
Frequentist Prediction Sets for Species Abundance using Indirect Information

[To be read before The Royal Statistical Society at the Discussion Meeting on the 'Analysis of citizen science data' held at the Society's 2024 annual conference in Brighton on Tuesday, 3 September 2024, the President, Dr Andrew Garrett, in the Chair]

Elizabeth Bersson[✉]* and Peter D. Hoff

Department of Statistical Science, Duke University, Durham, 27701, NC, USA

*Corresponding author. elizabeth.bersson@duke.edu

Abstract

Citizen science databases that consist of volunteer-led sampling efforts of species communities are relied on as essential sources of data in ecology. Summarising such data across counties with frequentist-valid prediction sets for each county provides an interpretable comparison across counties of varying size or composition. As citizen science data often feature unequal sampling efforts across a spatial domain, prediction sets constructed with indirect methods that share information across counties may be used to improve precision. In this article, we present a nonparametric framework to obtain precise prediction sets for a multinomial random sample based on indirect information that maintain frequentist coverage guarantees for each county. We detail a simple algorithm to obtain prediction sets for each county using indirect information where the computation time does not depend on the sample size and scales nicely with the number of species considered. The indirect information may be estimated by a proposed empirical Bayes procedure based on information from auxiliary data. Our approach makes inference for under-sampled counties more precise, while maintaining area-specific frequentist validity for each county. Our method is used to provide a useful description of avian species abundance in North Carolina, USA based on citizen science data from the eBird database.

Key words: categorical data, conformal prediction, empirical Bayes, exchangeability, frequentist coverage, nonparametric

Introduction

Understanding species abundance across heterogeneous spatial areas is an important task in ecology. Citizen science databases that consist of observations of species counts gathered by volunteers are increasingly regarded as one of the richest sources of data for such a task. One of the largest such data sources is the eBird database in which citizen scientists throughout the world input counts of bird sightings [Sullivan et al., 2009]. In addition to its use for describing avian species abundance, eBird is a principal resource for understanding global biodiversity and is widely used in constructing and implementing conservation action plans [Sullivan et al., 2017].

More generally, analyses from such databases may be used for informing policy, conservation efforts, habitat preservation, and more, for which understanding species prevalence for non-overlapping geographic areas, such as counties across a state or country, is important. In practice, species abundance from citizen science data are commonly summarised within areas such as counties by empirical proportions from a sample, as in, e.g., Arnold et al. [2021], Camerini and Groppali [2014]. Such proportions can be used

to construct a prediction set for each county that provides a description of species prevalence for that county with guaranteed frequentist coverage.

Given the impact on policy design, corresponding uncertainty quantification is of particular import [Lele, 2020], and so it is desirable that precise prediction sets maintain a target coverage rate regardless of the county's size or composition. This is challenging as a common feature of citizen science data is unequal sampling efforts that results in some counties with large amounts of data information and others with very little. Using direct procedures that only make use of within-county information, a prediction set may be imprecise in these counties with low sampling efforts. This suggests using indirect information such as data from neighboring counties to improve prediction set precision for a given county.

In this article, we describe species abundance across sampling areas such as counties with frequentist-valid prediction sets that are constructed to contain an unobserved bird with $1 - \alpha$ probability. That is, a valid prediction set for a given county is a set of avian species such that an unobserved bird will belong to one of those species with $1 - \alpha$ probability in a frequentist sense. We

develop a valid nonparametric prediction method that allows for information to be shared across counties. Specifically, our approach results in prediction sets with guaranteed frequentist coverage for each county that are constructed with the incorporation of indirect or prior information. We detail and provide code for an empirical Bayes procedure to estimate such prior information from auxiliary data such as neighboring counties. If this indirect information used to construct the prediction sets is accurate, the prediction sets will be smaller than direct sets that only make use of within-county information.

Identifying avian species that are present across a spatial domain is of general interest in ecology [Twedt et al., 2010, Lebrun et al., 2012, Shanahan and Possingham, 2009]. To this aim, in Section 4, we detail the usefulness of the proposed approach in summarising the eBird citizen science data. For this task, frequentist-valid prediction sets provide useful summaries of the data that may be used to compare information across subregions and better inform policy. In particular, a frequentist-valid prediction set consists of species that are likely to be present in a given area and can be used to draw statistically robust conclusions regarding future observations. Such a prediction set is constructed based on observed species counts, and, by maintaining a specified frequentist coverage rate, reflects the uncertainty in the data. Commonly reported summaries of observed species abundance data, such as a list of observed proportions or counts do not reflect the uncertainty in the data. Moreover, the approach we propose constructs frequentist-valid prediction sets with the incorporation of auxiliary information. This is desirable as sharing information across counties generally results in smaller prediction sets as compared to direct prediction approaches, particularly so in counties with low sampling efforts. In contrast, indirect approaches constructed with auxiliary information such as a Bayesian prediction set may reflect uncertainty in the data, but are not calibrated.

Methodology

Background and Notation

For county $j \in \{1, \dots, J\}$, let \mathbf{X}_j be a vector of length K where $X_{j,i} = x_{j,i}$ is the observed count of species i over some set sampling period that may vary across counties. We model \mathbf{X}_j with a K -dimensional multinomial distribution with $N_j = \sum_{i=1}^K x_{j,i}$ trials and population proportions vector $\boldsymbol{\theta}_j$,

$$\mathbf{X}_j \sim MN_K(\boldsymbol{\theta}_j, N_j). \quad (1)$$

We construct a prediction set $A_\alpha(\mathbf{X}_j)$ for an observation of a new bird arising from the same distribution, $\mathbf{Y}_j \sim MN_K(\boldsymbol{\theta}_j, 1)$ where $\mathbf{Y}_j \in \mathcal{Y}$ for $\mathcal{Y} = \{(y_1, \dots, y_K) : \sum_{i=1}^K y_i = 1, y_i \in \{0, 1\}\} (i = 1, \dots, K)\}$. Let $\mathbf{y}_j^{(k)} \in \mathcal{Y}$ denote a prediction of category k , that is, let $\mathbf{y}_j^{(k)}$ be a vector of length K with a one at index k and zeros elsewhere. In particular, we are interested in a prediction set for \mathbf{Y}_j that maintains frequentist validity for some error rate α . Formally, we refer to this as an α -valid prediction set:

Definition (α -Valid Prediction Set) An α -valid prediction set for a predictand $\mathbf{Y}_j \in \mathcal{Y}$ is any subset A_α of the sample

space \mathcal{Y} that contains \mathbf{Y}_j with probability greater than or equal to $1 - \alpha$,

$$P_\theta(\mathbf{Y}_j \in A_\alpha(\mathbf{X}_j)) \geq 1 - \alpha, \quad \forall \theta, \quad (2)$$

where the probability is taken with respect to \mathbf{Y}_j and \mathbf{X}_j .

