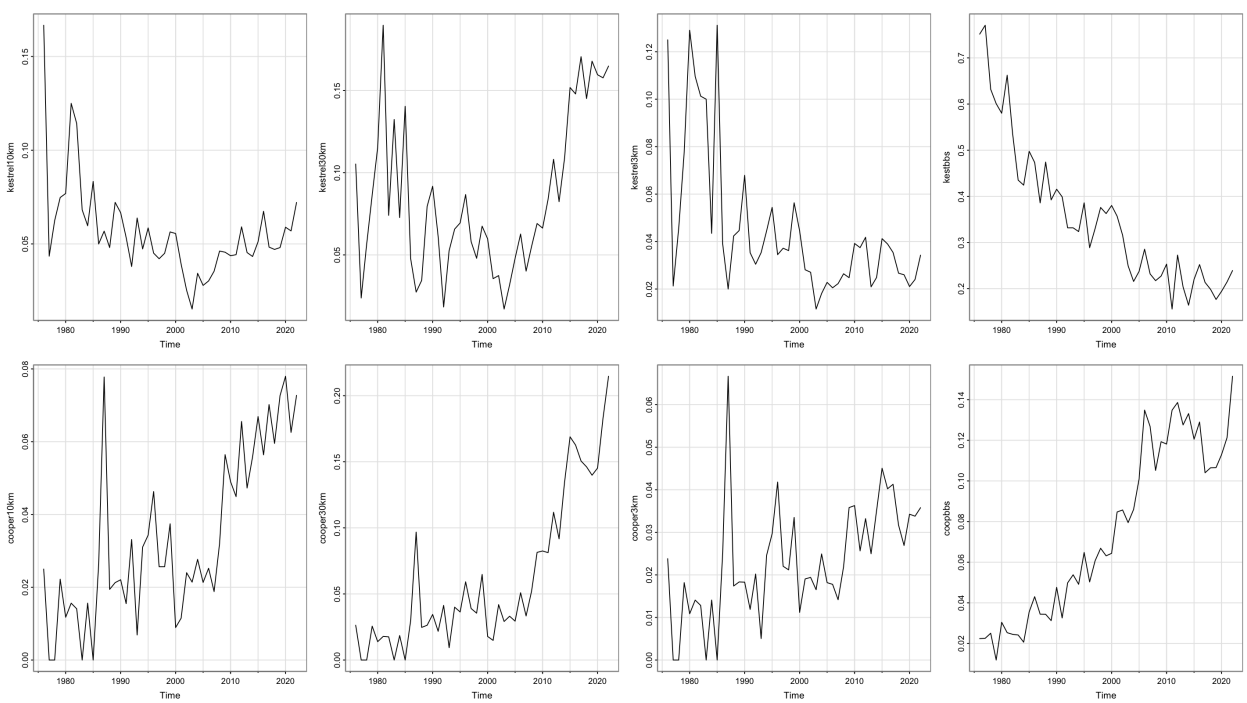

Figure 7: Series from subsampled checklists

4.2 Model Estimation and Forecasts

JK: Here, we describe the process of fitting the time series model to the subsampled eBird data. We provide details on model estimation, including the specific parameter settings and considerations for capturing seasonal and trend components. Using the fitted model, we perform a forecast to predict future population trends, interpreting the results in the context of American Kestrel population dynamics.

After subsampling, the amount of Kestrels and Cooper's Hawks found each year was simply divided by the number of checklists in the sample each year to get an accurate description of the study region based on the verification in the simulation study.

In eBird, users have the option to report that a bird was present without mentioning the count. In this case, the all observations with "X" were shifted to the median observation count of 1 for both Kestrels and Cooper's Hawks.

4.2.1 Regression Results (AICc) and Diagnostics for Top 3 Models

The regression with ARMA errors was done by using the Hyndman-Khandakar algorithm [12] to choose by AICc. Each of the three sampled series and the BBS Kestrel series were modeled with the corresponding Cooper's Hawk series as a covariate. The results of the modeling are shown in Table 3 below.

