

A method for empirically assessing small area estimators via bootstrap-weighted k-nearest-neighbor artificial populations, with applications to forest inventory

Grayson W. White^{1,2,*}, Jerzy A. Wieczorek³, Zachariah W. Cody³, Emily X. Tan³, Jacqueline O. Chistolini³, Kelly S. McConville⁴, Tracey S. Frescino⁵, Gretchen G. Moisen⁵

¹Department of Mathematics and Statistics, Reed College, 3203 Woodstock Blvd, Portland, OR 97202, United States

²Department of Forestry, Michigan State University, 480 Wilson Rd, East Lansing, MI 48824, United States

³Department of Statistics, Colby College, 5841 Mayflower Hill, Waterville, ME 04901, United States

⁴Dominguez Center for Data Science, Bucknell University, One Dent Dr, Lewisburg, PA 17837, United States

⁵Rocky Mountain Research Station, USDA Forest Service, 4919 S 1500 W, Riverdale, UT 84405, United States

*Corresponding author: Department of Mathematics and Statistics, Reed College, 3203 Woodstock Blvd, Portland, OR 97202, United States. E-mail: gwhite@reed.edu

Abstract

National Forest Inventories monitor forest attributes across a variety of spatial and temporal scales in a given country. Increased interest in reporting and management at smaller scales has driven National Forest Inventories to investigate and adopt small area estimation (SAE) due to the promise of increased precision at these scales. However, comparing and evaluating SAE models for a given application is inherently difficult. Typically, many areas lack enough data to check unit-level modeling assumptions or to assess unit-level predictions empirically; and no ground truth is available for checking area-level estimates. Design-based simulation from artificial populations can help with each of these issues, but only if the artificial populations realistically represent the application at hand and are not built using assumptions that inherently favor one SAE model over another. In this paper, we borrow ideas from random hot deck, approximate Bayesian bootstrap, and k nearest neighbor imputation methods to propose a k nearest neighbor-based approximation to approximate Bayesian bootstrap, for generating an artificial population when rich unit-level auxiliary data are available. We introduce diagnostic checks on the process of building the artificial population, and demonstrate how to use it for design-based simulation studies to compare and evaluate SAE models, using real data from the Forest Inventory and Analysis program of the United States Department of Agriculture Forest Service (the National Forest Inventory of the United States).

Keywords: small area estimation; simulated populations; approximate Bayesian bootstrap; national forest inventory; k nearest neighbors; nonparametric methods

Introduction

In recent years, National Forest Inventories (NFIs) have been tasked with providing precise estimates of forest attributes at progressively finer spatial and temporal scales. For example, in the United States (US), the United States Department of Agriculture (USDA) 2014 and 2018 Farm Bills ([Agricultural Act of 2014](#), 2014, [Agriculture Improvement Act of 2018](#), 2018) direct the Forest Inventory and Analysis (FIA) program (<https://www.fs.usda.gov/research/programs/fia>) to “implement procedures to improve the statistical precision of estimates at the sub-State level” and “find efficiencies in the operations of the FIA Program under section 3(e) of the Forest and Rangeland Renewable Resources Research Act of 1978 (16 U.S.C. 1642(e)) through the improved use and integration of advanced remote sensing technologies to provide estimates for State- and national-level inventories, where appropriate.” The FIA program provides data and tools for delivering standardized estimates at state levels over 5–10 year periods, but there is currently

no standardized methodology to generate estimates at these sub-State levels or shorter time intervals. In response to the Farm Bill initiatives, FIA is working toward a strategic plan to define FIA user needs, evaluate small area estimation (SAE) methodology, and deliver statistically sound, standardized estimates to users at the sub-State level.

Motivation for producing estimates at sub-State levels extends beyond congressional directives, as producing estimates of forest attributes at these finer scales is of the utmost importance for management and ecological monitoring given global climate change and our dependence on natural resources. Specifically, the need for more precise information regarding forest attributes at sub-State levels is essential for field specialists and managers to make tactical assessments for local and broad-scale monitoring and planning. With these needs in mind, combined with the availability of high-quality auxiliary data from satellite imagery, climate records, elevation models, etc., scientists and managers

Handling editor: Dr. Fabian Fasnacht

Received 16 May 2025; revised 7 September 2025; accepted 7 October 2025

© The Author(s) 2025. Published by Oxford University Press on behalf of the Institute of Chartered Foresters.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

alike are turning to SAE methods, which are designed to address the limitations of small sample sizes in unplanned domains, to better understand forest landscapes across the globe.

The increased interest in and use of SAE methods from NFIs and their users has the potential for substantial benefits to forest management and monitoring. However, much caution must be taken when utilizing SAE methods, especially those of the model-based variety. This is because model-based methods do not have the same statistical properties as many design-based methods (e.g. design-unbiasedness). Furthermore, caution must be taken when implementing SAE methods from the design-based paradigm that only exhibit asymptotic statistical properties (e.g. asymptotic design-unbiasedness) such as model-assisted methods including the post-stratification and generalized regression estimators. For model-based estimators, and for design-based estimators with only asymptotic statistical guarantees, it is imperative to empirically evaluate estimators to ensure accurate and unbiased estimation. SAE methods may produce estimates and estimated uncertainty intervals that appear precise for some forest parameter of interest in a small region, but these estimates are only reliable if the underlying model has met the necessary modeling assumptions. Thus, without sufficient evaluation of modeling assumptions, SAE methods can lead decision-makers astray. Further, countless SAE methods exist for estimating a particular forest parameter of interest in a particular small area, but understanding the best SAE method to use for each particular NFI application is still an open question. Furthermore, empirically comparing and evaluating SAE models is an inherently challenging problem (Brown et al., 2001, Rao and Molina, 2015, Dorfman, 2018, Tzavidis et al., 2018). While there exist many useful model-comparison tools for unit-level predictive models (for instance, cross-validation, or information criteria such as AIC and BIC), these tools are often of limited use for SAE. Sometimes we do not have unit-level data, only area-level. Even when we do have unit-level data, what we really want to assess are the area-level estimates, not the unit-level predictions themselves. In both cases, area-level direct estimates from the sample itself are not precise enough to be treated as ground truth for model comparisons; we cannot rely on checking whether the model-based estimate is close to the direct estimate for each area.

Dorfman (2018) lists a “Variety of Inadequate Methods of Evaluation” for SAE, including “(7) large scale *simulation studies* from administrative, census or large samples—these can give useful insights but satisfactory extrapolation to the case at hand has to be assumed[...].” From the design-based perspective, such simulation studies are most straightforward to justify if they consist of repeated sampling from a real and relevant population (Lehtonen and Veijanen, 2009, Tzavidis et al., 2018), but a complete real population is rarely available. Alternately, simulation studies can consist of repeated sampling from an *artificial* population (Alfons et al., 2011, Wieczorek et al., 2012, Wieczorek and Franco, 2013, Templ et al., 2017), but this requires extra care to justify why the artificial population resembles the real one in the relevant sense: How can we trust that model assessment across samples drawn from the artificial population should be informative about model performance on our *real* sample from the *real* population?

