Optimizing Sensor Placement with Greedy Algorithms: A Case Study in Wildlife Camera Trapping for Spatial Capture-Recapture Population Estimation

Hannah Murray[1] , **Amrita Gupta**[2] , **Arielle W. Parsons**[3] , **Justin P. Suraci**[3] , **Bistra Dilkina**[1]
[1]University of Southern California
[2]Microsoft AI for Good Lab
[3]Conservation Science Partners
hannahmu@usc.edu, amritagupta@microsoft.com, arielle@csp-inc.org, justin@csp-inc.org, dilkina@usc.edu

Abstract

Estimating wildlife populations is central to conservation planning, yet designing sensor deployments that produce reliable data for such estimates remains challenging. Spatial capture-recapture (SCR) models, widely used to estimate animal population sizes, are highly sensitive to sensor layout, where poor placement can substantially increase uncertainty in population estimates. We present a novel approach that formulates sensor placement for SCR as a scenario-based optimization problem under real-world resource constraints. Using collections of simulated animal capture histories spanning ecologically plausible parameter ranges, candidate sensor placements are evaluated via closed-form, SCR-derived design criteria linked to the precision of population estimates and optimized using both genetic algorithms and a greedy search strategy. We demonstrate our approach through an example American marten camera trapping study in British Columbia's South Chilcotin Mountains, achieving lower relative standard error and bias in population estimates than single-scenario and grid-based baselines. Our method was co-developed with conservation practitioners and is currently being used in real-world monitoring programs. This framework offers a general approach for designing wildlife surveys that support reliable population estimation across a range of realistic ecological scenarios.

1 Introduction

Wildlife monitoring is the basis for estimating the population size and spatial distribution of species, which inform assessments of ecosystem health and guide conservation management strategies [Oliver *et al.*, 2023]. These estimates are derived from observations collected by spatially distributed sensors such as camera traps, which capture evidence of wildlife and their movements across a landscape [Gibb *et al.*, 2019; Chen *et al.*, 2022; Steenweg *et al.*, 2017]. While such non-invasive technologies can increase detection rates and reduce field effort and costs, the accuracy and precision of resulting population estimates depend on how sensors are arranged spatially. In practice, sensor deployments are constrained by limited budgets and site accessibility, creating a need for principled approaches to sensor placement that maximize inferential power under these constraints [Kerry *et al.*, 2022].

Spatial capture-recapture (SCR) models are commonly used to estimate wildlife population sizes from sensor observations. In these models, detection probability explicitly depends on the distance between sensors and individual animals' activity centers, which are not directly observed. As a result, the number and spatial distribution of detections in a study depend strongly on how traps are placed relative to animals on the landscape. In practice, ecologists choose sensor spacing based on the assumed movement ranges of the target species. This ensures that animals moving through the landscape are likely to encounter nearby sensors while avoiding large unsampled gaps between traps.

The movement distances used for this spacing choice are typically informed by pilot studies or published estimates of species-specific movement parameters. However, sensor placements in practice are rarely designed with perfect knowledge of species-specific parameters, which can vary substantially across populations or individuals. While simple spacing heuristics rely primarily on a single movement scale parameter, more advanced approaches for optimizing sensor placement in SCR studies, including those based on genetic algorithms, require specifying additional habitat and detection-related parameters in order to generate candidate sensor layouts [Durbach *et al.*, 2021; Dupont *et al.*, 2021]. In both cases, layouts are optimized for a single "best guess" parameter set. This single-scenario approach can lead to sensor layouts that fail to yield informative observations needed for precise population estimation if the true species ecology deviates from those assumptions. Addressing this sensitivity requires sensor placement strategies that explicitly account for uncertainty in species-specific parameters.

We propose a novel simulation-based optimization approach for designing sensor placements, simultaneously evaluating candidate sensor layouts across multiple plausible ecological scenarios. It combines ecological simulation with either greedy search or genetic algorithms, both of which take into account a range of plausible species-specific parameters. Our approach generates simulated capture histories and can incorporate real-world constraints such as budget or landscape limitations. Notably, the environmental data it requires

(habitat covariates, species parameters, and candidate trap locations) are the same inputs already used for standard SCR analyses, even in the absence of optimization. Therefore, our framework leverages existing practitioner knowledge to derive principled camera placements rather than relying on ad hoc or purely grid-based designs. We evaluate candidate configurations based on their ability to estimate the true population size and the relative precision of those estimates.

We apply our framework in a real study area, here based on a simulated population of American martens in the South Chilcotin Mountains, Canada, to illustrate the utility of this approach. Beyond this case study, we are actively working with three wildlife conservation partners to pilot our approach to monitoring grizzly bears in the South Chilcotin Mountains, Canada, pumas in the Olympic Peninsula of Washington, USA, and Grévy's zebras in the Laikipia region of Kenya. By enabling more effective use of deployed cameras and more reliable population estimates, our proposed framework supports sustained wildlife monitoring and conservation decision-making at scale, consistent with UN Sustainable Development Goal 15 (Life on Land).

2 Background

2.1 Camera Trapping in Biodiversity Surveys

Camera traps are one of the most commonly used sensors due to their noninvasive nature and ability to provide temporal and spatial data on the whereabouts of animals [Glover-Kapfer *et al.*, 2019]. Camera trap surveys generally assume that animals occupy the landscape according to an underlying spatial distribution, with detections occurring when animal movement through the landscape intersects with sensors and triggers motion-activated captures. Consequently, survey design plays a central role in determining detection probability.

Camera traps are commonly deployed using grid-based configurations [Sun *et al.*, 2014]. One common design is the uniform layout, with cameras evenly spaced at regular intervals (e.g., every 3 kilometers) (Figure 1.A). Clustered designs are also common, and group cameras within smaller, concentrated areas. This reduces overall spatial coverage, but can increase the probability of detection of elusive species in key habitats (Figure 1.B) [Sun *et al.*, 2014]. Sequential designs are another approach, where cameras are deployed in grid-based patterns, but moved after a period of time to new locations (Figure 1.C). Camera placement is also often opportunistic, with traps placed primarily along accessible features such as hiking trails or roads [Tanwar *et al.*, 2021].