Series	ARMA Order (p, q)	AR1	AR2	MA1	Intercept	Xreg	Xreg SE	AICc
3 km resolution series	(1, 1)	0.9157	-	-0.6536	0.0548	-0.2982	0.3153	-202.77
10 km resolution series	(1, 0)	0.4246	-	-	0.0699	-0.2682	0.2542	-184.93
3 km Kestrel + 30 km resolution Cooper	(1, 1)	0.8894	-	-0.6248	0.0557	-0.1215	0.1247	-204.18
30 km resolution series	(2, 0)	0.3637	0.2564	-	0.0665	0.3259	0.1800	-178.24
BBS	(1, 1)	0.9889	-	-0.2560	0.3967	0.7794	0.6861	-126.01

Table 3: Regression with ARMA Errors Results with Cooper's Hawk Coefficients

The top two sampling buffers, as shown by the AICc values in Table 3, were the 10 km buffered sample with an AICc of -184.93 and the 3 km buffered sample with a lower AICc of -202.77. We validate the findings with the ARMA diagnostic plot.

In Table 4, we find that the eBird data at 3 and 10 km kestrel resolution are more relevant to BBS data. Thus, in addition to models at each resolution, one more model with the Kestrel 3 km and Cooper 30 km (3, 30) will be fit since they have the highest relation to BBS data.

Bird	**3 km**	**10 km**	**30 km**	**Simulated**
Kestrel	0.66	0.63	-0.13	0.65
Cooper	0.49	0.65	0.71	0.93

Table 4: Correlation analysis

The diagnostic plot for the model with the second and third-best AICc is shown in Figure 8, generated using the sarima function in Shumway and Stoffer [11]. The Ljung-Box statistic p values areg above the significance value, meaning that we fail to reject the null hypothesis that residuals are uncorrelated, indicating that the regression with ARMA errors has accounted for autocorrelation in the data. The trend was also picked up as can be seen in the Normal Q-Q Plot. The Q-Q plot also shows that the residuals are normally distributed as they follow the line. The Standardized residuals and their ACF also don't show much deviation, and the residual ACF stays close to zero, showing that the model extracted the trend in the data.

The diagnostic plot for the model with the best AICc shows similar conclusions with the Ljung-Box plot, normal Q-Q plot, and the residual and ACF plot in Figure 9. The Ljung-

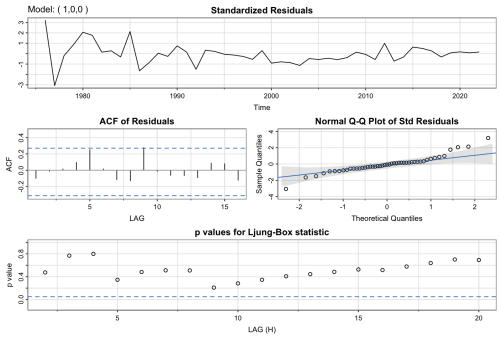

(a) Forecast for Cooper's Hawk

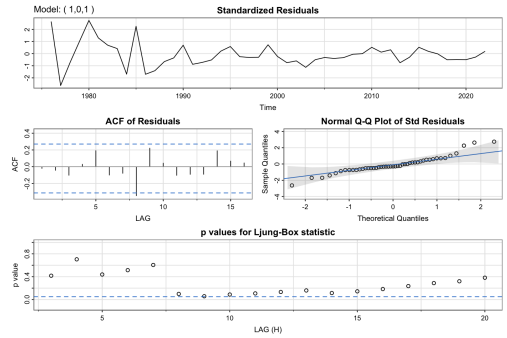

(b) Kestrel Forecast

Figure 8: Diagnostic plot for 10 and 3 km resolution models

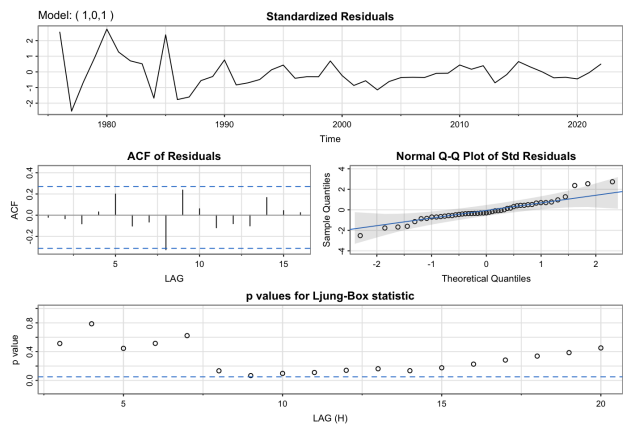

Figure 9: Diagnostic for best fit model

Box p-values and the normal Q-Q plot show that residuals are stationary and normally distributed, showing that the model has captured the trend.