In the present paper, we describe a novel approach to generating one or more artificial populations: a “kNN-based approximation to the approximate Bayesian bootstrap” (KBAABB). We justify why KBAABB can allow for “satisfactory extrapolation” from simulation studies. We illustrate how to use our artificial population for model evaluation and comparison, following approaches and metrics similar to those in Dorfman (2018) and

Wieczorek et al. (2012). The objectives of this study were (1) to develop a novel nonparametric method for generating artificial populations, KBAABB; (2) to demonstrate empirical diagnostics for assessing the properties of artificial populations generated by KBAABB; and (3) to illustrate how KBAABB can be used to compare SAE models on a real dataset.

Section 2.1 introduces the specific FIA data and auxiliary variables used for our analyses. Briefly, in our setting we have complete auxiliary data, denoted X , for the entire population of interest, the M333 Northern Rocky Mountain Forest-Steppe-Coniferous Forest-Alpine Meadow Province, but a much smaller survey sample of the joint auxiliary and response variables at each FIA plot, which we denote (X, Y) . To generate a complete artificial population, we wish to impute response values from the observed survey data (donor dataset) to every row of the full-population auxiliary data (recipient dataset). We introduce methodology to generate the artificial population in Section 2.2. We begin with the approximate Bayesian bootstrap (ABB), a multiple imputation tool that approximates the process of drawing from a posterior distribution for the missing data given the observed data (Rubin and Schenker, 1986). We use bootstrap-weighted sampling with k nearest neighbor (kNN) to implement KBAABB, a computationally cheaper approximation of ABB. The resulting artificial population can be interpreted as one draw from a posterior distribution for the population’s response values. Alternatives to KBAABB are discussed below in this Introduction.

A single run of KBAABB produces one artificial population, from which we repeatedly draw samples using the FIA’s quasi-systematic sampling design. This allows us to carry out design-based simulations in a finite-population framework. In addition, repeated runs of KBAABB could easily produce multiple artificial populations, allowing for simulations under a superpopulation framework (Isaki and Fuller, 1982) and/or Bayesian posterior inference, although we do not do so in this paper. In Section 3.1 we assess the qualities of the artificial population generated with KBAABB from FIA data. Empirically, we show that the output of KBAABB on our FIA dataset is a reasonable population to compare small area estimators on. Further, in [Supplementary material S.1](#) we show the artificial population is comparable with or better than artificial populations built via other imputation methods commonly used in practice: single nearest neighbor (NN) imputation, and unweighted kNN imputation where each donor is selected uniformly at random from a donor pool of the recipient’s k closest neighbors (Andridge and Thompson, 2023). KBAABB avoids the worst-case risks of single NN and the *ad hoc* nature of unweighted kNN.

Our artificial population is designed to have the correct marginal distribution of X and a reasonable estimate of the conditional distributions of each $Y|X$, and our samples from the artificial population accurately represent the real sampling design. However, we do not claim that the area-level means from the artificial population are correct for the real population. Rather, we will use the artificial population to seek SAE methods that tend to produce good estimates—close to the “artificial truth” of the artificial population’s area-level means—when fitted to samples drawn from the artificial population. To demonstrate the use of our methodology in practice, Section 3.2 evaluates a set of common SAE models for estimation of average basal area per acre on the design-based samples taken from the artificial population generated with KBAABB for M333. We assess estimator quality from insights taken from the KBAABB artificial population, finding substantial differences in estimator performance in this FIA setting. This provides a valuable example




Figure 1. Simulation study area.

of the use case of KBAABB in practice for forest managers, SAE practitioners, ecologists, and other like-minded researchers. Finally, we conclude our study with Section 4, a discussion of the implications of our findings, and future work.

By contrast with KBAABB, other work on realistic artificial populations for design-based simulations typically relies only on the survey dataset itself, or on the survey data combined with population summaries of the auxiliary variables rather than a complete unit-level dataset. Alfons *et al.* (2011) generate fully synthetic artificial populations through a combination of weighted sampling from the real survey data and model-based simulation from parametric models fit to the real survey data. Wieczorek and Franco (2013) treat a large sample as the full artificial population and mimic the real sampling design to take smaller subsamples from it for the purpose of evaluating SAE models. Beyond the survey dataset alone, if full-population marginal totals or cross-tabulations for the auxiliary data are also available, e.g. tables from a recent census, Templ *et al.* (2017) summarize approaches to generating an artificial population by resampling entire rows from the survey data until the artificial population meets these marginal or cross-tabulated constraints. However, none of the above approaches make use of the rich unit-level auxiliary data that are available in our FIA setting.

Another approach is to make up a parametric model (or fit one to the sample); make a prediction $\hat{Y}(X_i)$ for each row of the X s; and add random noise to each prediction to generate the artificial population's Y s (Morris *et al.*, 2014). However, under this approach, the artificial population would inherently favor SAE models of the same parametric form. If some of the models we wish to evaluate are correct for the artificial population but mis-specified for the real population, then simulation studies on the artificial population would be overly optimistic for these models. Instead, KBAABB's nonparametric kNN-based matching ensures that the artificial population will not have a bias toward being fit well by any of the common parametric SAE models, yet our imputed Y values will still have realistic conditional distributions for $Y|X$. Of course, if we also planned to evaluate the use of kNN-based models for small area estimators, then we would want to choose a different non-kNN model to generate our artificial population (Baffetta *et al.*, 2009).

Methods

Data

Our simulation study area is the M333 Northern Rocky Mountain Forest–Steppe–Coniferous Forest–Alpine Meadow Province from the US Forest Service's national hierarchical framework of ecological units (ECOMAP, 1993, Cleland *et al.*, 2007). EcoMap provinces are delineated by climatic zones and broad vegetation types with further delineations into sections and subsections, representing different geomorphic and topographic features (ECOMAP, 1993). The M333 province stretches across western Montana, through Northern Idaho, to eastern Washington and is about 98 700 km² in size. The province encompasses four mountainous sections, with an average of six subsections each, 23 in total. These are our 23 domains (Fig. 1).

The FIA program surveys a sample of field plot locations across the US based on a hexagonal grid. Each plot samples ~2.5 acres (1.0 ha) of land and represents ~6000 acres (2428 ha) of land within each hexagon (Bechtold and Patterson, 2005, McRoberts *et al.*, 2005). Attributes are measured and observed at each survey plot location and stored in FIA's national database (Burrill *et al.*, 2021). We extracted data from FIA's database (last updated 8 July 2021) on 20 July 2021. There are 3946 survey plots across the 23 domains, with 98% of the plots measured between 2010 and 2019. From these plots, we compile our Y -variables and associated X -variables.

Our Y -variables were all aggregated to the plot level and calculated on a per-acre basis. These variables were basal area (BA), net cubic-foot volume (VOLCFNET), net board-foot volume (VOLBFNET), number of trees (COUNT), and aboveground dry biomass (DRYBIO). The VOLCFNET response variable measures the average net volume of trees per acre, where the VOLBFNET response variable measures the average net volume of usable timber per acre. See Table 1 for further details.

We also have auxiliary X -variables from satellites, climate records, and digital elevation models that were thought to relate to and discriminate our Y variables (Table 2). The auxiliary information includes surface reflectance images, including spectral indices and classified imagery to get a picture of vegetation on the ground (Schwager and Berg, 2021) and both broad-scale and localized climatic characteristics to understand ecological

Table 1. Y-variables or survey variables recorded by field crews at each surveyed plot.