Additionally, small or precise α -valid prediction sets are of particular interest, where prediction set size is measured by expected cardinality, that is, expected number of the K categories in the sample space included in the prediction set.

Order-based prediction for a single area

A standard approach to construct α -valid prediction sets for each county or area is with a direct method that only makes use of within-area information. As such, we first consider construction of a prediction set for a single area j , using only data from area j . For ease of notation, we drop the area-identifying subscript in this subsection.

For multinomial data in general, if the event probability vector $\boldsymbol{\theta}$ is known, an α -valid prediction set is any combination of categories such that their event probabilities cumulatively sum to be greater than or equal to $1 - \alpha$. Equivalently stated, an α -valid prediction set may be constructed by excluding categories such that the cumulative sum of the excluded categories' event probabilities is less than α . Such a prediction set may be constructed by admitting categories in some prespecified order into the prediction set until the cumulative sum of their event probabilities is at least $1 - \alpha$. The resulting prediction set will have $1 - \alpha$ coverage regardless of the ordering used to admit categories. In fact, the class of all α -valid prediction sets may be constructed by following this procedure for non-strict total orderings of categories.

Perhaps intuitively, constructing such a prediction set by including categories with the largest event probabilities will result in the smallest α -valid prediction set. In the terminology of ordering, this corresponds to constructing a prediction set based on an ordering of categories that matches the ordering of the elements in $\boldsymbol{\theta}$. We refer to this optimal ordering as the oracle ordering:

Theorem 1 (Oracle order-based prediction) *Let $\mathbf{Y} \sim MN_K(\boldsymbol{\theta}, 1)$ for $\boldsymbol{\theta}$ known. Then,*

1. *the class of all α -valid prediction sets for a given $\boldsymbol{\theta}$ consists of prediction sets of the form,*

$$A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}} = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1}(o_k \geq o_l) \theta_l \right] > \alpha \right\}, \quad (3)$$

for some vector $\mathbf{o} \in \mathbb{R}^K$, and

2. *the **oracle ordering** is that which corresponds to the increasing order statistics of $\boldsymbol{\theta}$,*

$$\mathbf{o}^\theta = \{\mathbf{o} : \theta_m < \theta_n \Rightarrow o_m < o_n \forall m, n \in \{1, \dots, K\}, m \neq n\},$$

and $A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}^\theta}$ has the smallest cardinality among all orderings.

In practice, $\boldsymbol{\theta}$ is unknown, but a prediction set may be constructed based on an observed sample $\mathbf{X} = \mathbf{x}$. It

turns out, in fact, that any conditional α -valid prediction set can be written similarly to the previous construction (Equation 3) where the cumulative sum is computed with respect to the empirical proportions given by \mathbf{x} and \mathbf{y} . This is a generalization of the conformal prediction framework, a popular machine learning approach to construct prediction regions based on measuring conformity (or non-conformity) of a predictand to an observed sample [Vovk et al., 2005].

Theorem 2 (α -valid order-based prediction) *Let $\mathbf{X} \sim MN_K(\boldsymbol{\theta}, N)$, $\mathbf{Y} \sim MN_K(\boldsymbol{\theta}, 1)$. Then, every conformal α -valid prediction set based on observed data \mathbf{x} can be written*

$$A_\alpha(\mathbf{x}) = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1}(o_k \geq o_l) \frac{x_l + y_l^{(k)}}{N+1} \right] > \alpha \right\}, \quad (4)$$

for some vector $\mathbf{o} \in \mathbb{R}^K$.

Note that the prediction set depends on the vector \mathbf{o} only through the order of its elements. Such a vector may be defined based on the data, or otherwise. We elaborate in the remainder of this section.

For any ordering of the K categories, constructing a prediction set following Theorem 2 results in a prediction set with guaranteed finite-sample $1 - \alpha$ frequentist coverage. The choice of ordering, however, will impact prediction set precision, that is, the set's cardinality. For inference for a single area, a natural approach is to order the categories with respect to their empirical proportions. The empirical proportions are unbiased for population proportions, so, if the area has a large sample size, an ordering based on the empirical proportions will approximate the oracle ordering well. It turns out this approach is well-motivated by classical prediction approaches. Specifically, a standard direct prediction method constructs a prediction set separately for an area based on an area-specific conditional pivotal quantity [Faulkenberry, 1973, Tian et al., 2022]. For a multinomial population, $\mathbf{Y}|\mathbf{X} + \mathbf{Y}$ is such a quantity that follows a multivariate hypergeometric distribution which does not depend on the event probability vector. See Thatcher [1964] for work on prediction sets of this type for binomial data. A prediction set constructed to contain species belonging to a highest mass region of this pivotal distribution is obtained by including species with the largest empirical counts until their cumulative proportion sum exceeds $1 - \alpha$,

$$A_\alpha^D(\mathbf{x}) = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1} \left((x_k + y_k^{(k)}) \geq (x_l + y_l^{(k)}) \right) \frac{x_l + y_l^{(k)}}{N+1} \right] > \alpha \right\}. \quad (5)$$

This direct prediction set based on an ordering of the empirical proportions is appealing as it is easy to interpret and has finite-sample guaranteed $1 - \alpha$ frequentist coverage. For an area with low sampling effort, though, the empirical proportions will not precisely estimate the true proportions. As a result, a prediction set may have prohibitively large cardinality such that it is not practically useful. For such an area, incorporating indirect information from neighboring

counties can improve the estimates of the county proportions and thereby increase the precision of a prediction set.

Order-based prediction for multiple areas

In general, in analyzing small area data, that is, areal data featuring small within-area sample sizes in some areas, it is common to utilize indirect methods that share information across areas [Rao and Molina, 2015]. The eBird database is a rich data source, and inference in any given county may be improved upon by taking advantage of auxiliary data using an indirect method. In this subsection, we detail how information from neighboring counties may be used in estimating an ordering of categories to improve prediction set precision.