4.2.2 Forecasts

Forecasts were fit on these two models based on the projected value of the Cooper's Hawk forecast. Using the same Hyndman-Khandakar algorithm for the covariate, time-based forecasts for the Cooper's Hawk abundance can be passed in as a parameter for Kestrel abundance. Figures 10, 12, and 14 shows different ways the exogenous Cooper's Hawk variable could change, and its forecasted impact on the Kestrel abundance for the 10 km and 3 km samples respectively, along with the same for the Kestrel 3 km - Cooper 30 km sample. Figures 11, 13, and 15 shows a forecast for the abundance of Kestrels using the regression with ARMA errors model based on a time-based forecast of Cooper's Hawks for the 10 km, 3 km samples respectively, and the sample with different resolutions respectively.

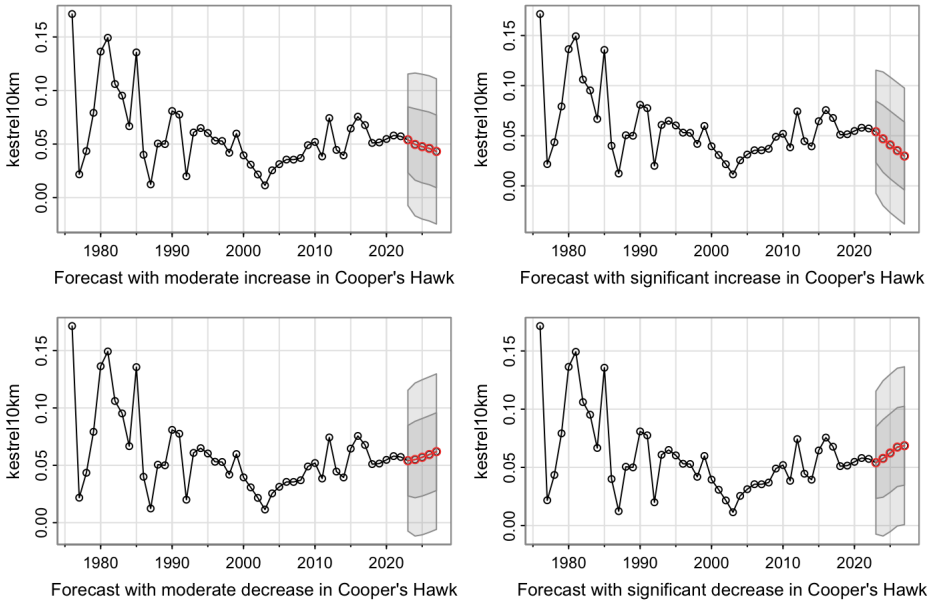

Figure 10: Forecasts for Kestrel 10 km Series based on different Cooper's Hawk changes

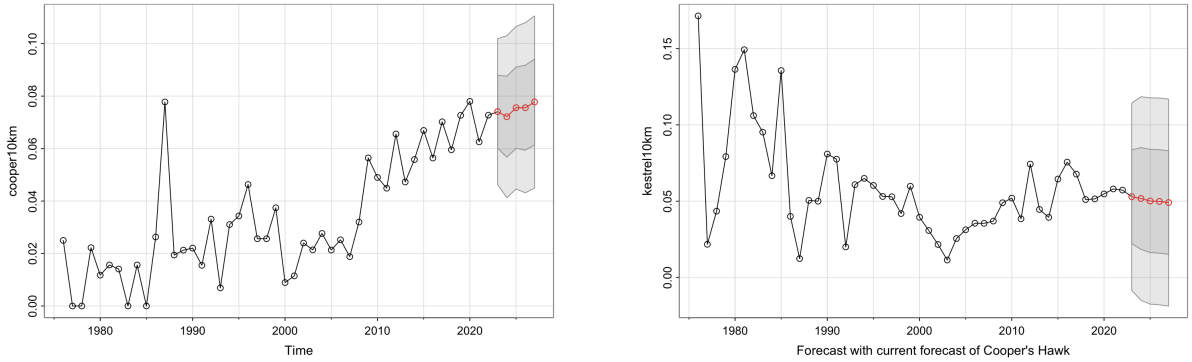

(a) Forecast for Cooper's Hawk

(b) Kestrel Forecast

Figure 11: Forecast for Kestrel Series at 10 km Resolution Based on Cooper's Hawk Forecast

Forecast Accuracy Figure 16 shows the RMSE of the regression with ARMA errors depending on different forecast steps. As expected, the forecasting accuracy declined as the number of forecast steps improved, but in comparison to the standard linear model, the acknowledgment of the time-dependent structure led to an increase in forecast accuracy. For the time series models, the RMSE ranged from 0.05-0.0875 while the linear regression model had an RMSE of $0.085 - 0.1$ for forecasts 5+ years ahead.