Survey design strongly influences the accuracy and precision of capture–recapture estimates [Efford and Boulanger, 2019; Sun *et al.*, 2014; Wilton *et al.*, 2014]. If detectors are spaced too widely relative to animal movement or placed disproportionately in certain habitats, the resulting detections may not reflect the true spatial distribution of individuals, leading to sparse data, reduced precision, and biased density estimates. These issues are especially pronounced when detection probability or animal density varies spatially with habitat preferences.

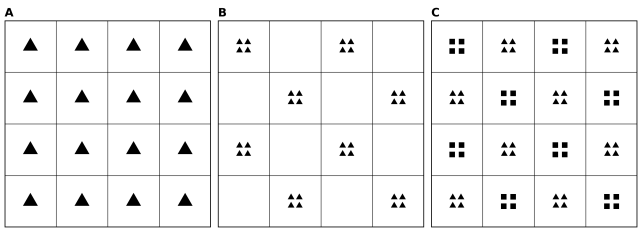

Figure 1: Traditional sensor deployment configurations. A represents a uniform approach. B represents a clustered design. C represents a sequential design, with triangles representing cameras at one time period, and squares at another.

2.2 Spatial Capture-Recapture Models

Spatial capture-recapture (SCR) models are the standard for estimating wildlife population size $\hat{N}$ using detection data from identifiable animals captured from sensors, known as capture histories. Capture histories are records of individual detection events, specifying which animal was captured at which camera and during which sampling time period.

These models are based on the concept that each animal has a unique activity center, which can be loosely interpreted as the center of its home range. The location of the activity center l relative to the sensor array influences the probability of an animal being detected [Royle and Young, 2008; Borchers and Efford, 2008]. Detection probability declines with increasing distance between an animal's activity center and a given sensor location j ($d(l, j)$), and this spatial decay is described by a detection function, which has an intercept g_0 and scale parameter σ (Eq. 1) [Royle and Young, 2008; Borchers and Efford, 2008]. The intercept g_0 represents the sensor's expected detection probability if an animal's activity center was exactly where the sensor is placed, and σ captures the distance between an activity center and a sensor at which detection probability declines to near zero [Royle and Young, 2008; Borchers and Efford, 2008]. We restrict attention to this standard distance-based SCR formulation; targeted placement at watering holes or other attractors would require different models that can account for resource-mediated variation in encounter rates or specialized covariates.

$$Pr[l, j] = g_0 \cdot e^{-\frac{1}{2\sigma^2} d(l,j)} \tag{1}$$

SCR models estimate $\hat{N}$ and other parameters by linking individual detections to the spatial distribution of activity centers. The model defines a likelihood function $L(\alpha, H_\alpha)$, expressing the probability of the observed capture histories H given parameters α and trap configuration $\mathbf{x}$. Here, α includes g_0 and σ, as well as the average density of activity centers D and parameters governing relative density, such as habitat covariates (β) like tree cover or human footprint. Parameters are estimated by maximizing this likelihood.

Notably, SCR models may fail to converge to reliable estimates when sensor spacing exceeds 2.5σ, resulting in insufficient recaptures to identify activity centers [Efford and Boulanger, 2019; Durbach *et al.*, 2021]. Wide spacing fails to overlap home ranges adequately, inflating $\hat{N}$ and compromising reliability, underscoring the need for optimized detec-

tor placement tailored to species-specific movement ecology [Efford and Boulanger, 2019; Durbach *et al.*, 2021]. Another key element of SCR model specification is the definition of the state space, or the area in which potential activity centers are distributed. More information on this can be reviewed in Section § A.1.

SCR model performance is typically evaluated using precision metrics, such as relative standard error of the estimated population size ($RSE(\hat{N})$), but accuracy metrics, such as bias ($\hat{N} - N$), can further quantify the difference between the estimated and true population sizes of the target species.

2.3 Prior Algorithmic Methods

Genetic algorithms are the primary algorithmic approach available for optimizing camera trap configurations in SCR studies for wildlife monitoring [Dupont *et al.*, 2021; Durbach *et al.*, 2021]. These methods are inspired by evolution and work by maintaining a set of possible sensor configurations that, through selection, crossover, and mutation, maximize a fitness function over several iterations [Wolters, 2015]. This method has been implemented in R packages such as `kofnGA` and its extension `SCRdesign::GAoptim`, which is specific to sensor placement for wildlife population estimation [Efford, 2025; Wolters, 2018].

While genetic algorithms represent a more sophisticated approach to sensor placement than grid-based methods, they have practical limitations. Evaluating camera configurations across multiple budgets is computationally expensive: with a budget of B cameras, a layout optimized for B cannot be reused to assess any smaller budget $b < B$ without rerunning the algorithm for each b. This makes it difficult for practitioners to quantify the marginal benefit of adding cameras, an important consideration when each additional unit is costly, and genetic algorithm methods further amplify this cost on the computational side by requiring separate optimizations for every budget level of interest. Secondly, these methods optimize configurations under a single-parameter scenario α. These parameters vary with species ecology, season, habitat, and individual behavior, creating uncertainty that a single-point evaluation cannot capture. Therefore, SCR models may fail to reach a finite and reliable maximum likelihood estimate if the true parameters do not closely match the parameters used in training. Given these limits, we introduce a modified genetic algorithm method that averages the fitness of each candidate design over a set of parameter draws, so selection is driven by its mean precision across multiple plausible scenarios rather than a single estimate (Section 3.2).

3 Methods

3.1 Stochastic Optimization Problem Formulation

We formulate trap placement as a stochastic optimization problem, minimizing expected loss under uncertain ecological parameters that follow distributions informed by prior literature.