Variable	Units	Description
BA	square feet/acre	Basal area of live trees
VOLCFNET	cubic feet/acre	Net cubic-foot volume in sawlog portion of a live sawtimber tree
VOLBFNET	board feet/acre	Net board-foot volume in sawlog portion of a live sawtimber tree
COUNT	number of trees/acre	Number of live trees per acre unadjusted
DRYBIO	tons/acre	Aboveground biomass of live trees (Heath et al., 2008)

tolerances. All auxiliary data were resampled from their original resolution to 90x90m pixels. For each plot, we have the auxiliary data corresponding to that plot's location. Although more auxiliary variables could have been used, our kNN matching steps only used the categorical variable `wc2c1` for stratification and the quantitative variables in Table 2 for matching. Some of our analyses in Supplementary material S.3 used `tnt` for post-stratification.

As a basis for the artificial population, besides the survey sample we also have complete auxiliary-only data for the full population: 11 752 067 pixels in the 4286 hexagons that fall within or overlap with EcoProvince M333. The largest hexagon had 2979 pixels, and the median was 2964, although several hexagons along the edge of the province had just a few pixels in the target population. For each of these auxiliary-only pixels, we had the same X variables as for the 3946 pixels corresponding to the real survey plots.

We use an analysis function in the `FIESTAnalysis` R package (which is an extension of the `FIESTA` R package), `anGetPixelDat()`, to generate pixel-level data tables for each subsection (Frescino and White, 2022, Frescino et al. 2023a). The function creates a table of a combination of longitude and latitude extracted from each 90x90m pixel across each subsection, including a unique identifier of each pixel. This table is then used to extract values of our auxiliary information to create a set of data for our simulation models.

Implementation of KBAABB

Artificial population generation

With rich auxiliary data for the entire population, we can treat the unobserved response values in our population as missing data that need to be imputed. Several popular approaches to imputation of missing data are based on NN. In the single NN approach, each recipient row is matched to the “nearest” donor row—the donor row whose auxiliary variables are most similar to the recipient's by some distance metric—and that donor's response variables are imputed to that recipient. However, this approach risks overusing the same donor for all recipients with similar X values, which could result in unrealistically low variability in the imputed Y values (Andridge and Thompson, 2023). For the same reason, imputing the means or medians of larger donor pools (Crookston and Finley, 2008) is also unsuitable for our purposes.

An alternative is to use a recipient's k nearest neighbors (kNN) as the donor pool, for a k such as 5, 10, or 20 (Morris et al., 2014), and then to choose one case in this donor pool *uniformly at random* to be imputed to the recipient (Andridge et al., 2021). This avoids single NN's problem of always imputing the same donor—but

it can also lead to bias in the imputed conditional distributions of $Y|X$, since a larger donor pool may include Y values that are not realistic for the recipient's X values. It also requires careful choice of k . (In the extreme case, $k = n$ would replace every imputed conditional distribution of $Y|X$ with an imputed marginal distribution of Y.)

Our proposed method, KBAABB, works like kNN imputation but with uniform random sampling of donors replaced by weighted sampling, with the weights based on bootstrap selection probabilities. The j th NN is selected with probability $p_b \times (1 - p_b)^{j-1}$, where $p_b = 1 - e^{-1} \approx 0.632$ is the asymptotic probability of an observation being selected into a bootstrap sample (Andridge and Thompson, 2023). Since these probabilities become negligible quickly, in practice we only find each recipient's $k = 10$ NNs to use as its donor pool, then use bootstrap-weighted sampling from this pool to select one donor for that recipient. See inner loop of Algorithm 1.

Furthermore, to reduce bias and variance, KBAABB can be applied independently within strata; and to avoid sensitivity to units of measurement, the auxiliary variables can be centered and scaled before calculating distances to find the k NNs for each recipient. See outer loop of Algorithm 1. By default we use Euclidean distance; alternatives are discussed in Section 2.2.3.

If there are several response Y variables to be imputed for each recipient, KBAABB imputes all of them from the same donor to their recipient. If instead we were to impute each response variable independently, our artificial population would not accurately reflect the associations among these variables (Nur et al., 2005, Andridge and Little, 2010).

In principle, by selecting nearer neighbors with higher probability, KBAABB should avoid the pitfalls of single NN (too little variability in Y) and kNN (bias due to large k) while also removing the need to choose tuning parameter k . Empirically, in Supplementary material S.1, we report a sensitivity analysis regarding the probabilities of selection and the choice of k , comparing KBAABB with unweighted kNN approaches on our FIA dataset. For each response variable Y, we found the same results: as kNN's k grows, the distribution of Y in each domain becomes more similar across domains, and unweighted $k = 20$ or larger makes the domains too similar to each other. Furthermore, when looking at spatial maps of the imputed Y values, the spatial patterns look realistically smooth for single NN and for KBAABB, but begin to look unrealistically noisy for unweighted $k = 5$ or larger.

Heuristic rationale for bootstrap weighting

KBAABB's use of bootstrap weights is heuristically inspired by the ABB (Rubin and Schenker, 1986). ABB is popular for multiple imputation because it can be proven to be “valid,” a technical term in the imputation literature, loosely meaning that it allows asymptotically unbiased design-based inferences under both the sampling and the nonresponse mechanisms; for details, see chapter 4 of Rubin (2004). In Bayesian inference, one run of ABB can also be treated as one draw from a posterior distribution for the population's response values, given the survey sample and the full-population auxiliary values and using a nonparametric model for the likelihood. Neither “validity” nor Bayesian inference are central to our goal of creating an artificial population, but ABB inspired us to try bootstrap weighting, and empirically we have found it to be a success in our application; see Supplementary material S.1.

Briefly, our bootstrap weights arise from mimicking true ABB, where we would first draw a bootstrap sample from the

Table 2. X-variables or auxiliary variables.

Variable	Units	Description
tcc	%	National Land Cover Dataset (NLCD) Analytical Tree Canopy Cover (Yang et al., 2018)
elev	meters	LANDFIRE 2010 DEM—elevation (Geological Survey, 2019)
eastness	(−100 to 100)	Transformed aspect degrees to eastness (Geological Survey, 2019)
northness	(−100 to 100)	Transformed aspect degrees to northness (Geological Survey, 2019)
tri	index	Terrain Ruggedness Index (Geological Survey, 2019)
tpi	index	Topographic Position Index (Geological Survey, 2019)
ppt	mm×100	PRISM mean annual precipitation—30yr normals (1991–2020) (Daly et al., 2002)
tmin01	°C×100	PRISM mean minimum temperature (Jan)—30 yr normals (Daly et al., 2002)
wc2c1	class	Eur. Space Agency (ESA) 2020 WorldCover global land cover (Zanaga et al., 2021)
tnt	class	LANDFIRE 2014 tree/non-tree lifeform mask (Rollins, 2009 , Picotte et al., 2019)