As opposed to a direct prediction set based on an ordering corresponding to within-county empirical proportions, an indirect prediction set can be constructed similarly whereby species are admitted into the prediction set based on an ordering corresponding to empirical posterior proportions estimated from a hierarchical model. Such an estimate may be obtained based on a conjugate Dirichlet prior distribution parameterized with a common concentration hyperparameter for the J areas,

$$\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_J \sim \text{Dirichlet}_K(\boldsymbol{\gamma}). \quad (6)$$

Given a hyperparameter $\boldsymbol{\gamma} \in \mathbb{R}^K$, the posterior expectation of the proportions $\boldsymbol{\theta}_j$ in county j is $\tilde{\mathbf{x}}_j / (N_j + \sum_{i=1}^K \gamma_i)$ where $\tilde{\mathbf{x}}_j = \mathbf{x}_j + \boldsymbol{\gamma}$. In this way, $\tilde{\mathbf{x}}_j$ may be interpreted as a posterior vector of counts for county j . Then, an α -valid prediction set based on $\tilde{\mathbf{x}}_j$ is,

$$A_\alpha^I(\mathbf{x}_j) = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1} \left((\tilde{x}_{j,k} + y_k^{(k)}) \geq (\tilde{x}_{j,l} + y_l^{(k)}) \right) \frac{x_{j,l} + y_l^{(k)}}{N_j + 1} \right] > \alpha \right\}. \quad (7)$$

By Theorem 2, $A_\alpha^I(\mathbf{x}_j)$ is an α -valid procedure, and it is constructed based on prior information. Specifically, it differs from the direct set given in Equation 5 in that categories are admitted into the prediction set based on an ordering determined by posterior counts that incorporate indirect information $\boldsymbol{\gamma}$, as opposed to an ordering based on the observed sample. Moreover, it has been shown that if the indirect information used is accurate, $A_\alpha^I(\mathbf{x}_j)$ may be more precise than a direct prediction set with the same coverage rate [Hoff, 2023, Berrson and Hoff, 2024].

In total, $A_\alpha^D(\mathbf{x}_j)$ and $A_\alpha^I(\mathbf{x}_j)$ are both α -valid prediction procedures. They differ in the order in which species are admitted into the prediction sets, as species are admitted into the direct set in terms of decreasing empirical proportions and into the indirect set in terms of decreasing posterior counts. As a result, for an area with a small sample size, incorporating accurate prior information can result in an ordering used to construct a prediction set that more accurately approximates the oracle ordering as the empirical proportions might be too unstable. Of note, these two approaches are equivalent for a uniform prior $\boldsymbol{\gamma} = c\mathbf{1}$, for any constant c . This includes, for example, a standard noninformative prior $c = 1$, a standard objective Bayes Jeffrey's prior $c = 1/2$, and an improper prior $c = 0$.




Fig. 1: Monte Carlo approximations (+/- 1 standard deviation) of the expected cardinality ratio of (red) indirect to direct methods and (blue) α -valid prediction set given the oracle ordering to direct method.

Empirical Bayes estimation of indirect information

To obtain an α -valid indirect prediction set for county $j \in \{1, \dots, J\}$, all that is required is an estimate of the prior concentration parameter γ . We propose an empirical Bayesian approach whereby values of γ to be used for county j are estimated from data collected in neighboring counties. Specifically, we use the maximum likelihood estimate of the marginal likelihood based on the conjugate hierarchical model given by Equations 1 and 6,

$$\begin{aligned} \gamma_j &= \arg \max_{\gamma} \log p \left(\bigcup_{l \in L_j} \mathbf{X}_l \mid \gamma \right) \\ &= \arg \max_{\gamma} \log \prod_{l \in L_j} \left[\frac{\Gamma(\sum_{i=1}^K \gamma_i)}{\Gamma(\sum_{i=1}^K x_{l,i} + \gamma_i)} \times \prod_{i=1}^K \frac{\Gamma(x_{l,i} + \gamma_i)}{\Gamma(\gamma_i)} \right], \end{aligned} \quad (8)$$

where $L_j \subseteq \{1, \dots, J\} \setminus \{j\}$ is a non-empty set containing the indices of counties neighboring county j . Information is shared across neighboring counties to inform an estimate of the prior for county j , and, when estimated in this way, the prior concentration represents an across-county pooled prior concentration. This optimization problem can be solved numerically with a Newton-Raphson algorithm. See Appendix B for details and derivation of such an algorithm. Code to implement this procedure in the R Statistical Programming language is available online, see Section 5.

When γ_j is estimated using data independent of area j and used to construct $A_{\alpha}^I(\mathbf{x}_j)$, the finite sample coverage guarantee of $A_{\alpha}^I(\mathbf{x}_j)$ holds regardless of the accuracy of the estimated prior hyperparameter. If the estimated vector γ_j is accurate, then $A_{\alpha}^I(\mathbf{x}_j)$ may also be more precise than direct prediction approaches.

Simulation study

To illustrate how the incorporation of indirect information can affect precision of prediction sets, we compare expected set cardinality obtained from the indirect and direct prediction methods for a single simulated area. In contrast to the eBird data, for example, the analysis of this

section corresponds to that of one county. Because citizen science data such as these often feature unequal sampling efforts across counties, we are particularly interested in demonstrating the difference in cardinality between these two approaches for a range of sample sizes $N = 10, 100, 1000$. Moreover, we compare results for varying number of categories K . Throughout, we consider a low entropy regime in which $\lceil K/4 \rceil$ categories unequally split nearly all of the probability mass, and the rest of the categories have nearly probability 0. While we do not necessarily expect real populations in practice to have such a distribution, it is chosen to clearly demonstrate the benefit of including indirect information in the construction of prediction sets that maintain frequentist coverage.

In one construction of indirect prediction sets, we consider a prior based on full information with moderate prior precision $\gamma = \theta \times 10$. We compare with direct prediction sets given by Equation 5, or, equivalently, indirect prediction sets constructed with a uniform prior $\gamma = \mathbf{1}$. Finally, we compare the approaches to α -valid order-based prediction sets obtained based on an oracle ordering. Results comparing Monte Carlo approximations of the expected prediction set cardinality ratios between the various approaches obtained from 25,000 replications are displayed in Figure 1.

As all methods considered are α -valid procedures, the crucial difference between them is the incorporation of indirect information. Utilizing accurate prior information in the construction of prediction sets generally results in prediction sets distinctly smaller than direct sets, particularly so if there are a large number of categories relative to the sample size. This is evidenced by the red dashes in Figure 1 showing the expected cardinality ratios of the indirect to direct prediction sets are always at or below a value of 1. An accurate prior may be one that approximates the true probability mass vector well with large precision relative to sample size, as seen in the left plot of Figure 1 for sample size $N = 10$. More generally, though, all that is needed is a prior that results in posterior counts that accurately approximate the oracle ordering of categories. We discuss the three sample size regimes in detail below.

For a small sample size of $N = 10$, the prior γ used to construct the indirect prediction sets is an informative prior with strong precision in that the scale used is equal to the sample size in this case. As a result, the posterior distributions contain notably more information than what is in each simulated dataset. As a result, the ordering of categories induced by the posterior counts, used to construct the indirect prediction sets, are accurately approximating the oracle ordering of categories. This is evidenced by the nearly identical behavior of the two cardinality ratios explored. In conjunction with the instability of the direct method in the presence of such a small sample size, this results in notably smaller cardinality of the indirect set as compared to the direct set, even for relatively small total numbers of categories. At its best, the indirect prediction set is about 80% smaller than the direct set.