Let α denote a vector of parameters, including detection probability at zero distance g_0, spatial scale σ, habitat covariates β, and overall density D. Given a finite set of candidate camera trap locations $\mathcal{J} = \{1, \ldots, J\}$ and a camera budget B, we seek a configuration $\mathbf{x}^\star \in \{0,1\}^J$ with $\sum_{j=1}^{J} x_j = B$ that minimizes a stochastic objective defined over plausible values of α. We define the per-parameter loss $f(\mathbf{x}; \alpha)$ as the relative standard error of the SCR population size estimate, $RSE(\hat{N})$, consistent with SCR study-design practice [Efford and Boulanger, 2019]. In other monitoring or decision-making settings, the same stochastic optimization framing could be paired with other loss functions tailored to different objectives.

The SCR model computes $f(\mathbf{x}; \alpha)$ for each parameter vector α by taking as input the chosen trap configuration $\mathbf{x}$ and the capture histories generated under that configuration $H(\mathbf{x}, \alpha)$, and returns an estimate of population size $\hat{N}(\mathbf{x}, \alpha)$ whose relative standard error defines the loss $f(\mathbf{x}; \alpha)$. We then define the stochastic objective as

$$Z(\mathbf{x}) = \mathbb{E}_\alpha \left[f(\mathbf{x}; \alpha) \right] \tag{2}$$

where α is a random element with a distribution that does not depend on $\mathbf{x}$, and $f(\mathbf{x}; \alpha)$ is a (deterministic) real-valued function. Therefore, the stochastic program is

$$\mathbf{x}^\star \in \arg \min_{\mathbf{x} \in \{0,1\}^J} Z(\mathbf{x}) \quad \text{subject to} \quad \sum_{j=1}^{J} x_j = B. \tag{3}$$

Sample average approximation (SAA) is a technique for solving stochastic optimization problems by sampling from the underlying distribution to generate a finite set of scenarios, thereby reducing the stochastic optimization problem to a deterministic analogue [Shapiro, 2003; Verweij *et al.*, 2003; Kleywegt *et al.*, 2002]. We consider a set of G i.i.d. parameter samples $\{\alpha^{(g)}\}_{g=1}^{G}$ and approximate $Z(\mathbf{x})$ by the average relative standard error of the estimated population size $\frac{1}{G} \sum_{g=1}^{G} RSE\big(\hat{N} \mid \mathbf{x}, \alpha^{(g)}\big)$ across each plausible α given selected sensor locations $\mathbf{x}$. In SAA, we repeatedly solve M finite-sample approximations of the stochastic problem. In run $m = 1, \ldots, M$ we first draw a training set of T parameter scenarios $\alpha^{(m,1)}, \ldots, \alpha^{(m,T)}$ from the simulation pool G (with replacement) and solve

$$\mathbf{y}^{(m)} \in \arg \min_{\mathbf{x} \in \{0,1\}^J} \hat{Z}^{(m)}(\mathbf{x}) \quad \text{subject to} \quad \sum_{j=1}^{J} x_j = B, \tag{4}$$

where the sample-average objective is

$$\hat{Z}^{(m)}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^{T} RSE\big(\hat{N} \mid \mathbf{x}, \alpha^{(m,t)}\big). \tag{5}$$

We then evaluate each $\mathbf{y}^{(m)}$ directly using SCR on a separate validation set of unseen parameter scenarios in G by recording the average $RSE(\hat{N})$ across those draws. We select the $\mathbf{y}^{(m)}$ that achieves the lowest average $RSE(\hat{N})$, and test it on another unseen set of parameters in G using SCR, constituting our test results.

The SAA method provides theoretical guarantees on the optimality gap between the true stochastic optimum and solutions obtained from finite samples, with the gap shrinking

as the number of training and test scenarios increases [Verweij *et al.*, 2003]. While in real-world wildlife population estimation, each SAA subproblem is solved approximately rather than to true optimality, using the SAA procedure of optimizing over multiple independent samples and evaluating on large held-out test sets still offers a valuable approach.

3.2 Optimizing Sensor Placement

This subsection describes how we search for camera trap configurations that achieve low average $RSE(\hat{N})$. Our methodologically diverse approaches operate using the SAA objective, evaluating each candidate configuration on a finite set of training parameter scenarios and using average precision error as the quantity to minimize.

Greedy Algorithm Given the computational complexity associated with evaluating the large number of possible sensor configurations, we propose using forward greedy algorithms to efficiently search the solution space. The greedy algorithm (Algorithm 1) selects the sensor location, one at a time, that produces the greatest expected reduction in average $RSE(\hat{N})$ across a set of α parameter draws when added to the current partial configuration, starting from the empty set.

The algorithm is initialized with an empty configuration $x = [0, 0, \ldots, 0]$, indicating no traps have been selected (line 1). The outer loop (line 2) iterates B times, adding one trap per iteration until the budget is exhausted. At the start of each iteration, we initialize a tracker for the minimum observed average $RSE(\hat{N})$ across T scenarios and the best candidate location (lines 3–4). The inner loop (line 5) then considers every candidate location j not yet included in the current configuration. For each j, we temporarily add it to obtain a candidate configuration x_{temp} (line 6) and evaluate the sample-average objective $\hat{Z}^{(m)}(\mathbf{x}_{\text{temp}})$ over the current SAA replicate's T parameter scenarios (line 7). If this candidate configuration achieves a lower average $RSE(\hat{N})$ than any previously evaluated in this iteration, it is recorded as the current best (lines 8–10). After all unselected locations have been evaluated, the location yielding the greatest reduction in $RSE(\hat{N})$ is permanently added to the configuration (line 13).

A key advantage of this forward selection strategy is that a *single* run produces layouts for all budgets from 1 to B: the first chosen location is optimal for a budget of one camera, the first two locations define the layout for a budget of two cameras, and so on, until B cameras have been placed. This eliminates the need to rerun the optimization procedure separately for each budget level, reducing computational cost.

Modified Genetic Algorithm We adapt the existing genetic algorithm method to our multi-scenario setting. The existing genetic baseline is designed to optimize performance under a single parameter setting and, therefore, cannot directly account for variation across multiple α scenarios. To address this, we define fitness in terms of the same sample-average objective used in our SAA framework. Configurations with lower average error across T scenarios yield higher fitness. SAA-Genetic runs the modified genetic algorithm M times using the same groups of T parameter scenarios in training and evaluates using the same validation and test sets as the greedy method to enable one-to-one comparison.