Algorithm 1 KBAABB, with stratification and standardization

```

k ← 10                                     ▷ Choose k and distance metric, or use defaults
DistMetric ← Euclidean distance

for j = 1 to k − 1 do                       ▷ Define the selection probabilities
  P(jth ranked NN selected) ← exp(−j + 1) × (1 − exp(−1))
end for
P(kth ranked NN selected) ← 1 − ∑j=1k−1 P(jth ranked NN selected)

for stratum ∈ {strata} do                 ▷ Process each stratum separately
  recipients ← {all population units ∈ stratum}
  donors ← {all sample units ∈ stratum}
  Center/scale auxiliary data in recipients and donors, using recipients means and SDs

  for unit ∈ recipients do                ▷ kNN matching and ABB imputation
    Dists ← distances to unit from all donors, via DistMetric in auxiliary data space
    NNs ← k nearest neighbors to unit from donors, according to Dists
    Rank NNs by distance to unit in auxiliary data space from j = 1 to k
    Randomly sample one element of NNs based on P(jth ranked NN selected)
    Impute all response values from selected element of NNs to unit
  end for
end for

```

complete-data cases, then draw with replacement from this bootstrap sample to impute the missing cases. In order to reduce both bias and variance, typically ABB imputes from a smaller donor pool based on some auxiliary information. If we use a single NN from the bootstrap sample to define the donor pool for each recipient, around p_b of recipients will have their NN from the original sample selected as their donor ([Andridge and Thompson, 2023](#)). However, NN with true ABB is still computationally expensive for our task of generating an artificial population, as we would need to take a new bootstrap sample of the complete-case survey data for every recipient, then redo the NN search for that recipient. KBAABB is a computational shortcut: use the original sample (not a resample) to find a donor pool of k NNs for each recipient row separately, then use bootstrap-weighted sampling to choose one donor for that recipient. Under true ABB with NN, each recipient has approximately a $(1 - p_b)^{j-1} \times p_b$ probability that the first $j - 1$ NNs would not be in the bootstrap sample but the j th would be. (This probability becomes negligible quickly, so in practice we only find the $k = 10$ NNs and select the 10th NN with probability $1 - \sum_{j=1}^9 (1 - p_b)^{j-1} \times p_b \approx 0.000124$.) Thus, KBAABB's bootstrap weights approximate the process of running NN with true ABB once per recipient.

Considerations in implementation

To implement KBAABB, as with any other kNN-based method, we must decide how to define “nearest” for the kNN search: Which variables, transformations, and distance metric should be used to rank the donor pool for each recipient? For a detailed discussion of how we chose to define “nearest” for our study, see [Supplementary material S.2.1](#). We now discuss the considerations to be made in general.

First, if the population consists of strata or post-strata for which units tend to be homogeneous within strata but quite different across strata, KBAABB will have less bias and less unnecessary variability if kNN matching is carried out separately in each stratum.

Next, kNN is known to suffer from the curse of dimensionality: if we use Euclidean distance in a high-dimensional space, the several NNs might all be very far away ([James et al., 2021](#)). When there are many auxiliary variables available for matching, NN imputation can be asymptotically biased and some analysts prefer to summarize the auxiliary variables into a univariate matching variable ([Morris et al., 2014](#), [Yang and Kim, 2019](#)). However, to keep the matching process as model-free as possible, others may prefer to simply reduce the matching space to a subset of the variables,

using subject matter expertise. Such parsimony can also reduce unnecessary variance in the matching process (D’Orazio, 2015).

Furthermore, if matching variables are heavily skewed, it may be worthwhile to transform them to be more symmetric. Otherwise, donors with distant outlier values might never be one of the NNs for any recipient. Chen and Shao (2000) show that for imputation purposes, NN imputation can be biased when the matching variables are not symmetric.

We may also center and scale each of the matching variables to have mean 0 and standard deviation (SD) 1 (within each stratum separately) on the full-population auxiliary dataset. The same centering and scaling constants (by class and variable) should be applied to the smaller survey sample dataset. This removes the effect of different variables having different units and scales, ensuring that each variable is given equivalent weight in the kNN matching process (James et al., 2021).

Finally, a distance metric must be chosen and computed between each recipient and each donor using the selected and transformed variables. We use Euclidean distance as the KBAABB default and as the distance for the case study in Section 3.2. Euclidean distance should suffice when the selected variables are not strongly correlated with each other and they have been centered and scaled. However, if necessary in other applications, alternatives such as e.g. Mahalanobis distance could account for covariances between variables (Crookston and Finley, 2008).

Sampling for design-based simulations

For each replication (“rep”) of a survey sample from the KBAABB-generated artificial population, we take a sample using the same sampling design that was used to draw the real sample. This lets us carry out design-based simulations that respect the sampling design of the actual survey data. When we fit a given SAE model form to each rep, and then summarize the model’s empirical properties across reps, these properties are reasonable stand-ins for what we could expect to see under repeated sampling from the real sampling design. For our application, we took design-based samples using FIA’s quasi-systematic sampling design (see [Supplementary material S.2.2](#)).

Case study: estimators and performance metrics

We have chosen four commonly applied estimators in the SAE literature to assess in M333, to give a sense of what can be done with KBAABB through a basic yet illustrative example, with results discussed in Section 3.2. In particular, we implemented two design-based estimators, the Horvitz–Thompson (Horvitz and Thompson, 1952) and modified generalized regression (GREG) (Woodruff, 1966, Wojcik et al., 2022), along with two model-based estimators, the area-level empirical best linear unbiased prediction (EBLUP) Fay–Herriot (Fay and Herriot, 1979) and the unit-level EBLUP Battese–Harter–Fuller (Battese et al., 1988). Model-based and -assisted estimators used auxiliary variables *tcc*, *tri*, and *elev*, which are defined in [Table 2](#).

These estimators were implemented in R using the `FIES-TAutils` function `SAest.large()` (Frescino et al. 2023b). The Horvitz–Thompson estimator was fit without weights and hence its estimate and variance estimate are the sample mean and sample variance, respectively. The modified GREG estimator’s regression coefficients were fit with ordinary least squares and the estimate and variance estimate were created based on the standard equations described in Rao and Molina (2015). The Fay–Herriot and Battese–Harter–Fuller estimators were fit with restricted maximum likelihood to ensure unbiased estimation of

the random effects. For further details regarding the estimators and their statistical properties, see Rao and Molina (2015).

We considered four evaluation metrics in this case study: relative bias, MSE, MSE ratio, and 95% confidence interval coverage. The relative bias is defined as, for a particular estimator (*i*) and domain (*j*),

$$\frac{\mathbb{E}[\hat{\mu}_{ij}] - \mu_j}{\mu_j},$$

where μ_j is the true average basal area per acre in subsection *j* as generated by KBAABB and $\mathbb{E}[\hat{\mu}_{ij}]$ is defined as follows:

$$\mathbb{E}[\hat{\mu}_{ij}] = \frac{1}{K} \sum_{k=1}^K \hat{\mu}_{ijk}, \quad (1)$$

i.e. the empirical expected value for estimator *i* of average basal area per acre in subsection *j* across all reps (*k* from 1 to *K*). The MSE for a given estimator and subsection is given by

$$\text{MSE}_{ij} = \frac{1}{K} \sum_{k=1}^K (\hat{\mu}_{ijk} - \mu_j)^2.$$

The MSE ratio is defined as the ratio of the empirical expected value of the MSE estimator and the empirical MSE for a particular estimator and area of interest,

$$\frac{\mathbb{E}[\widehat{\text{MSE}}_{ij}]}{\text{MSE}_{ij}},$$

where $\mathbb{E}[\widehat{\text{MSE}}_{ij}]$ is the empirical expected value of the MSE estimator for estimator *i* in subsection *j* across all *K* reps, defined analogously to Equation (1). The 95% confidence interval coverage is defined as, for a particular estimator and subsection, the proportion of samples across all reps where a given estimator’s 95% confidence interval contains μ_{ij} .