5 Discussion

The results from the analysis and the forecasts supported the initial hypothesis that the quick rebound of predatory species, combined with the habitat changes by the Kestrel population

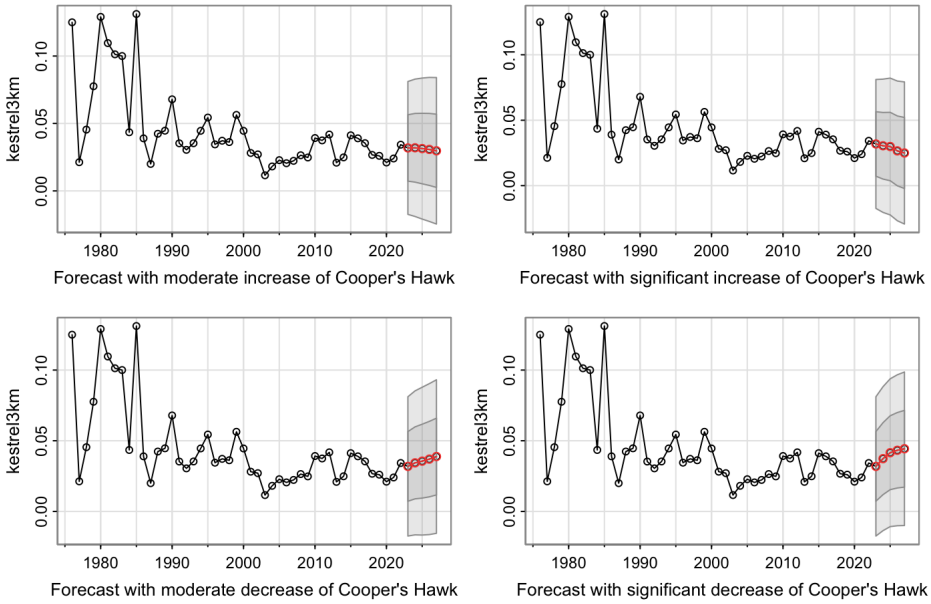

Figure 12: Forecasts for Kestrel 3 km Series based on different Cooper's Hawk changes

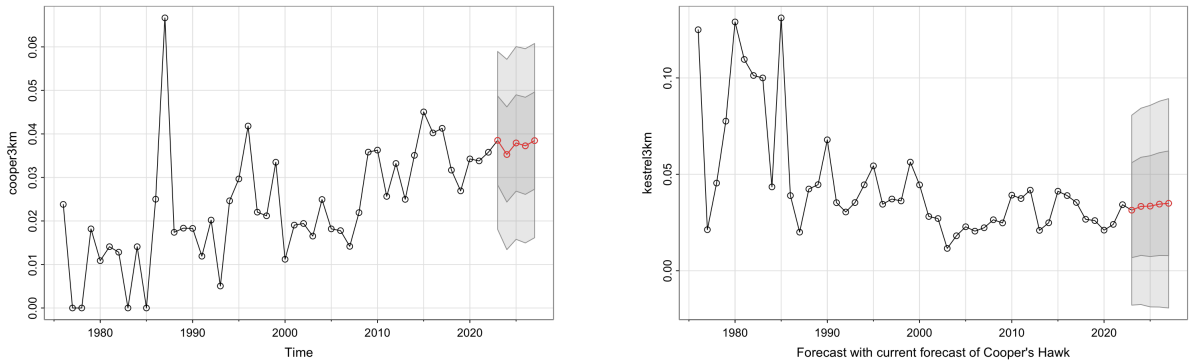

(a) Forecast for Cooper's Hawk

(b) Kestrel Forecast

Figure 13: Forecast for Kestrel Series at 3 km Resolution Based on Cooper's Hawk Forecast

after the ban of DDT, has continued to drive Kestrel numbers down. The analysis also supports the conclusion that, while the ban on DDT was helpful for conservation purposes, it may not be enough to save certain bird species such as the Kestrels due to other compounding factors.