Algorithm 1 Forward Greedy Algorithm for Trap Selection

Input: Candidate trap locations, budget, scenarios $\alpha^{(m,t)}$
Output: Selected trap configuration $\mathbf{x}$

1: $\mathbf{x} \leftarrow [0, 0, \ldots, 0]$ (no traps selected)
2: **for** $s = 1$ to budget **do**
3: $\quad min_{RSE} \leftarrow \infty$
4: $\quad add \leftarrow$ None
5: $\quad$ **for** each trap j with $x_j = 0$ **do**
6: $\quad\quad$ Temporarily set $x_j \leftarrow 1$ to obtain $\mathbf{x}_{temp}$
7: $\quad\quad$ Evaluate objective $\hat{Z}^{(m)}(\mathbf{x}_{temp})$
8: $\quad\quad$ **if** $\hat{Z}^{(m)}(\mathbf{x}_{temp}) < min_{RSE}$ **then**
9: $\quad\quad\quad add \leftarrow j$
10: $\quad\quad\quad min_{RSE} \leftarrow \hat{Z}^{(m)}(\mathbf{x}_{temp})$
11: $\quad\quad$ **end if**
12: $\quad$ **end for**
13: $\quad$ Set $x_{add} \leftarrow 1$ (permanently add trap with largest decrease in RSE)
14: **end for**
15: **return** $\mathbf{x}$

Integer Programming We also attempted a formulation using integer programming. The objective commonly used by ecologists of $RSE(\hat{N})$ is highly nonlinear (Section 3.3). To leverage this method, we constructed a truncated integer programming formulation based on simplifying the nonlinear terms in the $RSE(\hat{N})$ expression. Although this approximation led to faster solve times than our greedy and modified genetic methods, the resulting configurations did not improve on the greedy solutions. They generally performed worse than our modified genetic method, but better than uniform baselines. Because the truncation did not yield better empirical performance, we did not pursue this approach further.

3.3 Fast Computation of Population Estimation Precision

Evaluating the performance of each candidate layout requires fitting the SCR model (solving a maximum likelihood estimation problem) over the simulated capture histories under that sensor layout, which is computationally prohibitive in real-world settings. For context, evaluating one 10-camera and one 80-camera layout on 150 simulations with SCR takes roughly 40 and 210 minutes of wall-clock time, respectively, on a remote computing cluster. To overcome this barrier, Efford and Boulanger introduced an approximation that can estimate the performance of camera trap configurations without repeatedly fitting SCR models.

$$\overline{RSE}(\hat{N}|\mathbf{x}, \alpha) = \frac{1}{\sqrt{\min(E_n(\mathbf{x}, \alpha), E_r(\mathbf{x}, \alpha))}} \qquad (6)$$

They proposed a non-linear, closed-form approximation for the relative standard error of the estimated population size (Eq. 6), where E_n represents the expected number of unique individuals observed and E_r represents the expected number of recaptures (Eq. 7, 8) and $D(l)$ is the density in location l under the habitat parameters β in the parameter vector α [Efford and Boulanger, 2019]. We expand the expressions for

E_n and E_r to take into account the sensor placement decisions encoded by a binary vector $\mathbf{x}$, where x_j represents the decision of whether a camera is placed at each location j (1) or not (0) under a given parameter scenario α.

$$\begin{aligned} E_n(\mathbf{x}, \alpha) &= \text{Expected \# of Unique Individuals } i \text{ Detected} \\ &= \sum_l D(l) \cdot \Pr[\text{detect } i \text{ with activity center } l] \\ &= \sum_l D(l) \cdot \{1 - \Pr[\text{never detect } i \text{ with } l]\} \\ &= \sum_l D(l) \cdot \left\{ 1 - \prod_{j=1}^{J} (1 - x_j \Pr[l, j;\ \alpha])^K \right\} \end{aligned} \tag{7}$$

$$\begin{aligned} E_r(\mathbf{x}, \alpha) &= \text{Expected Number of Recaptures} \\ &= \text{Expected Number of Captures} - E_n \\ &= \sum_l \sum_j \sum_k D(l) \cdot x_j \cdot \Pr[l, j, \alpha] - E_n \end{aligned} \tag{8}$$

Given the runtime improvements provided by this closed form approximation over using maximum likelihood estimation of SCR, we use $\overline{RSE}(\hat{N})$ when optimizing sample averages across the training parameter scenarios.

4 Case Study Data Generation

We define five habitat preference and detection parameters, which together form the α parameter vector, and whose values are treated as random inputs to generate the scenario set in our American marten case study – g_0, σ, β_1 (tree cover), β_2 (human footprint), and average density D. To define a realistic range of values, we drew upon ecological and behavioral information from published literature. Detection parameters were varied over a plausible range of baseline detection probabilities ($g_0 \in [0.01, 0.5]$) and movement scales ($\sigma \in [200, 1000]$ m), with the σ range chosen to match home range sizes observed in telemetry studies across North America [Feldhamer *et al.*, 2003; Connor *et al.*, 2023]. We simulated spatial variation in marten density in response to percent tree cover [Sexton *et al.*, 2013] and human footprint index [Wildlife Conservation Society, 2005]. We assumed a strongly positive relationship with tree cover ($\beta_1 \in [0.1, 2]$) and a negative to weakly positive relationship with human footprint ($\beta_2 \in [-1, 0.5]$) [Feldhamer *et al.*, 2003]. Average population density was varied over $D \in [0.0015, 0.004]$ animals/ha, consistent with estimates of marten abundance across North American forested landscapes, yielding true $N \in [178, 561]$ for our study area [Feldhamer *et al.*, 2003]. We assumed a study with $K = 5$ sampling time periods of indeterminate length, consistent with the `SCR` package default. Our greedy and genetic methods optimize over this entire range, so performance depends on the assumption that the true parameters lie somewhere within these bounds rather than on precise point estimates. Before using simulated data for design optimization, we verified that the simulation produces capture rates and detection patterns that are consistent with expectations from prior work or experts.