Results

We now turn to discussing the results of generating our artificial population and the comparison of small area estimators in estimation of average basal area per acre in Sections 3.1 and 3.2, respectively. We discuss the results and diagnostics of the artificial population generation in order to properly evaluate that our artificial population is well-suited for evaluating estimators.

Empirical diagnostics for KBAABB in our FIA setting

To check whether the artificial population generated with KBAABB looked reasonable, we compared the statistical distributions of the imputed *Y* variables to the original survey. We also looked for imbalances in how often each donor was used, and in how often each recipient’s domain matched the donor’s domain. We repeated these checks for each of the *Y*-variables and found similar results, so for brevity most examples in this section are illustrated using a single *Y* variable: average basal area per acre.

First, we confirmed that marginal distributions from the artificial population resembled the real survey data by comparing the original and imputed values of each variable using histograms or eCDFs (empirical cumulative distribution functions), as in [Fig. 2](#). Histograms are more familiar to many readers, but eCDFs do




Figure 2. Histograms (left) and empirical CDFs (right) of original vs imputed average basal area per acre.

not require us to choose a bin size and can also make it easier to compare behavior in the tails. With either graph type, we found that each Y variable's marginal distribution did not change substantially between the survey and the artificial population.

Next, we checked how well the imputed domain-level SDs match the real survey SDs in each domain. Figure 3 shows that for each Y-variable, the KBAABB artificial population's domain-level SDs are strongly but not perfectly correlated with their original-sample counterparts. This is a good sign, because if they were perfectly correlated we might simply be reproducing the original sample, while if they had low correlations we might be imputing unrealistic donors to each recipient. For a similar example using kNN with varying values of k , please see [Supplementary material S.1](#).

Next, we checked how often each of the 3946 donor rows were used (Fig. 4). Recall that for each of the roughly 12 million recipient rows, KBAABB finds its 10 NNs among the donor rows and then uses weighted sampling to choose one donor. Hence, first we checked how often each donor was included in a donor pool of some recipient's 10 NNs (left side of the figure). For both *wc2c1* classes, most donors were chosen to be in 10 000 to 50 000 donor pools, though some were chosen for as many as 69 302 donor pools, while others for as few as 1 or 0 donor pools. Every one of the 3101 donors in class 1 (tree) was in a donor pool at least once, and 811 of the 845 donors in class 2 (non-tree) were in a donor pool at least once. Next, we repeated the analysis but to see how often each donor row was actually used as a donor (right side of the figure). For both classes, most donors were actually used for 1000 to 5000 recipients, though some were used as many as 9173 times, while others were only used 1 or 0 times. Again every one of the 3101 donors in class 1 was used as a donor at least once, but 790 of the 845 donors in class 2 were used as a donor at least once. In short, only 55 of the possible donors were never used (all in *wc2c1* class 2); 34 of them were never in a donor pool, and the other 11 were in a pool but never happened to be chosen by the weighted sampling. We are not concerned about these few never-used donors, because they all came from the non-tree class, in which most of the Y-variables are all 0s. The other possible donors

were typically used a moderate number of times, with few over-used or under-used donors.

Recall that our 23 domains are “subsections,” nested within 4 “sections,” ecologically defined so that pixels should tend to be more similar within sections than across them. We cross-tabulated the domain of each donor vs. each recipient to see whether recipients tended to get donors from the same subsection; from a different subsection but still the same section; or from elsewhere (Fig. 5). Within each column (recipient), the highest frequencies tend to be in the corresponding row (donor) or another row in the same section (regions separated by black lines), showing that recipients often but not always tended to get donors from the same subsection or at least the same section. If all counts had been only on the main diagonal (each domain's imputed values all came from the same domain), we would have been concerned that the artificial population is too similar to simply taking copies of the survey sample. On the other hand, if counts were completely uniform, we would have been concerned that recipients are getting donors which do not actually resemble them. As it is, we feel comfortable seeing that our results fall between these two extremes.

Outside of these diagnostic checks, we also qualitatively assessed the spatial distribution of the artificial population generated in KBAABB compared with unweighted kNN approaches. For a discussion of this analysis, please see [Supplementary material S.1](#).

Case study: using KBAABB to evaluate SAE methods in our FIA setting

The artificial population generated with KBAABB allows for an extensive empirical evaluation of SAE methods and strategies that may be implemented by researchers and others interested in the performance of small area estimators under realistic conditions for the population and sampling design of interest. In order to evaluate such small area estimators for average basal area per acre for M333, we created 2500 samples from the artificial population, where we sampled one pixel per hexagon for each hexagon that overlaps M333, as described in [Supplementary material S.2.2](#).




Figure 3. Original vs imputed SDs of each Y-variable for each subsection, using KBAABB. All plots show $y = x$ line as a common reference.




Figure 4. Count of how often each donor row was included in a donor pool of some recipient's 10 NNs (left) or actually used as a donor (right), for $wc2cl$ classes 1 (tree, top) and 2 (non-tree, bottom).

We assess the small area estimators described in Section 2.3 based on their estimates and MSE estimates of average basal area per acre within each subsection in M333.

In order to assess the metrics described in Section 2.3, we turn to Fig. 6. Each subfigure displays a metric for each subsection

and estimator combination; in particular subfigure A displays the relative bias, subfigure B displays the MSE, subfigure C displays the MSE ratio, and subfigure D displays the confidence interval coverage. For the two design-based estimators, the relative bias remains low, as we would expect. However, for the model-based




Figure 5. Count of how often donors from each subsection were used by recipients in each subsection. Box widths (fixed within each column) are proportional to the total number of pixels in each recipient subsection. Box heights (varying within each column) are proportional to the number of times a donor from that subsection was used for a recipient in that subsection, scaled so that boxes along the main diagonal would be squares if each conditional proportion (how often did that recipient use donors from its own subsection?) was equal to the overall proportion (31% of recipient pixels got donors from their own subsection).

estimators we see a larger spread of values for relative bias. On median, the area- and unit-level EBLUP estimators exhibit negative and positive relative bias, respectively. Further, the distribution of the area-level EBLUP's bias is primarily negative, with its interquartile range entirely following below the 0 on the y-axis. While examining the empirical relative bias can be informative, it can also be useful to look at the distribution of estimates produced over each simulation rep in order to get a fuller sense of the estimator's performance across samples. This is shown in Fig. 7. From Fig. 6 alone we might be concerned with some of the biases that the Horvitz–Thompson estimator is exhibiting, but upon examination in Fig. 7 we can see how variable the estimator is, and that it does not appear to be producing concerning bias in any subsection. Contrastingly, the model-based estimators display some more concerning characteristics: in subsections such as M333Ah (Okanogan Semi-Arid Foothills) and M333Cc (Mission–Swan Valley–Flathead River), we see the model-based estimators producing precise estimates that are consistently over- or under-estimating the parameter of interest, leading to bias in the estimator. This bias may be attributed to a variety of causes, but upon ecological inspection of the subsections, we see that, compared with the rest of the M333 province, these are lower elevation regions that likely exhibit very different qualities than the otherwise mountainous majority of M333, making it difficult for the model's random effect

to pool accurately. The modified GREG estimator is both moderately precise and exhibits less concerning distributions than the model-based approaches, although the modified GREG does still have trouble capturing the true mean in M333Ah.