For a moderate sample size of $N = 100$, the prior precision used to construct the indirect prediction sets is not overwhelming as compared to the sample size, and hence the posterior counts do not approximate the oracle ordering as well as in the regime with a smaller sample size. This is evidenced by the divergence of the red and blue dashes in the middle plot of Figure 1. Still, particularly as the number of categories increases for fixed N , the benefit of utilizing prior information of this type is highlighted by the decline of the cardinality ratio of the indirect to direct prediction sets (red lines). For example, in the case of $N = 100$ and $K = 150$, the indirect prediction set constructed with γ is about 15% smaller than the direct prediction set.

A similar but less pronounced pattern is seen in the presence of a larger sample size of $N = 1000$. For this sample size with $K \leq 150$, all methods considered perform relatively similarly. However, as the number of categories increases, there is a distinct gain in prediction set precision given the input of indirect information in prediction set construction.

Comparison with Bayesian prediction sets

For multi-group count data such as we are considering, a standard Bayesian prediction set may be constructed from the posterior predictive distribution (see, for example, Gelman et al. [2014] §1.3) based on the hierarchical model given by Equations 1 and 6. Specifically, the posterior predictive distribution for group j is a multinomial distribution for a single trial with proportions vector $\tilde{\mathbf{x}}_j / (N_j + \sum_{i=1}^K \gamma_i)$. Then, a Bayesian prediction set A_α^B may be taken to be the categories corresponding to the highest mass region of this predictive distribution such that the categories' cumulative predictive probability sum exceeds $1 - \alpha$,

$$A_\alpha^B(\mathbf{x}_j) = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1}(\tilde{x}_{j,k} \geq \tilde{x}_{j,l}) \frac{\tilde{x}_{j,l}}{N_j + \sum_{i=1}^K \gamma_i} \right] > \alpha \right\}. \quad (9)$$

Such a prediction set may have smaller cardinality than the indirect or direct approach, but it will not be α -valid. In fact, the frequentist coverage may arbitrarily be nearly zero under a misspecified prior parameter with a large scale.

To illustrate the relationship between cardinality and coverage, we compare the Bayesian prediction method to

the proposed indirect approach (7) in simulation. We follow the same simulation set-up as before and, for both methods, consider a prior based on full information with moderate prior precision $\gamma = \theta \times 10$. Results for a moderate sample size of $N = 100$ are plotted in Figure 2. By construction, the indirect approach is α -valid and maintains frequentist coverage rates at or above the nominal level. In contrast, the Bayesian prediction sets may be narrower than the indirect sets, but, correspondingly, the frequentist coverage rate of the Bayesian sets may fall below the nominal rate.

Summarising eBird species abundance data

In this section, we describe avian species abundance in North Carolina, USA from eBird data obtained from citizen-uploaded complete checklists of species observations in the first week of May 2023. Across the 99 counties, 393 unique species were identified. Some species such as the Northern Cardinal, Carolina Wren, and American Robin were identified frequently. Many others like the Northern Saw-whet Owl and the Solitary Sandpiper were rarely seen; in fact, 50% of species were seen fewer than 100 times each across the entire state. Moreover, within-county sample sizes vary drastically (Figure 3) from approximately 50,000 individual birds identified in Wake County, one of the most populous counties in NC that contains the state's capital, to only 8 in Pasquotank County, a small coastal county consisting of about 1/30th of the human population of Wake County.

As motivated in the Introduction, describing such data with α -valid prediction sets for each county provides a useful summary with unambiguous statistical interpretation. That is, with at least probability $1 - \alpha$, an unobserved bird in a given county will belong to a species contained in the specified prediction set, where the probability is taken with respect to the random sample and the predictand. Here, we demonstrate the usefulness of this approach in gaining better understanding of species abundance. Moreover, we elaborate on the benefit of utilizing indirect information in the construction of practically useful sets that are precise, particularly for counties with small within-county sample sizes.

For each county in NC, we construct an indirect prediction set based on a prior hyperparameter estimated from data in the five nearest neighboring counties, following the procedure described in Section 2.4. The eBird data consist of independent samples collected across the state, so samples are independent across counties. As a result of this independence, finite-sample coverage of the indirect prediction approach is guaranteed. We compare the cardinality of these indirect prediction sets to that of direct prediction sets, both of which maintain at least 95% coverage for each county. The cardinality ratios of the indirect to direct prediction sets across the counties in NC are plotted in Figure 4. To highlight the impact of within-county sample size, the lower quantile sample sizes are overlaid on their respective county.

In general, the incorporation of indirect information in the construction of prediction sets results in notably smaller cardinality of the indirect prediction sets as compared with that of the direct prediction sets. Of the 99 counties in NC,

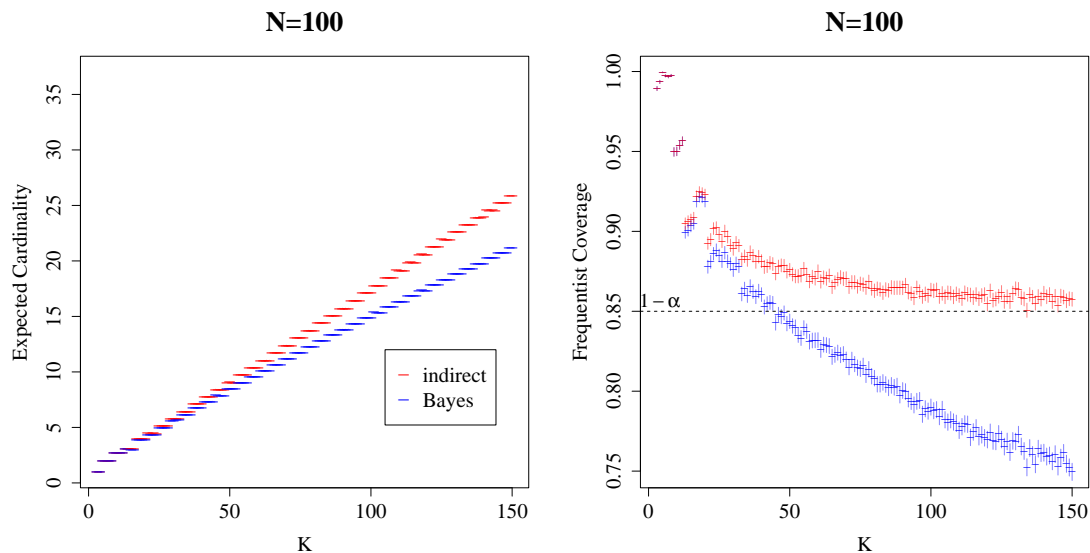


Fig. 2: (Left) Monte Carlo approximations (± 1 standard deviation) of the expected cardinality of the (red) indirect α -valid prediction set and (blue) Bayesian posterior predictive set, both constructed with prior concentration $\gamma = \theta \times 10$. (Right) Monte Carlo approximations (with 95% Coppler-Pearson intervals) of the frequentist coverage rate of each method, both constructed for an error rate of $\alpha = 0.15$.