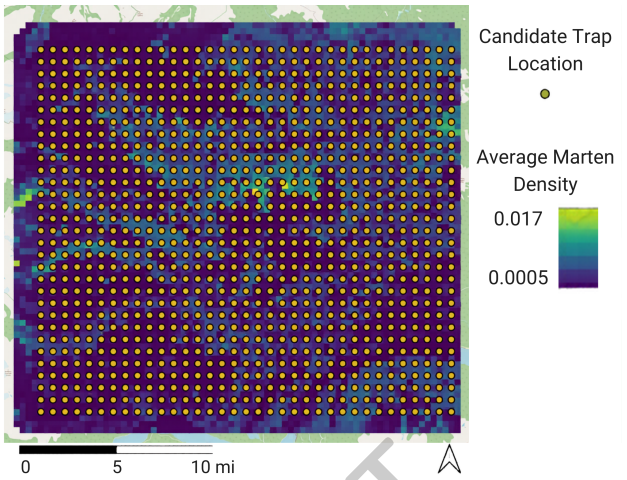

Figure 2: Simulated Average Marten Density (lower = purple, higher = yellow) & Candidate Locations for Camera Trap Placement (yellow circles) within Study Area in South Chilcotin Mountains

We generated 300 parameter realizations covering the set of possible combinations of parameter values G using Latin Hypercube Sampling (LHS) [Helton and Davis, 2003]. LHS divides each parameter's range into equal-probability intervals and ensures that each interval is sampled exactly once per parameter across all samples, representing an efficient way to achieve stratified, space-filling coverage with a minimal number of samples. We reserved 100 parameter draws as our training bank, of which $M = 20$ repetitions optimize on $T = 10$ randomly selected parameters with replacement in our SAA framework. We reserved a further 50 draws for validation of our optimized layouts using SCR, and the remaining 150 as an independent test set for SCR. More details on our selection of the size of our training, validation, and test parameter sets can be found in Section § A.2.

Candidate camera trap sites were defined as a uniform grid at 1 km spacing across the study area. This spacing is $\approx 67\%$ finer than the 3 km grids typically deployed in this landscape, providing ecologists with a denser set of candidate placements to choose from. When necessary, experts could refine this candidate set further to remove locations that are inaccessible by humans due to terrain or other logistical constraints.

We evaluate performance using our SAA framework across eight trap budget levels (10, 20, . . . , 80). These budgets were selected as they include the number of traps needed to satisfy the ideal 2.5σ spacing in the study area for our species across the range of potential sigma values sampled, from $\sigma = 200$ ($\approx$ 75 traps) to $\sigma = 1000$ ($\approx$ 15 traps). The simulated average density across all 300 LHS samples and the candidate trap locations in our dense grid can be viewed in Figure 2.

5 Results & Discussion

Figure 3 showcases the true relative standard error of the predicted population size, $RSE(\hat{N})$, and mean absolute bias, $|N - \hat{N}|$, when evaluated on our test parameter set using SCR, and compares our SAA-Greedy and SAA-Genetic methods to the Uniform baseline. These results exclude parameter real-

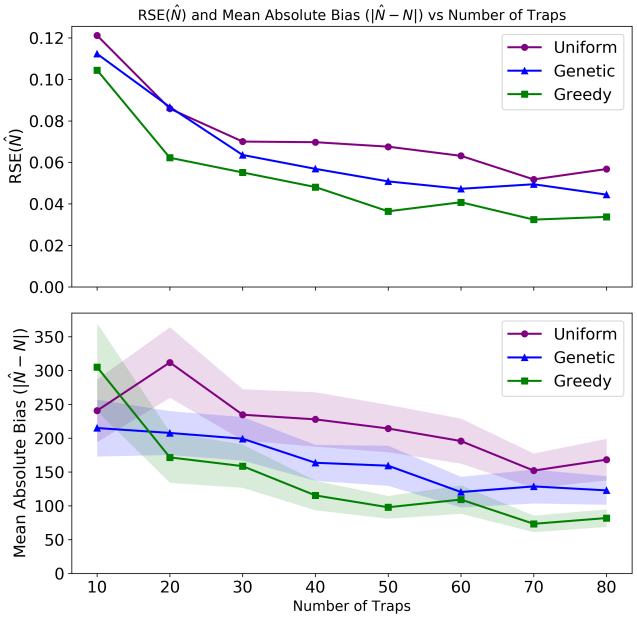

Figure 3: Relative Standard Error and Mean Absolute Bias per Camera Budget by Method (Uniform - purple, SAA-Genetic - blue, SAA-Greedy - green) Configuration Design Methods. Shading represents 95% confidence interval of values across test parameter set.

Number of Traps	Method	RSE($\hat{N}$)†	Mean Abs. Bias†
10	Uniform	0.121	240.59
	Genetic	0.112	**214.89**
	Greedy	**0.104**	304.95
20	Uniform	0.086	311.84
	Genetic	0.087	207.65
	Greedy	^{UG}**0.062**	^{U}**171.51**
30	Uniform	0.070	234.58
	Genetic	0.064	199.07
	Greedy	^{U}**0.055**	^{U}**158.59**
40	Uniform	0.070	227.69
	Genetic	0.057	163.55
	Greedy	^{UG}**0.048**	^{UG}**115.33**
50	Uniform	0.068	214.21
	Genetic	0.051	159.22
	Greedy	^{UG}**0.036**	^{UG}**97.82**
60	Uniform	0.063	195.68
	Genetic	0.047	120.32
	Greedy	^{U}**0.041**	^{U}**109.20**
70	Uniform	0.052	151.98
	Genetic	0.049	128.62
	Greedy	^{UG}**0.032**	^{UG}**73.35**
80	Uniform	0.057	168.21
	Genetic	0.044	122.76
	Greedy	^{UG}**0.034**	^{UG}**81.88**

Table 1: Relative standard error RSE($\hat{N}$) and mean absolute bias $|\hat{N}-N|$ by trap budget and method. †Superscripts on greedy method rows imply statistically significantly lower RSE or bias at $\alpha = 0.05$ relative to other methods: ^{U} vs Uniform and ^{G} vs Genetic.

izations that were total failures of the SCR model, where SCR fails to reach a finite and reliable maximum likelihood estimate. When comparing methods, total failures occurred most commonly in the Uniform baseline. When comparing budgets, failures exhibited a linear decreasing trend as the budget increased before plateauing around 40 traps for all methods (Section § A.3). Total failures were also less common in our SAA-Genetic approach than in the single α scenario genetic baseline, indicating further improvements in our genetic method (Section § A.4).