Returning to Fig. 6, we see in subfigure B the MSE of model-based and -assisted estimators are competitive in this case study, with the two model-based approaches exhibiting the lowest MSE on median, but the modified GREG estimator displaying similar values. Subfigure C displays the MSE ratio, i.e. how well an estimator's MSE is estimated, and we see here the Horvitz–Thompson and modified GREG estimators with values right around 1, and the model-based estimators show an overestimation of MSE, on median. Further, in subfigure D we see the 95% confidence interval coverage looks very sensible for the Horvitz–Thompson and modified GREG estimators, and slightly more instances of over- or under-coverage for the model-based estimators. This makes sense in light of subfigure C where we see an over-estimation of the uncertainty in these estimators.

In light of Fig. 6, subfigure C, we wanted to understand if the poor MSE estimation of some estimators can be attributed to any systematic properties in the areas of interest, and to assess this we turn to Fig. 8, which displays the MSE ratio by proportion of zero-valued pixels in the population for the response variable, with points sized by MSE. The Horvitz–Thompson, modified GREG, and




Figure 6. Performance metrics of each estimator for each subsection displayed as side-by-side boxplots. Within each subfigure, from left to right the boxplots represent the Horvitz–Thompson estimator, the modified GREG estimator, the area-level EBLUP estimator based on the Fay–Herriot, and the unit-level EBLUP estimator based on the Battese–Harter–Fuller. Subfigure A displays the relative bias metric, subfigure B displays the empirical MSE, subfigure C displays the ratio of estimated and empirical MSE, and subfigure D displays the 95% confidence interval coverage rate. For subfigures A–D, the horizontal dashed line has an intercept of 0, 0, 1, and 0.95, respectively, representing the ideal value for each metric.

area-level EBLUP estimator do not exhibit a relationship between the displayed variables; however, we see that the area-level EBLUP does over-estimate the MSE for subsections with small MSEs, and under-estimates MSE for subsections with large MSEs. This is somewhat expected with the pooling that occurs when fitting a mixed model. We also see this trend for the unit-level EBLUP estimator. Strikingly though, the unit-level EBLUP estimator's MSE ratio has a strong positive linear relationship with the proportion of zero-valued pixels in the population for the response variable. We see that when the proportion of zero-valued pixels is greater than ~ 0.1 , the MSEs tend to be smaller but are also over-estimated. This is likely due to model-misspecification that can occur at the unit-level due to zero-inflation in the data. The forest inventory data used here are positive and continuous sans zeros, and we expect, as the proportion of zeros in a subsection increases, that the model will be increasingly misspecified. Zero-inflated estimators have been proposed to deal with the model-misspecification discussed here (Finley et al., 2011, White et al., 2025), but this systematic issue in estimating the MSE of small area estimates is not discussed. In order to obtain the precision of a unit-level, model-based approach without the perils of the unit-level EBLUP shown here, it may be worthwhile to consider the approaches taken by Finley et al. (2011) or White et al. (2025).

From this case study, we are able to obtain much information about the estimators considered for average basal area per acre in M333. First of all, the Horvitz–Thompson estimator appears to be performing as we expect, which is a good check that our simulation is producing sensible results. Next, the modified GREG

and area-level EBLUP both appear to be decent estimators for these data and areas of interest. Both estimators produce sensible estimates; however, in some outlying subsections we do see estimator bias appear. Finally, the unit-level EBLUP displays concerning results regarding poor and systematic over-estimation of the MSE in this setting. While the unit-level EBLUP was the most precise estimator in terms of MSE (on median, Fig. 6), its MSE estimation was unreliable. Given the choice between the four estimators, the modified GREG has been empirically shown to have the best balance of low bias, low MSEs, accurate MSE estimation, and correct confidence interval coverage on this artificial population.

Discussion

The need to properly assess small area estimators is a high priority for those making inferences and management decisions with forest inventory data. SAE methods promise precise estimates in constrained geographic and temporal regions, but merely fitting a small area model is insufficient to making well-informed inferences. We presented KBAABB to help empirically assess SAE methods. We find KBAABB to be well-suited for assessing estimators on forest inventory data in our case study, but KBAABB is suitable more generally when we have rich and complete auxiliary data for the entire population of interest. Of course, not all surveys have unit-level auxiliary data for the entire population of interest; but when they do, we argue that KBAABB's artificial populations are realistic, fine-grained, and not biased in favor of any particular




Figure 7. Distributions of estimates produced for each estimator and subsection (density curve), along with true KBAABB population parameter value (dot). The subsections are mapped to the x-axis, and the distribution estimates and true parameter values are mapped to the y-axis. Each panel corresponds to a different estimator. Subsections are ordered by true average basal area per acre from the KBAABB population, least to greatest.

SAE model, making KBAABB appropriate for “satisfactory extrapolation” from simulation studies to SAE estimator performance for the case at hand.

One direction for future work is nonresponse. FIA surveys do experience some nonresponse (for instance, land owners may not admit a survey field crew, or weather conditions may make it impossible to visit a survey plot on schedule), though less than a typical demographic survey and this nonresponse is usually treated as missing at random (MAR) within a given population (Bechtold and Patterson, 2005). So far, in generating and sampling from our artificial populations, we have assumed that nonresponse was missing completely at random (MCAR) and ignorable. However, in demographic applications we typically take extra care to ensure that imputation is used appropriately given the posited nonresponse mechanism, which is often MAR or missing not at random rather than MCAR. It should be possible to add a simulated nonresponse stage after KBAABB’s sampling stage, to mimic any nonresponse behavior seen in the real survey. Appropriate use of KBAABB for generating artificial populations when the survey data had non-MCAR nonresponse is an open question.

Another future direction is the interplay between variable selection for KBAABB matching vs. for SAE modeling. KBAABB is not biased toward any of the common parametric SAE model forms. However, it remains to be seen whether the choice of variables for KBAABB matching may bias the performance of SAE models, when comparing models whose explanatory variables do vs. do not include the matching variables. For the present paper, we side-step this issue by using subject-matter expertise to choose the KBAABB matching variables, and by only using

a subset of those variables in the SAE models in our case studies.