Fig. 3: Within-county log sample size.

indirect sets have smaller cardinality in 65, and the two approaches result in the same cardinality in 20 counties. The improvement in cardinality is particularly conspicuous in counties with small to moderate sample sizes, as evidenced by the sample sizes of counties with the brightest shade of red in Figure 4. Moreover, ten counties have trivial direct sets consisting of all K species, while only two counties with the smallest within-county sample sizes, 8 and 14, have trivial indirect prediction sets. For the county with the third smallest sample size (24), the indirect prediction set only includes 80 species, or about 20% of all possible species, while the direct prediction set is the trivial set.

Overall, even in counties with larger sample sizes, it is most common for the indirect and direct prediction sets to contain a different set of species. In fact, the indirect and direct prediction sets disagree for nearly every county in NC. They are equivalent for only six counties where they aren't both trivial sets. Commonly, this discrepancy corresponds with smaller indirect sets, and hence highlights the benefit of inclusion of indirect information in the construction of prediction sets.

Order-based prediction in Robeson County

To further compare the two approaches and elucidate the role of the ordering of the species, we elaborate on

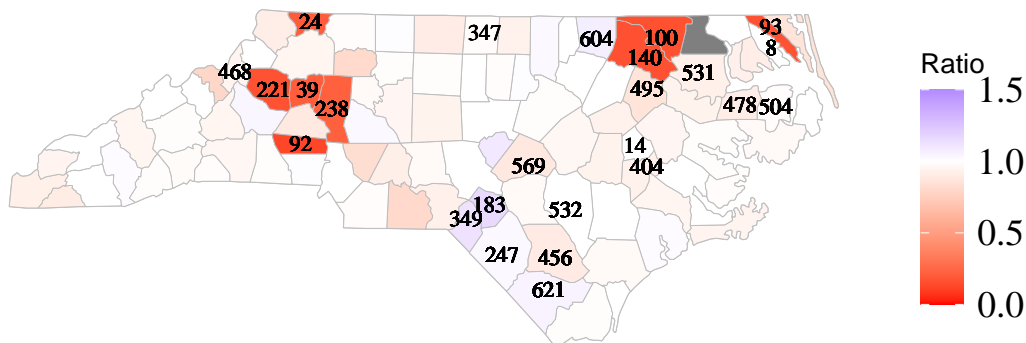


Fig. 4: Cardinality ratio of indirect to direct prediction sets. Prior hyperparameters estimated with an empirical Bayesian procedure based on five nearest neighbors for each county. The lowest quantile sample sizes are overlaid on their respective counties.




Fig. 5: Empirical probability masses of species included in only the indirect (red text), only the direct (blue text), or both (purple text) prediction sets, sorted by the posterior proportion. MLE plotted in blue, and the posterior proportion based on a prior estimated from data in nearest five neighboring counties.

the construction of indirect and direct prediction sets for Robeson County. Robeson is located near the southeastern border of NC and features a moderately small within-county sample size of 247 birds observed, with species-specific

observation counts ranging from zero to ten. The two prediction sets have nearly the same cardinality but contain differences in species inclusion. Specifically, the indirect

	D. Cormorant	E. Kingbird	Pine Warbler	C. Sparrow
Robeson	0.81%	0.4%	0.00%	0.00%
NC-017	0.00%	2.85%	1.97%	2.63%
NC-047	0.00%	0.00%	1.29%	0.64%
NC-051	0.00%	0.68%	2.62%	5.59%
NC-093	0.00%	0.00%	1.09%	0.55%
NC-165	0.00%	0.86%	2.58%	5.16%
γ	0.00	1.33	3.42	4.69

Table 1. Percentage of all birds observed within each respective county, for select species included in either the indirect set (red text) or the direct set (blue text). Estimated prior hyperparameter γ for Robeson County recorded in the last row.

	L. Flycatcher	R. Hawk	C. Yellowthroat	E. Kingbird	Bobolink
Haywood	0.21%	0.24%	0.21%	0.24%	0.24%
NC-021	0.06%	0.29%	0.18%	0.43%	0.14%
NC-099	0.00%	0.08%	0.16%	0.00%	0.00%
NC-115	0.07%	0.14%	0.21%	0.28%	0.00%
NC-173	0.18%	0.18%	0.66%	0.09%	1.41%
NC-175	0.00%	0.30%	1.00%	0.54%	0.84%
γ	0.4	1.29	2.43	1.49	2.15

Table 2. Percentage of all birds observed within each respective county for species included in either both prediction sets (purple text) or only the direct set (blue text). Estimated prior hyperparameter γ for Haywood County recorded in the last row. Species are sorted by posterior proportion.

prediction set contains 33 species, and the direct set contains 32, with an overlap of 27 species.

To illustrate the role of the ordering used in the construction of α -valid prediction sets, the empirical proportions based on the observed sample (MLE) and posterior proportions (Post.Pred) are plotted in Figure 5 for the union of included species in the two sets. In the figure, the species are sorted by increasing posterior proportions. The indirect and direct sets include species based on the posterior and empirical distributions, respectively. Discrepancies between the indirect and direct sets occur when these two distributions disagree. From Figure 5, it is easy to see the indirect prediction set consists of the species with the 33 largest posterior predictive proportions. In contrast, the direct set consists of species with the largest sample probability mass. Naturally, the ordering of these two estimates agree for species common to the region, and, as such, there is a fair amount of overlap of species inclusion.

As a result of our estimation procedure for the prior hyperparameter γ for Robeson County, the disparity between inclusion or exclusion of a species among the two prediction set methods is further elucidated by examining species presence in neighboring counties. In short, species with more frequent occurrence in neighboring counties will have a larger estimated prior count than those seen rarely in neighboring counties. Species occurrences in neighboring counties are displayed in Table 1 for a select few species along with the estimated γ for Robeson County, obtained by solving Equation 8 using data in these neighboring counties.

Intuitively, species that are seen in neighboring counties with some relative frequency, such as the Chipping Sparrow or Pine Warbler, are probably also present in Robeson County, and hence should be included in a prediction set. In practice, these species have a comparatively high estimated prior of about 5 and 4, respectively, and hence are included in the indirect prediction set even though they weren't recorded as being observed in Robeson County in the dataset. Alternatively, consideration of indirect information yields the conclusion that species like the Eastern Kingbird and Cormorant may be rare in the area in general, as reflected by small γ values, and thus these species are not included in the indirect prediction set.