Our SAA-Greedy approach achieved the lowest $RSE(\hat{N})$ across all camera budgets, demonstrating better precision in estimating the population size over the independent test set of scenarios than our SAA-Genetic method and the Uniform baseline. When comparing our SAA-Greedy method to the Uniform baseline, the SAA-Greedy method produced a statistically significant lower bias on all budgets but 10 cameras, which produced a higher bias that was not statistically significant (Table 1). The error at 10 cameras should be interpreted as an edge case rather than a realistic design choice, as a 10-trap configuration cannot maintain even the coarsest recommended spacing relative to the largest simulated σ values, and ecologists would not deploy so few cameras in this landscape for the American marten. This demonstrates our SAA-Greedy method's robust performance advantage over the Uniform baseline across a full range of plausible sensor budgets. The SAA-Genetic approach also improved performance relative to Uniform, but its gains leveled off as trap budget increased, whereas the SAA-Greedy method continued to improve. Again, the SAA-Greedy method produced lower bias across all budgets except 10 cameras, which showed a higher bias that was not statistically significant (Table 1). Of the other budgets, the SAA-Greedy method achieved a statistically significant improvement in bias compared to the SAA-Genetic method at 40, 50, 70, and 80 cameras. Additionally, our SAA-Greedy method is much more computationally feasible, with SAA-Greedy running 2.25 times faster than the SAA-Genetic method to optimize across all eight budgets. This highlights our greedy method as a practical, high-performing approach that captures most of the performance benefits at a fraction of the computational cost, without causing any statistically significant impact on the accuracy of predictions for the other parameters (Section § A.5).

Overall, the 70 camera layout achieved the lowest RSE and mean absolute bias across all the budgets (Figure 3, Table 1). This is not surprising, as 70 cameras is sufficient to maintain the ideal 2.5σ spacing along the study area for the majority of our parameter draws. Looking at the 70-camera configurations resulting from the Uniform, SAA-Genetic, and SAA-Greedy methods reveals clear differences (Figure 4). We hypothesize the clustering present in both the SAA-Genetic and SAA-Greedy approaches arises from a couple of sources. Firstly, the clusters of cameras border some of the most contiguous forests within the study area. This could symbolize our model is adjusted appropriately based on the marten's habitat preferences. Secondly, given

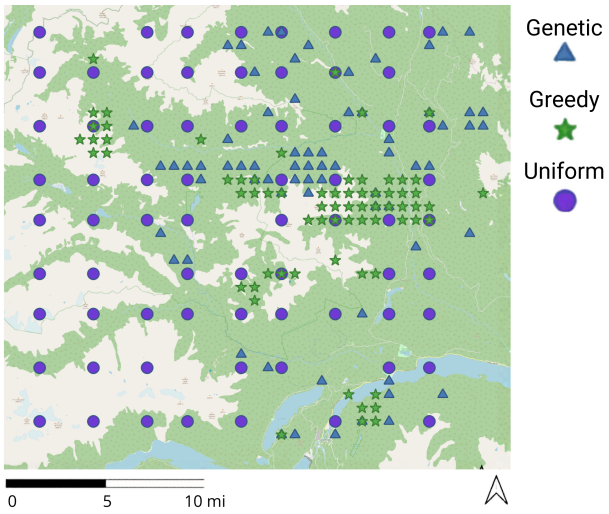

Figure 4: Baseline Uniform (Purple Circles), SAA-Genetic (Blue Triangles), and SAA-Greedy (Green Stars) Camera Trap Configurations with Budget of 70 Cameras.

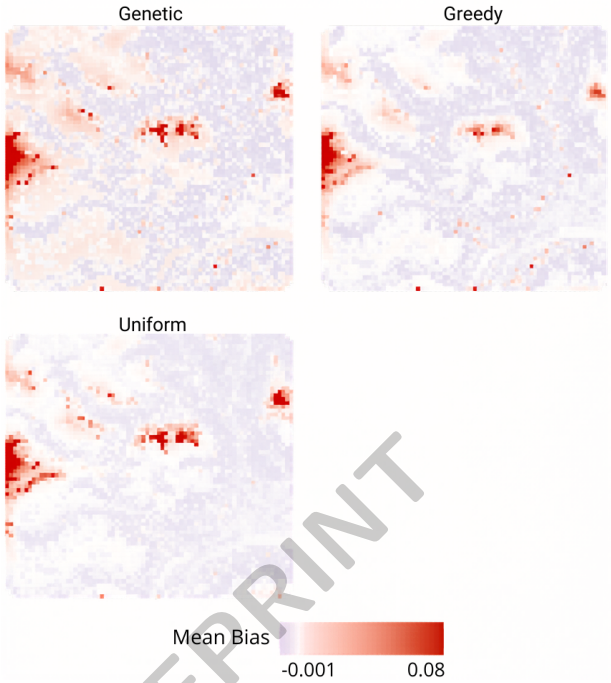

Figure 5: Mean bias in predicted cell-level density under our SAA-Genetic, (top left) and SAA-Greedy (top right), and uniform (bottom) 70-camera configurations. Warmer colors indicate positive bias (overestimation), cooler colors indicate negative bias (underestimation), and white cells are zero bias.

our candidate sensor grid locations' spacing of 1 km and the smallest marten σ being near 200 m, the algorithm likely prioritizes maximizing the expected number of recaptures (E_r) in the $\overline{RSE(\hat{N})}$ computation done in the optimization.