We have also not touched on a few complications specific to the FIA data: subplots and condition classes, data vintage, or spatial smoothing. Briefly, condition classes and subplots are more fine-grained levels of detail about the survey data collected at each plot. Our artificial population has been generated only at the plot level so far, and we may extend our methods to work at these other, smaller levels. Meanwhile, the issue of data vintage refers to accounting for the date at which each variable was measured. The survey data come from a panel design, in which different plots are sampled in different years; and the surveyed pixels’ X variables may not be from the same year as the corresponding Y . Although most of the X s change very slowly if at all from year to year (such as elevation or precipitation), others might be more timely (such as `tcc`, `wc2c1`, or `tnt`, which should change as forests grow, become harvested, experience wildfires, etc.). Both our artificial population and our SAE models might benefit from a closer match between the vintages of X and of Y , though prior work by FIA researchers suggests that common SAE models are fairly robust to data vintage in this population. Finally, in order to assess estimators with explicit spatial components, future work will incorporate spatial smoothing to the imputed Y values after running KBAABB.

Aside from such details that are specific to our FIA data sources, we believe that KBAABB is a useful approach that could be widely applied to generate artificial populations and evaluate small area estimators for any survey where rich, unit-level auxiliary data are available.




Figure 8. Ratio of estimated MSE and empirical MSE across every domain and estimator by proportion of zero-valued pixels for the response variable in the population, sized by empirical MSE. The MSE ratio is mapped to the y-axis, and the proportion of zero-valued pixels is mapped to the x-axis. Empirical MSE is mapped to point size. The dashed horizontal line has a slope of 0 and intercept of 1, and represents perfect MSE estimation. The solid lines are least-squares regression lines.

Software

We are developing an R package *kbaabb* that provides tools for generating an artificial population via KBAABB. The package is available at <https://github.com/graysonwhite/kbaabb/>. While we cannot share the proprietary FIA data used in this paper, the R package includes a worked example based on public data.

We have also created a draft R Shiny application (Chang et al., 2022) for exploring the case study's artificial population and several estimators. The app can be accessed at <https://civilstat.shinyapps.io/fia-simpop-app/> and the code is available at <https://github.com/ColbyStatSvyRsch/FIA-simpop-app>. For details and screenshots, see [Supplementary material S.3](#).

Author contributions

Grayson W. White (Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing—original draft, Writing—review & editing), Jerzy A. Wieczorek (Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Supervision, Validation, Visualization, Writing—original draft, Writing—review & editing), Zachariah W. Cody (FormalAnalysis, Investigation, Software, Visualization), Emily X. Tan (FormalAnalysis, Software, Visualization), Jacqueline O. Chistolini (Formal analysis, Investigation, Software, Visualization), Kelly S. McConville (Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Supervision,

Validation, Writing—original draft, Writing—review & editing), Tracey S. Frescino (DataCuration, Investigation, Resources, Software, Validation, Writing—original draft, Writing—review & editing), and Gretchen G. Moisen (Project administration, Resources, Supervision)

Supplementary material

[Supplementary data](#) is available at *Forestry Journal* online.

Conflict of interest: None declared.

Funding

This work was supported by the National Council for Air and Stream Improvement (grant number FO-SFG-2673 to J.A.W.).

Data availability

The data include confidential plot data, which cannot be shared publicly. FIA data can be accessed through the FIA DataMart (<https://apps.fs.usda.gov/fia/datamart/datamart.html>). Requests for data used here or other requests including confidential data should be directed to FIA's Spatial Data Services (<https://www.fs.usda.gov/research/programs/fia/sds>).