Inference among species with tied observed counts in Haywood County

In species abundance data, particularly for areas or counties with small sample sizes, it is common for multiple species to have the same observed count. A feature of the construction

of the direct order-based prediction approach as presented is that species with the same observed counts will either be jointly included or excluded from the prediction set. As a result, a direct prediction set constructed from a sample with tied species counts may have increased cardinality over an indirect prediction set that does not necessarily jointly admit all species with tied observed counts. If the direct set has increased cardinality for this reason, the direct set will also have increased coverage over the indirect set.

When constructing a prediction set based on the empirical proportions without consideration of indirect information, as in the construction of the direct set, this may commonly occur, and there is no clear approach to choose among the species with tied counts without further information than what is provided in the sample in that county. One could randomly choose to include one of the species from the set of species with tied counts, for example, but a more principled manner is to utilize indirect information to determine which species should be included. This is the mechanism used by the indirect prediction approach when the prior hyperparameter is a real valued vector estimated from indirect information. As such, a more nuanced benefit of utilizing indirect information in the construction of a prediction set is the capacity to include a select few categories with tied empirical proportions.

To demonstrate, we elaborate on species inclusion in the indirect and direct prediction sets in Haywood County. Haywood is popular destination in the Blue Ridge Mountains, located near the western border of North Carolina. It features a moderately large within-county sample size of roughly 4000 birds observed. In Haywood County, the indirect prediction set contains 70 species, and the larger direct set contains 74. In the construction of these prediction sets, the ordering of species with regards to the posterior proportions and the empirical proportions agree for most species. As a result, all 70 species included in the indirect set are also included in the direct set. The disparity in species inclusion occurs primarily as a result of tied counts of species occurrence in the sample.

Empirical proportions in Haywood and neighboring counties are reported in Table 2 for the five species included in Haywood County's prediction sets with the smallest posterior proportions. The species with the four smallest posterior proportions are included only in the direct set, and the other species, the Bobolink, is included in both the indirect and direct sets. The Bobolink was observed 9 times in the sample from Haywood County, or about 0.24% of the Haywood sample. For an ordering determined by either

the empirical counts or the posterior counts, this species is required to be included in the order-based prediction set to guarantee $1 - \alpha$ coverage. Two of the other species, the Red-shouldered Hawk and Eastern Kingbird, were each also observed 9 times in the sample from Haywood, and, by construction of the order-based prediction approach, must also be included in the direct set. When admitting the species into a prediction set by posterior counts based on the real-valued prior hyperparameter γ estimated from data in neighboring counties, as in the indirect approach considered, the ‘tie’ among these three species is broken, and only one, the Bobolink, is included in the indirect prediction set.

Comparison with Bayesian prediction sets

The α -valid prediction approaches may also be compared to the Bayesian prediction method detailed in Equation 9. As the Bayesian method and the indirect method are constructed based on the same hierarchical working model, the Dirichlet prior parameter in (9) may be estimated following the same empirical Bayes approach used for the indirect prediction method. While the Bayesian sets may be smaller than the direct or indirect sets, the Bayesian sets are not α -valid, and thus the frequentist coverage of a set for any given county may fall below the nominal level.

As conveyed in the cardinality comparison between the indirect and direct prediction sets, incorporating auxiliary information in the construction of the prediction sets results in improved cardinality overall. Specifically, the Bayesian prediction sets are smaller than the direct sets in 59 of the 99 counties in North Carolina. In comparing the two methods that utilize auxiliary information, the Bayesian sets are smaller than indirect sets in 33 counties, and they have the same cardinality in 15 counties. Overall, the indirect method outperforms both the Bayesian method and the direct method in terms of minimizing prediction set cardinality. Moreover, the Bayesian approach is not α -valid, and thus the frequentist coverage rates of the Bayesian prediction sets may fall below the nominal level.

Discussion

Species abundance data collected across heterogeneous areas is increasingly important in understanding biodiversity. Some of the largest sources of such data are citizen science databases for which volunteers spearhead the data collection. As a result of the civilian-led scientific effort, such data often feature unequal sampling across a spatial domain where some areas have large within-area sample sizes and others have much smaller within-area sample sizes.

In this article, we propose summarising species abundance data of this type with valid prediction sets that are constructed by sharing information across areas. Utilizing indirect information may result in smaller prediction sets than otherwise achievable with direct methods. Meanwhile, maintaining validity of the prediction sets for each area allows for an accessible interpretation that enables a straightforward comparison across areas. In particular, maintaining interpretable statistical guarantees on a descriptor of such data is important as analyses from such data often have far reaching policy implications. Smaller prediction sets may be attainable based on Bayesian inference of a spatial

hierarchical model such as that presented in Tang et al. [2023], for example, but these approaches introduce bias and a resulting prediction set would not retain the nominal frequentist coverage rate guarantee for each county.

The usefulness of our approach for summarising citizen science data is motivated in part to combat the common problem of varying sampling efforts across areas. We detail how α -valid prediction sets can be constructed with the incorporation of indirect information to improve within-county prediction set precision and propose an empirical Bayes procedure to do so. Incorporation of accurate indirect information results in a narrower prediction set for a given county than a direct prediction set by exploiting data in nearest neighboring counties. The proposed empirical Bayes procedure is based on a standard hierarchical model that is straightforward to understand, and the authors provide code for implementation.

There may, however, be a benefit to utilizing a more structured prior that incorporates indirect information in a more complex manner such as a prior that weights data from different parts of the state differently. For example, a model based on a learned intrinsic distance between counties was shown in Christensen and Hoff [2022] to fit a subset of the eBird data better than standard methods based on geographic adjacency structure. In the sample analyzed in Section 4, we found an indirect prediction set constructed with a hyperparameter estimated from five nearest neighbors results in overall narrower prediction sets than a direct approach, but it would be valuable to explore if this can be further improved upon with a more detailed prior. More broadly, different applications may warrant an alternative information sharing prior if, for example, there is no notion of spatial distance across the different areas. For example, it may be of interest to compare species abundance variation across different time frames for a given county.

All replication codes for this article, including functions to implement the empirical Bayes estimation procedure for the prior hyperparameter, are available at https://github.com/betsyberesson/FreqPredSets_Indirect.