Figure 5 shows the impact each 70-camera layout had on the error in predicted cell densities across the landscape. In all designs, clusters of positive bias (red cells) remain, particularly along the western border of the study area, which likely reflects large number of activity centers simulated in the buffer area just outside the deployable region (Section §A.1). Per-cell density errors ranged from about -0.001 to 0.080 animals/ha. Relative to the Uniform layout, both optimized designs substantially shrink the magnitude of error in high-bias regions: the maximum positive bias is reduced by roughly $\approx 30\%$ in the SAA-Greedy design and $\approx 25\%$ in the SAA-Genetic. The SAA-Greedy layout does so most effectively, with fewer and less intense hotspots than the SAA-Genetic layout. This pattern indicates that the SAA-Greedy design's reduction in bias arise from a broad suppression of spatial bias across the landscape, rather than from trading improvements in one region for worse performance in others.

6 Conclusion

We have presented a novel framework for optimizing camera trap placement across ecologically realistic simulations of animal behavior using greedy search and genetic algorithms. Across a wide range of budgets and parameter realizations, our SAA-Greedy method consistently achieved lower relative standard error and mean absolute bias than both the Uniform baseline and SAA-Genetic method. These gains appear not only in global summary metrics, but also in the spatial distribution of error, where algorithmic-derived layouts reduce high-bias hotspots and lessen estimation error across the landscape. Importantly, these gains persist when evaluated under parameter values different from those used during optimization. Together, these results indicate that our methods yield more reliable density estimates, enabling ecologists to better assess population status, evaluate conservation actions, and allocate limited monitoring resources more effectively.

We are now piloting these methods in real-world monitoring studies. We have collaborated with the University of British Columbia to deploy cameras for grizzly bear monitoring in the South Chilcotin Mountains, BC, and are currently in the process of analyzing the resulting imagery. In Washington State, we are working with Panthera to optimize puma monitoring in the Olympic Peninsula. In Kenya, we are collaborating with the Mpala Research Centre and MIT to apply our approach to Grévy's zebras in the Laikipia region. These pilots will enable us to refine the framework based on practitioner feedback.

Several extensions are possible for this work. Although our study focuses on allocating traps for new studies, the approach could also support existing deployments, including optimizing additional placements for partially deployed grids or identifying which traps can be removed with minimal performance loss. Our framework is also complementary to adaptive survey designs, in which results from one time period inform future camera placement. Lastly, it could be extended to optimize deployments for multiple species.

This work aims to provide a generalizable framework for enabling more efficient wildlife monitoring. To support reproducibility and reuse, our simulation, optimization, and evaluation code are available on GitHub.

A Appendix

A.1 Defining SCR Model State Spaces

A critical element of SCR modeling is defining the state space, or the area over which potential activity centers are distributed (A). The state space includes a buffer around the detector array wide enough to account for animals whose activity centers lie outside the trapping area but who may still be captured by sensors within it. Too narrow a buffer biases density estimates high by excluding these individuals from A; excessively large buffers add computational cost without improving estimates [Royle *et al.*, 2013]. We follow the common convention of applying a buffer of 2σ.

A.2 Sampling for Case Study Data Generation

Previous applications of Latin Hypercube Sampling (LHS) in ecological modeling suggest a few hundred samples are typically sufficient for robust coverage of joint parameter space [Marino *et al.*, 2008]. To verify this, we generated capture histories for each of our 300 LHS draws using a σ-spaced trap array, fit a SCR model to each, then subsampled 25, 50, 100, 150, and 250 simulations (repeated 100 times) and computed the relative root mean squared error (rRMSE) of each parameter. Error levels stabilized at 100–150 samples, so we used 150 parameter draws for training and validation and reserved the remaining 150 as an independent test set.

A.3 Total Number of SCR Model Failures

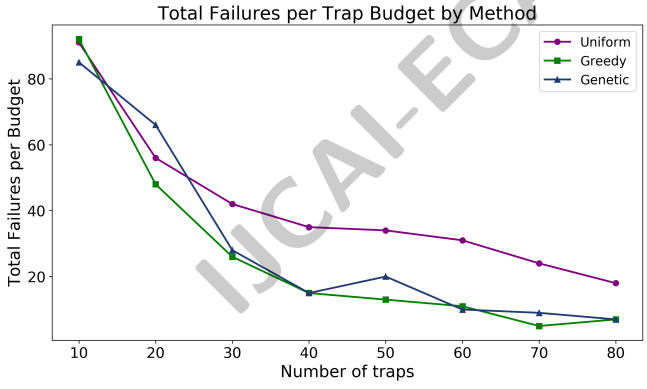

Figure 6: Total number of total failures of the SCR test evaluation per trap budget for three placement strategies. Lines show, for each camera budget, the total count of test-set simulations which failed to converge to reliable estimates for the uniform baseline (purple), SAA-Genetic (blue), and SAA-Greedy (green) methods. Failure counts decrease with increasing trap numbers for all methods, with the SAA-Greedy and SAA-Genetic methods generally producing fewer test failures than the uniform layout at moderate and high budgets.

A.4 SAA-Genetic Algorithm Improvements

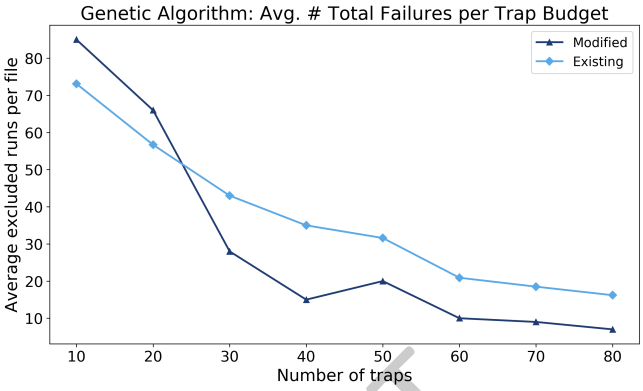

Figure 7: Average number of total failures per budget for the modified, SAA-Genetic (dark blue) and existing (light blue) genetic algorithms across trap budgets. For the existing method, each point is based on 20 independently trained genetic algorithms, where each optimization used a randomly selected single parameter scenario and was evaluated on a common test set of 150 parameter draws. For the SAA method, each point is based on 20 optimization replications, where each genetic replication was trained on 10 randomly selected parameter scenarios, and the plotted failures correspond to the selected layout for each budget evaluated on the same 150-draw test set. The large number of failures at the smallest trap budgets are likely a consequence of insufficient camera coverage over the study area.