In the samples of eBird data, at the representative 3 km and 10 km samples of the data, we can see that the increasing Cooper's Hawk abundance in the forecasts led to a decrease in Kestrel abundance, along with the external regressor showing a negative coefficient.

5.1 Analysis of Subsampling and Simulation Results

The subsampling was able to alleviate spatiotemporal biases for both species at the 3 and 10 km resolutions as shown in Figure 7. At these resolutions, the series resembled the general

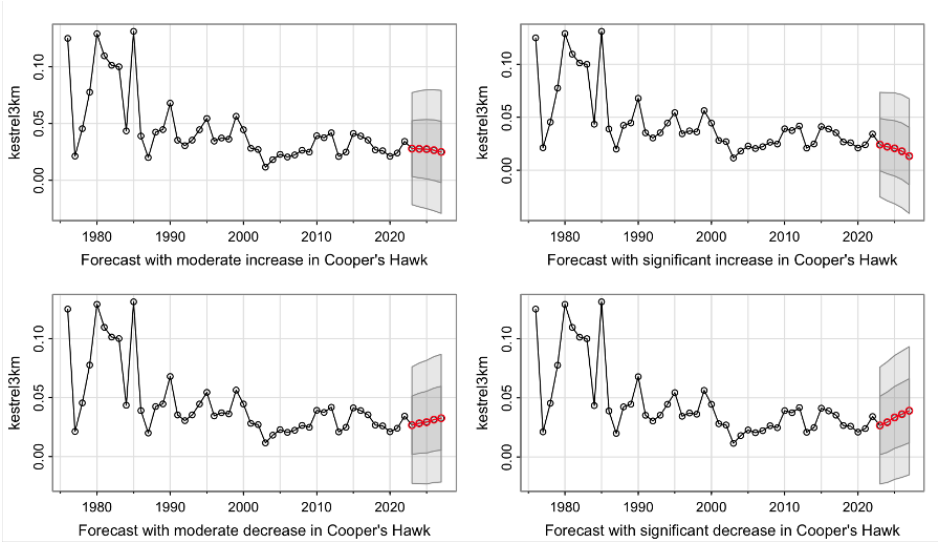

Figure 14: Forecast for Different Resolution Series at 3 km Resolution Based on Cooper's Hawk Forecast

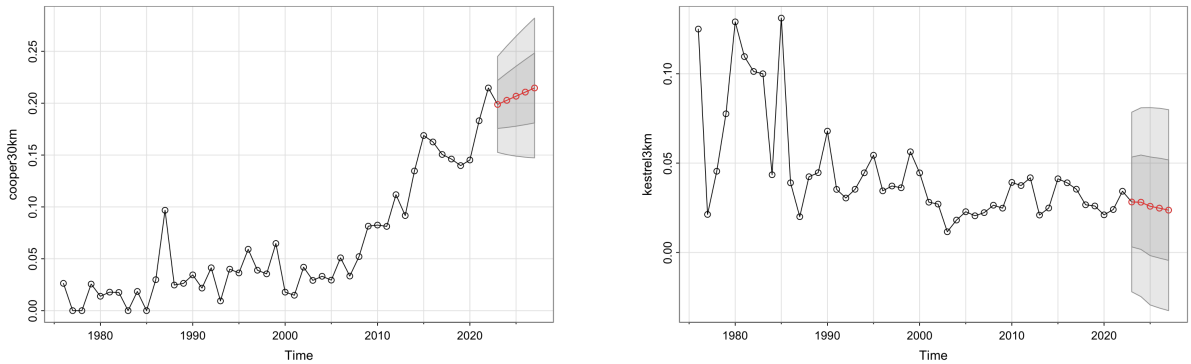

(a) Forecast for Cooper's Hawk

(b) Kestrel Forecast

Figure 15: Forecast for Kestrel Series at 3 km Resolution Based on Cooper's Hawk Forecast

BBS trend for their species. However, at the 30 km resolution, there were some limitations to the success of the sampling.

First, the sampling process uses case control to oversample detections. Since Kestrels and Cooper's Hawks are rare, oversampling detections addresses the class imbalance between detection and non-detection checklists. This oversampling worked well for 3 km and 10 km grid boxes, since most checklists in the grid box were representative of the grid box. However, when the boxes are larger, a spatiotemporal bias is reintroduced. As the number of checklists increased, the number of gridboxes that contained a checklist also increased. The new boxes bordered open areas because of their large size (30 km) and Kestrels are more frequent in these areas. Checklists outside of Kestrel habitat could be boxed into the large grid size and are not represented creating a bias in this sample. This spatiotemporal bias is illustrated in Figure 17.