Appendix A: Proofs

Remark 1 (Concerning Theorem 1.) *We first elaborate on the construction of an order-based prediction set following Equation 3. To test if an element $\mathbf{y}^{(k)}$ in the sample space \mathcal{Y} is included in a prediction set for a given vector \mathbf{o} and known event probability vector $\boldsymbol{\theta}$, the cumulative sum of event probabilities of the categories corresponding to the minimum element of \mathbf{o} up to element k , following the ordering of \mathbf{o} , is computed. If this cumulative sum is greater than the error rate α , then element k is included in the prediction set. As a result, all elements with such cumulative sums greater than α are included in the prediction set. The elements with such cumulative sums less than or equal to α are not included. Therefore, by construction, $P(\mathbf{Y} \notin A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}} | \boldsymbol{\theta}) \leq \alpha$. Consequently,*

$$P(\mathbf{Y} \in A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}} | \boldsymbol{\theta}) = 1 - P(\mathbf{Y} \notin A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}} | \boldsymbol{\theta}) \geq 1 - \alpha,$$

so $A_{\alpha}^{\boldsymbol{\theta}, \mathbf{o}}$ is α -valid.

Proof of Theorem 1. (1): Note first that \mathbf{o} enters Equation 3 only through the ordering of its elements. As such, without loss of generality, consider vectors of the form $\mathbf{o} \in \{0, 1\}^K$. When constructing a prediction set following Equation 3 based on such a vector \mathbf{o} , the category space is effectively divided into two disjoint subsets,

$$\begin{aligned}\mathcal{Y}_0 &= \{\mathbf{y}^{(k)} \in \mathcal{Y} : o_k = 0\} \\ \mathcal{Y}_1 &= \{\mathbf{y}^{(k)} \in \mathcal{Y} : o_k = 1\},\end{aligned}$$

such that, by construction,

$$A_{\alpha}^{\theta, \mathbf{o}} = \begin{cases} \mathcal{Y} & \text{if } \sum_{\{j: \mathbf{y}^{(j)} \in \mathcal{Y}_0\}} \theta_j > \alpha \\ \mathcal{Y}_1 & \text{else} \end{cases}.$$

For a given error rate α , clearly any α -valid prediction set may be constructed under considerations of permutations of the vector \mathbf{o} of the form $\mathbf{o} \in \{0, 1\}^K$.

(2): We wish to show no other ordering results in an α -valid prediction set with strictly smaller cardinality than \mathbf{o}^{θ} . In words, consider switching the ordering of one element at a time. This will always result in a prediction set with the same cardinality or greater cardinality than that under \mathbf{o}^{θ} . More formally, let $\tilde{\mathbf{o}} = \{o_1^{\theta}, o_2^{\theta}, \dots, o_{k^*}^{\theta}, o_{k^*-1}^{\theta}, o_{k^*+1}^{\theta}, \dots, o_K^{\theta}\}$, that is, equivalent to \mathbf{o}^{θ} with the (k^*) th and $(k^* - 1)$ th ordering flipped. But, by construction, $\theta_{k^*} \geq \theta_{k^*-1}$, so $|A_{\alpha}^{\theta, \tilde{\mathbf{o}}}| \leq |A_{\alpha}^{\theta, \mathbf{o}^{\theta}}|$. \square

Proof of Theorem 2. Let $\mathbf{Y}_1, \dots, \mathbf{Y}_{N+1} \sim \text{i.i.d. } MN_K(\boldsymbol{\theta}, 1)$. Then, we wish to construct a conformal prediction set for \mathbf{Y}_{N+1} based on an observation of $\mathbf{X} = \sum_{i=1}^N \mathbf{Y}_i$, and some conformity measure C .

To determine if a candidate category $k \in \{1, \dots, K\}$ is included in a $1 - \alpha$ conformal prediction set, the conformal algorithm proceeds as follows, see Section 2.1 of Bersson and Hoff [2024] for more details:

1. Set $\mathbf{y}_{N+1} = \mathbf{y}^{(k)}$ where $\mathbf{y}^{(k)}$ is a vector of length K with a 1 in the k^{th} index and 0s elsewhere.
2. For $j = 1, \dots, N + 1$, compute conformity scores $c_j = C(\mathbf{x} - \mathbf{y}_j + \mathbf{y}^{(k)}, \mathbf{y}_j)$.
3. Set

$$p_k = \frac{|\#\{j \in \{1, \dots, N + 1\} : c_{N+1} \geq c_j\}|}{N + 1}$$

More compactly, and by symmetry in the problem, this conformal p -value may be equivalently written as:

$$p_k = \sum_{l=1}^K \mathbb{1}(o_k^* \geq o_l^*) \frac{x_l + y_l^{(k)}}{N + 1},$$

where $o_k^* = c_{N+1}$ and $o_l^* = c_j$ for a $j \in \{1, \dots, N\}$ such that $\mathbf{y}_j = \mathbf{y}^{(l)}$. Then, for prediction mis-coverage rate α , the category k is included in the prediction set if $p_k > \alpha$. A prediction set constructed from this procedure may be concisely written as follows:

$$A_{\alpha}(\mathbf{x}) = \left\{ \mathbf{y}^{(k)} \in \mathcal{Y} : \left[\sum_{l=1}^K \mathbb{1}(o_k \geq o_l) \frac{x_l + y_l^{(k)}}{N + 1} \right] > \alpha \right\},$$

for some $\mathbf{o} \in \mathbb{R}^K$.

\square

Appendix B: Maximization of the marginal multinomial-Dirichlet likelihood

In this section, we detail a Newton-Raphson algorithm to maximize the log marginal likelihood of a conjugate multinomial-Dirichlet model, sometimes referred to as the Dirichlet-multinomial compound distribution:

$$\begin{aligned}\mathbf{X}_j &\sim MN_K(\boldsymbol{\theta}_j, N_j), \text{ independently for } j = 1, \dots, J \\ \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_J &\sim \text{Dirichlet}_K(\boldsymbol{\gamma}).\end{aligned}$$

The log likelihood of the marginal likelihood is as follows,

$$\begin{aligned}\mathcal{L}(\boldsymbol{\gamma}) &\propto \sum_{j=1}^J \left[\log \Gamma\left(\sum_{i=1}^K \gamma_i\right) - \log \Gamma(N_j + \sum_{i=1}^K \gamma_i) + \right. \\ &\quad \left. \sum_{i=1}^K \log \Gamma(x_{j,i} + \gamma_i) - \sum_{i=1}^K \log \Gamma(\gamma_i) \right].\end{aligned}$$