References

- Agricultural Act of 2014. H.R.2642, 113th Congress (2013–2014), 2014. <https://www.congress.gov/113/plaws/publ79/PLAW-113publ79.pdf> (29 October 2025, date last accessed).
- Agriculture Improvement Act of 2018. H.R.2, 115th Congress (2017–2018), 2018. [https://www.agriculture.senate.gov/imo/media/doc/Agriculture Improvement Act of 2018.pdf](https://www.agriculture.senate.gov/imo/media/doc/Agriculture%20Improvement%20Act%20of%202018.pdf) (29 October 2025, date last accessed).
- Alfons A, Kraft S, Templ M. et al. Simulation of close-to-reality population data for household surveys with application to EU-SILC. *Stat Meth Appl* 2011; **20**:383–407. <https://doi.org/10.1007/s10260-011-0163-2>.
- Andridge RR, Bechtel L, Thompson KJ. Finding a flexible hot-deck imputation method for multinomial data. *J Surv Stat Methodol* 2021; **9**:789–809. <https://doi.org/10.1093/jssam/smaa005>.
- Andridge RR, Little RJ. A review of hot deck imputation for survey non-response. *Int Stat Rev* 2010; **78**:40–64. <https://doi.org/10.1111/j.1751-5823.2010.00103.x>.
- Andridge RR, Thompson KJ. Adapting nearest neighbor for multiple imputation: advantages, challenges, and drawbacks. *J Surv Stat Methodol* 2023; **11**:213–33. <https://doi.org/10.1093/jssam/smab058>.
- Baffetta F, Fattorini L, Franceschi S. et al. Design-based approach to k-nearest neighbours technique for coupling field and remotely sensed data in forest surveys. *Remote Sens Environ* 2009; **113**:463–75. <https://doi.org/10.1016/j.rse.2008.06.014>.
- Battese GE, Harter RM, Fuller WA. An error-components model for prediction of county crop areas using survey and satellite data. *J Am Stat Assoc* 1988; **83**:28–36. <https://doi.org/10.1080/01621459.1988.10478561>.
- Bechtold WA, Patterson PL. The Enhanced Forest Inventory and Analysis Program—National Sampling Design and Estimation Procedures. Gen. Tech. Report SRS-80. Asheville, NC: U.S. Department of Agriculture, Forest Service, Southern Research Station, 2005.
- Brown G, Chambers R, Heady P. et al. Evaluation of small area estimation methods—an application to unemployment estimates from the UK LFS. *Proceedings of Statistics Canada Symposium*, Hull, Quebec, 2001; **2001** Statistics Canada:1–10.
- Burrill E, DiTommaso A, Turner J., et al. *The Forest Inventory and Analysis Database: Database Description and User Guide Version 9.0.1 for Phase 2*. USDA Forest Service. 2021. <http://www.fia.fs.fed.us/library/database-documentation/> (17 August 2022, date last accessed).
- Chang W, Cheng J, Allaire J., et al. shiny: web Application Framework for R. R package version 1.7.3. 2022. <https://CRAN.R-project.org/package=shiny> (13 April 2023, date last accessed).
- Chen J, Shao J. Nearest neighbor imputation for survey data. *J Off Stat* 2000; **16**:113–131.
- Cleland DT, Freeouf JA, Keys JE. et al. *Ecological Subregions: sections and Subsections of the Conterminous United States, Presentation Scale [1:3,500,000] [CD-ROM]*. A.M. Sloan, Cartographer, Gen. Tech. Report WO-76. Washington, DC: U.S. Department of Agriculture, Forest Service, 2007.
- Crookston NL, Finley AO. yaImpute: an R package for kNN imputation. *J Stat Softw* 2008; **23**:1–16. <https://doi.org/10.18637/jss.v023.i10>.
- D'Orazio M. Integration and imputation of survey data in R: the StatMatch package. *Rom Stat Rev* 2015; **2015**:57–68.
- Daly C, Gibson WP, Taylor GH. et al. A knowledge-based approach to the statistical mapping of climate. *Climate Res* 2002; **22**:99–113. <https://doi.org/10.3354/cr022099>.
- Dorfman AH. Towards a routine external evaluation protocol for small area estimation. *Int Stat Rev* 2018; **86**:259–74. <https://doi.org/10.1111/insr.12248>.
- ECOMAP. National Hierarchical Framework of Ecological Units. Unpublished Administrative Paper. Washington, DC: U.S. Department of Agriculture, Forest Service, 1993.
- Fay RE, Herriot RA. Estimates of income for small places: an application of James–Stein procedures to census data. *J Am Stat Assoc* 1979; **74**:269–77. <https://doi.org/10.1080/01621459.1979.10482505>.
- Finley AO, Banerjee S, MacFarlane DW. A hierarchical model for quantifying forest variables over large heterogeneous landscapes with uncertain forest areas. *J Am Stat Assoc* 2011; **106**:31–48. <https://doi.org/10.1198/jasa.2011.ap09653>.
- Frescino TS, Moisen GG, Patterson PL. et al. FIESTA: a forest inventory estimation and analysis R package. *Ecography* 2023a; **2023**:e06428. <https://doi.org/10.1111/ecog.06428>.
- Frescino TS, Toney C, White GW. FIESTAutils: utility functions for Forest inventory estimation and analysis. R package version 2023b; **1.1.5**. <https://CRAN.R-project.org/package=FIESTAutils>.
- Frescino TS, White GW. FIESTAanalysis: analysis functions for the FIESTA package. R package version 2022; **0.2.1**.
- Heath LS, Hansen M, Smith JE. et al. Investigation into calculating tree biomass and carbon in the FIADB using a biomass expansion factor approach. In: *Proceedings of the Forest Inventory and Analysis (FIA) Symposium*, U.S. Department of Agriculture Forest Service, Rocky Mountain Research Station, Park City, UT, 2008; **2009**:21–3.
- Horvitz DG, Thompson DJ. A generalization of sampling without replacement from a finite universe. *J Am Stat Assoc* 1952; **47**:663–85. <https://doi.org/10.1080/01621459.1952.10483446>.
- Isaki CT, Fuller WA. Survey design under the regression superpopulation model. *J Am Stat Assoc* 1982; **77**:89–96. <https://doi.org/10.1080/01621459.1982.10477770>.
- James G, Witten D, Hastie T. et al. *An Introduction to Statistical Learning with Applications in R*. 2nd edn. New York: Springer, 2021. <https://doi.org/10.1007/978-1-0716-1418-1>.
- Lehtonen R, Veijanen A. Design-based methods of estimation for domains and small areas. In: *Handbook of Statistics*, Vol. **29B**. Amsterdam: Elsevier, 2009. Chap. 31, 219–49.
- McRoberts RE, Bechtold WA, Patterson PL. et al. The enhanced forest inventory and analysis program of the USDA Forest Service: historical perspective and announcement of statistical documentation. *J For* 2005; **103**:304–8. <https://doi.org/10.1093/jof/103.6.304>.
- Morris TP, White IR, Royston P. Tuning multiple imputation by predictive mean matching and local residual draws. *BMC Med Res Methodol* 2014; **14**:1–13. <https://doi.org/10.1186/1471-2288-14-75>.
- Nur UAM, Longford NT, Cade JE. et al. Dealing with incomplete data in questionnaires of food and alcohol consumption. *Stat Transition* 2005; **7**:111–34.
- Picotte JJ, Dockter D, Long J. et al. LANDFIRE remap prototype mapping effort: developing a new framework for mapping vegetation classification, change, and structure. *Fire* 2019; **2**:35. <https://doi.org/10.3390/fire2020035>.
- Rao JNK, Molina I. *Small Area Estimation*. 2nd edn. New Jersey: John Wiley & Sons, 2015. <https://doi.org/10.1002/9781118735855>.
- Rollins MG. LANDFIRE: a nationally consistent vegetation, wildland fire, and fuel assessment. *Int J Wildland Fire* 2009; **18**:235–49. <https://doi.org/10.1071/WF08088>.
- Rubin DB. *Multiple Imputation for Nonresponse in Surveys*. New Jersey: John Wiley & Sons, 2004.
- Rubin DB, Schenker N. Multiple imputation for interval estimation from simple random samples with ignorable nonresponse.

- J Am Stat Assoc* 1986; **81**:366–74. <https://doi.org/10.1080/01621459.1986.10478280>.
- Schwager P, Berg C. Remote sensing variables improve species distribution models for alpine plant species. *Basic Appl Ecol* 2021; **54**: 1–13. <https://doi.org/10.1016/j.baee.2021.04.002>.
- Templ M, Meindl B, Kowarik A. et al. Simulation of synthetic complex data: the R package simPop. *J Stat Softw* 2017; **79**:1–38. <https://doi.org/10.18637/jss.v079.i10>.
- Tzavidis N, Zhang LC, Luna A. et al. From start to finish: a framework for the production of small area official statistics. *J R Stat Soc A* 2018; **181**:927–79. <https://doi.org/10.1111/rssa.12364>.
- Geological Survey US. LANDFIRE Elevation. Sioux Falls, South Dakota: USGS EROS, 2019.
- White GW, Yamamoto JK, Elsyad DH. et al. Small area estimation of forest biomass via a two-stage model for continuous zero-inflated data. *Can J For Res* 2025; **55**:1–19. <https://doi.org/10.1139/cjfr-2024-0149>.
- Wieczorek J, Franco C. An empirical artificial population and sampling design for small-area model evaluation. In: *Proceedings of the 2013 Joint Statistical Meetings*, American Statistical Association, Alexandria, VA, 2013.
- Wieczorek J, Nugent C, Hawala S. A Bayesian zero-one inflated beta model for small area shrinkage estimation. In: *Proceedings of the 2012 Joint Statistical Meetings*, American Statistical Association, Alexandria, VA, 2012.
- Wojcik OC, Olson SD, Nguyen PHV. et al. GREGORY: a modified generalized regression estimator approach to estimating Forest attributes in the interior western US. *Front For Global Change* 2022; **4**. <https://doi.org/10.3389/ffgc.2021.763414>.
- Woodruff RS. Use of a regression technique to produce area breakdowns of the monthly national estimates of retail trade. *J Am Stat Assoc* 1966; **61**:496–504. <https://doi.org/10.1080/01621459.1966.10480885>.
- Yang L, Jin S, Danielson P. et al. A new generation of the United States National Land Cover Database: requirements, research priorities, design, and implementation strategies. *ISPRS J Photogramm Remote Sens* 2018; **146**:108–23. <https://doi.org/10.1016/j.isprsjprs.2018.09.006>.
- Yang S, Kim JK. Nearest neighbor imputation for general parameter estimation in survey sampling. In: *The Econometrics of Complex Survey Data*, Vol. **39**. West Yorkshire: Emerald Publishing Limited, 2019, 209–34. <https://doi.org/10.1108/S0731-905320190000039012>.
- Zanaga D, Van De Kerchove R, De Keersmaecker W. et al. ESA WorldCover 10 M 2020 v100. Paris, France: European Space Agency, 2021.