The simulation process had high accuracy at discerning the spatial inhomogeneity in the

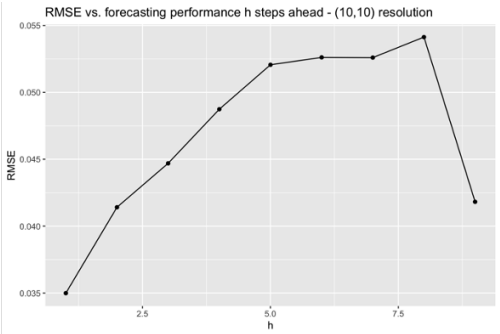

(a) RMSE for 10 km

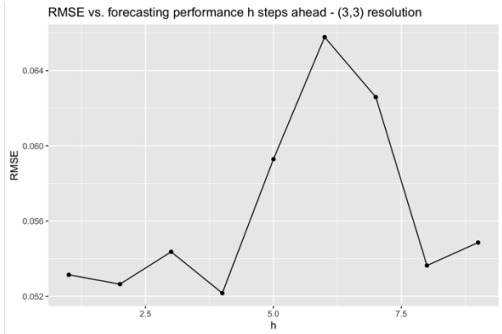

(b) RMSE for 3 km

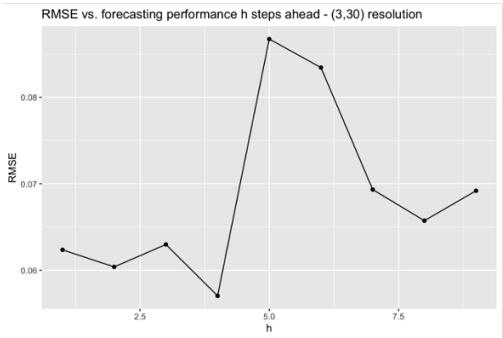

(c) RMSE for (3, 30) resolution

Figure 16: Forecast for Kestrel Series at 10 km Resolution Based on Cooper's Hawk Forecast

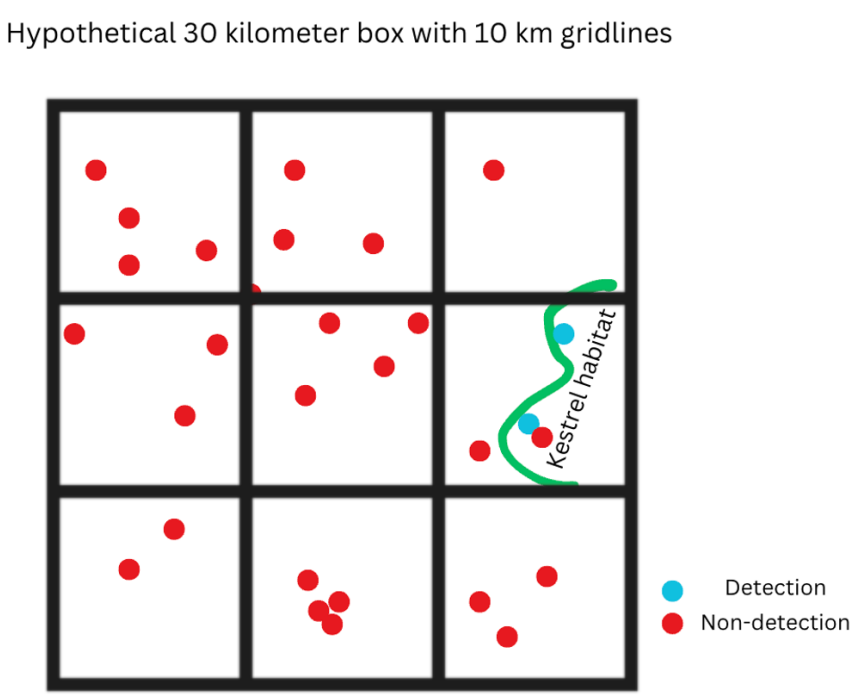

Figure 17: Case Control Sampling on a Hypothetical Grid

data. The simulated coefficient of -0.23 was similar to the real model coefficients of -0.26 and -0.22 in the (3, 3) and (10, 10) resolution subsample. The simulated series' spatial dependence on the GRF made spatial variation important to models at individual grid cells, showing that the reconstruction of the data was accurate as a spatial consensus.