Define

$$\Psi(s) = \frac{d}{ds} \log \Gamma(s) = -\xi + \sum_{n=0}^{\infty} \left[\frac{1}{n+1} - \frac{1}{n+s} \right],$$

where ξ is the Euler-Mascheroni constant. Then, it is straightforward to obtain the first and second derivatives of the marginal log likelihood,

$$\begin{aligned}\frac{d}{d\gamma_k} &= \sum_{j=1}^J \left[\Psi\left(\sum_{i=1}^K \gamma_i\right) - \Psi\left(N_j + \sum_{i=1}^K \gamma_i\right) + \Psi(x_{j,k} + \gamma_k) - \right. \\ &\quad \left. \Psi(\gamma_k) \right] \\ \frac{d}{d\gamma_k^2} &= \sum_{j=1}^J \left[\Psi'\left(\sum_{i=1}^K \gamma_i\right) - \Psi'\left(N_j + \sum_{i=1}^K \gamma_i\right) + \Psi'(x_{j,k} + \gamma_k) - \right. \\ &\quad \left. \Psi'(\gamma_k) \right] \\ \frac{d}{d\gamma_k d\gamma_{k'}} &= \sum_{j=1}^J \left[\Psi'\left(\sum_{i=1}^K \gamma_i\right) - \Psi'\left(N_j + \sum_{i=1}^K \gamma_i\right) \right],\end{aligned}$$

where Ψ' is the trigamma function. Let \mathbf{g} be the gradient vector of length K and \mathbf{H} the Hessian matrix. Finally, Newton's method updates $\boldsymbol{\gamma}$ as follows:

$$\boldsymbol{\gamma}^{(t+1)} = \boldsymbol{\gamma}^{(t)} - \mathbf{H}^{-1}(\boldsymbol{\gamma}^{(t)}) \mathbf{g}(\boldsymbol{\gamma}^{(t)}),$$

where the algorithm is iterated until convergence.

Solving for the maximum likelihood estimate (MLE) of the unknown parameter in the Dirichlet-multinomial compound distribution is covered in literature dating back to at least Mosimann [1962], and it is stated in Wallach [2008] that the likelihood is concave. For a detailed overview of the Dirichlet-multinomial compound distribution, see Ng et al. [2011]. For derivations and comparisons of various algorithms to solve the maximization problem, including computational expense and initialization sensitivity, we refer the reader to Minka [2000] and Wallach [2008].

Data availability

The eBird project is managed by the Cornell Lab of Ornithology, and the raw data is publicly available online at <https://science.ebird.org/en/use-ebird-data>. The sample analyzed in Section 4 is available at https://github.com/betsybersson/FreqPredSets_Indirect/tree/main/data.

Competing interests

No competing interest is declared.

Author contributions statement

E.B. and P.H. conceived the experiments, E.B. conducted the experiments, E.B. and P.H. analysed the results. E.B. and P.H. wrote and reviewed the manuscript.

References

- Z. J. Arnold, S. J. Wenger, and R. J. Hall. Not just trash birds: Quantifying avian diversity at landfills using community science data. *PLoS ONE*, 16(9), 9 2021.
- E. Bersson and P. D. Hoff. Optimal conformal prediction for small areas. *Journal of Survey Statistics and Methodology*, 4 2024.
- G. Camerini and R. Groppali. Landfill restoration and biodiversity: A case of study in Northern Italy. *Waste Management and Research*, 32(8):782–790, 2014.
- M. F. Christensen and P. D. Hoff. A flexible and interpretable spatial covariance model for data on graphs. *Environmetrics*, 7 2022.
- G. D. Faulkenberry. A method of obtaining prediction intervals. *Journal of the American Statistical Association*, 68(342):433–435, 1973.
- A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*. CRC Press, Taylor & Francis Group, Boca Raton, FL, third edit edition, 2014.
- P. Hoff. Bayes-optimal prediction with frequentist coverage control. *Bernoulli*, 29(2):901–928, 5 2023.
- J. J. Lebrun, W. E. Thogmartin, and J. R. Miller. Evaluating the ability of regional models to predict local avian abundance. *Journal of Wildlife Management*, 76(6):1177–1187, 8 2012. ISSN 0022541X. doi: 10.1002/jwmg.374.
- S. R. Lele. How should we quantify uncertainty in statistical inference? *Frontiers in Ecology and Evolution*, 8, 3 2020. ISSN 2296701X. doi: 10.3389/fevo.2020.00035.
- T. P. Minka. Estimating a Dirichlet distribution. Technical report, 2000.
- J. E. Mosimann. On the compound multinomial distribution, the multivariate Beta-distribution, and correlations among proportions. *Biometrika*, 49(2):65, 1962.
- K. W. Ng, G.-L. Tian, and M.-L. Tang. Dirichlet–multinomial distribution. In *Dirichlet and Related Distributions : Theory, Methods and Applications*, volume 895, chapter 6, pages 199–225. John Wiley & Sons, Incorporated, Sussex, UK, 1 edition, 5 2011.
- J. N. K. Rao and I. Molina. *Small Area Estimation*. John Wiley and Sons, Inc., New York, NY, 2 edition, 2015.
- D. F. Shanahan and H. P. Possingham. Predicting avian patch occupancy in a fragmented landscape: Do we know more than we think? *Journal of Applied Ecology*, 46(5): 1026–1035, 10 2009. ISSN 00218901. doi: 10.1111/j.1365-2664.2009.01694.x.
- B. Sullivan, C. Wood, M. Iliff, R. Bonney, D. FInk, and S. Kelling. eBird: a citizen-based bird observation network in the biological sciences. *Biological Conservation*, 142: 2282–2292, 2009.
- B. L. Sullivan, T. Phillips, A. A. Dayer, C. L. Wood, A. Farnsworth, M. J. Iliff, I. J. Davies, A. Wiggins, D. Fink, W. M. Hochachka, A. D. Rodewald, K. V. Rosenberg, R. Bonney, and S. Kelling. Using open access observational data for conservation action: A case study for birds. *Biological Conservation*, 208:5–14, 4 2017.
- B. Tang, J. S. Clark, P. P. Marra, and A. E. Gelfand. Modeling community dynamics through environmental effects, species interactions and movement. *Journal of Agricultural, Biological, and Environmental Statistics*, 28 (1):178–195, 3 2023. ISSN 15372693. doi: 10.1007/s13253-022-00520-3.
- A. R. Thatcher. Relationships between Bayesian and confidence limits for predictions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2):176–210, 1964.
- Q. Tian, D. J. Nordman, and W. Q. Meeker. Methods to compute prediction intervals: A review and new results. *Statistical Science*, 37(4):580–597, 11 2022. ISSN 21688745. doi: 10.1214/21-STS842.
- D. J. Twedt, J. M. Tirpak, D. T. Jones-Farrand, F. R. Thompson, W. B. Uihlein, and J. A. Fitzgerald. Change in avian abundance predicted from regional forest inventory data. *Forest Ecology and Management*, 260(7):1241–1250, 8 2010. ISSN 03781127. doi: 10.1016/j.foreco.2010.07.027.
- V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic Learning in a Random World*. Springer US, 2005.
- H. M. Wallach. *Structured Topic Models for Language*. University of Cambridge, 2008.

Elizabeth Bersson. Elizabeth Bersson is a Ph.D. candidate of Statistical Science at Duke University.

Peter D. Hoff. Dr. Peter D. Hoff is a Professor of Statistical Science at Duke University.