5.2 Analysis of Regression & Forecasts

The AICc values in Table 3 favor the series with checklists sampled with 3 km buffers. Among similar datasets, this model performs the best in its ability to forecast. Both diagnostic plots for the models showed similar features in their Q-Q plots and Ljung-Box statistics, but the finer spatial resolution allowed for inclusion of checklists representative of their spatial grids more than the 10 km or 30 km buffered samples. All of the forecasts in the representative samples pointed to the increase in Cooper's Hawks pointing to decreases in Kestrel populations, and the strong negative coefficients in the models show that there is a degree of truth to the hypothesis.

5.3 Research Novelty and Potential for Further Exploration

Past research has focused on standard/traditional linear models to observe species trends. With the amount of data increasing exponentially, especially large, crowd-sourced citizen science datasets (such as eBird), analyzing the data through time-series analysis methods, while automatically accounting for biases and errors, is a novel addition that opens up many new possibilities. These methods can be used to both explain ecological relationships, as well as identify new dependent and independent variables that affect the environment, in this particular case, the real drivers impacting the decline in Kestrel population. Additionally, this research is focused on building models that can be enhanced, reused and built upon to understand impact on a broader scale.

Further research could focus on the inclusion of more covariates into the regression with ARMA errors, and that time-based models with more parameters such as General Autoregressive Conditional Heteroskedasticity (GARCH) models can aid in this analysis. Machine learning can also be added to this analysis, and Automated ML could build models quicker and efficiently. Further enhancements of the statistical analysis present can reshape understanding of our bird populations and subsequently determine impactful steps for their conversation.

6 Acknowledgements

We would like to thank Dr. Jim Bednarz for his insights on American Kestrels, and the University of North Texas for providing the platform to conduct this research on campus this year.

References

[1] The Peregrine Fund. American kestrel: Explore raptors.

[2] Brian L. Sullivan, Christopher L. Wood, Marshall J. Iliff, Rick E. Bonney, Daniel Fink, and Steve Kelling. ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 142:2282–2292, 10 2009.

[3] Rob J Hyndman and George Athanasopoulos. *Forecasting : Principles and Practice*. Otexts, 2 edition, 2018.

[4] Cornell Lab of Ornithology. ebird basic dataset. Version: EBD_relApr-2024.

[5] James C. Bednarz, Jr. Klem, Daniel, Laurie J. Goodrich, and Stanley E. Senner. Migration counts of raptors at hawk mountain, pennsylvania, as indicators of population trends, 1934-1986. *The Auk*, 107:96–109, 01 1990.

[6] D.J. Ziolkowski Jr., M. Lutmerding, W.B. English, V.I. Aponte, and M-A.R. Hudson. North american breeding bird survey dataset 1966 - 2022, 2023.

[7] John R. Sauer, Keith L. Pardieck, David J. Ziolkowski, Adam C. Smith, Marie-Anne R. Hudson, Vicente Rodriguez, Humberto Berlanga, Daniel K. Niven, and William A. Link. The first 50 years of the north american breeding bird survey. *The Condor*, 119:576–593, 08 2017.

[8] G. W. Luck, T. H. Ricketts, G. C. Daily, and M. Imhoff. Alleviating spatial conflict between people and biodiversity. *Proceedings of the National Academy of Sciences*, 101:182–186, 12 2003.

[9] Alison Johnston, Wesley M. Hochachka, Matthew E. Strimas-Mackey, Viviana Ruiz Gutierrez, Orin J. Robinson, Eliot T. Miller, Tom Auer, Steve T. Kelling, and Daniel Fink. Analytical guidelines to increase the value of community science data: An example using ebird data to estimate species distributions. *Diversity and Distributions*, 27:1265–1277, 05 2021.

[10] Nariaki Sugiura. Further analysts of the data by akaike' s information criterion and the finite corrections. *Communications in Statistics - Theory and Methods*, 7:13–26, 01 1978.

[11] Robert H. Shumway and David S. Stoffer. *Time series : a data analysis approach using R*. CRC Press, Taylor and Francis Group, 1 edition, 05 2019.

[12] Rob J. Hyndman and Yeasmin Khandakar. Automatic time series forecasting: The forecast package for r. *Journal of Statistical Software*, 27, 2008.