A.5 Full Parameter Performance

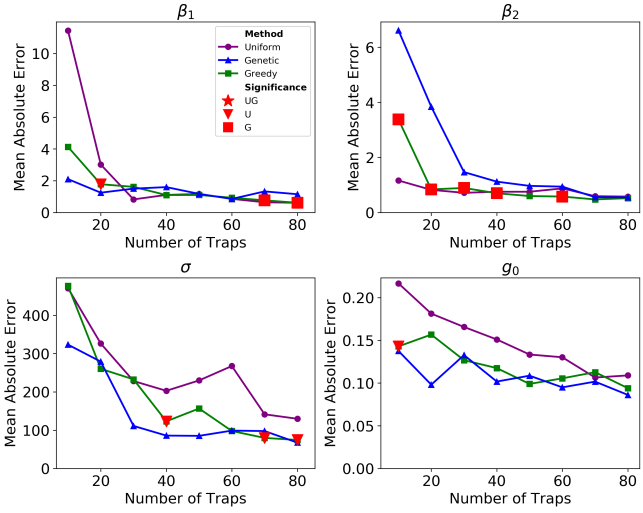

Figure 8: Mean absolute error in detection-function parameters. Panels show error in β_1, β_2, σ, and g_0 as a function of the number of traps, comparing the uniform baseline (purple), SAA-Genetic (blue), and SAA-Greedy method (green). Red symbols on the SAA-Greedy curve indicate budgets where the SAA-Greedy method produced a statistically significant better (Welch t-test, $\alpha = 0.05$) than the uniform (triangle, U), the SAA-Genetic (square, G), or both methods (star, UG).

Ethical Statement

The authors have no ethical issues to report.

Acknowledgments

This work was supported by Allen Family Philanthropies through the "Optimizing sensor placement for wildlife surveys to improve model inference and conservation outcomes" grant. We thank Angela Brennan from Save the Elephants and Robin Naidoo from the University of British Columbia for their guidance and assistance in deploying camera traps in the South Chilcotin Mountains for the grizzly bear monitoring component. We also thank Brett G. Dickson of Conservation Science Partners for his role in project conceptualization and general support. Lastly, we thank the reviewers for their valuable and constructive feedback.

References

[Borchers and Efford, 2008] D. L. Borchers and M. G. Efford. Spatially explicit maximum likelihood methods for capture-recapture studies. *Biometrics*, 64(2):377–385, June 2008.

[Chen *et al.*, 2022] Cheng Chen, Jedediah F. Brodie, Roland Kays, T. Jonathan Davies, Runzhe Liu, Jason T. Fisher, Jorge Ahumada, William McShea, Douglas Sheil, Bernard Agwanda, Mahandry H. Andrianarisoa, Robyn D. Appleton, Robert Bitariho, Santiago Espinosa, Melissa M. Grigione, Kristofer M. Helgen, Andy Hubbard, Cindy M. Hurtado, Patrick A. Jansen, Xuelong Jiang, Alex Jones, Elizabeth L. Kalies, Cisquet Kiebou-Opepa, Xueyou Li, Marcela Guimarães Moreira Lima, Erik Meyer, Anna B. Miller, Thomas Murphy, Renzo Piana, Rui-Chang Quan, Christopher T. Rota, Francesco Rovero, Fernanda Santos, Stephanie Schuttler, Aisha Uduman, Joanna Klees van Bommel, Hilary Young, and A. Cole Burton. Global camera trap synthesis highlights the importance of protected areas in maintaining mammal diversity. *Conservation Letters*, 15(2):e12865, 2022.

[Connor *et al.*, 2023] Thomas Connor, Emilio Tripp, B. J. Saxon, Jessica Camarena, Jesse A. Goodwin, William T. Bean, Dan Sarna, and Justin Brashares. A spatial, closed integrated population model to estimate wildlife population size and structure. *The Journal of Wildlife Management*, 87(7):e22459, 2023.

[Dupont *et al.*, 2021] Gates Dupont, J. Andrew Royle, Muhammad Ali Nawaz, and Chris Sutherland. Optimal sampling design for spatial capture–recapture. *Ecology*, 102(3):e03262, 2021.

[Durbach *et al.*, 2021] Ian Durbach, David Borchers, Chris Sutherland, and Koustubh Sharma. Fast, flexible alternatives to regular grid designs for spatial capture–recapture. *Methods in Ecology and Evolution*, 12(2):298–310, 2021.

[Efford and Boulanger, 2019] Murray G. Efford and John Boulanger. Fast evaluation of study designs for spatially explicit capture–recapture. *Methods in Ecology and Evolution*, 10(9):1529–1535, 2019.

[Efford, 2025] Murray Efford. *secrdesign: Sampling Design for Spatially Explicit Capture-Recapture*, 2025. R package version 2.9.3.

[Feldhamer *et al.*, 2003] George A. Feldhamer, Bruce C. Thompson, and Joseph A. Chapman. *Wild Mammals of North America: Biology, Management, and Conservation*. The Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 2003. Reviewed by Douglas A. Kelt, Journal of Mammalogy 85(6):1237–1239.

[Gibb *et al.*, 2019] Rory Gibb, Ella Browning, Paul Glover-Kapfer, and Kate E. Jones. Emerging opportunities and challenges for passive acoustics in ecological assessment and monitoring. *Methods in Ecology and Evolution*, 10(2):169–185, 2019.

[Glover-Kapfer *et al.*, 2019] Paul Glover-Kapfer, Carolina A. Soto-Navarro, and Oliver R. Wearn. Camera-trapping version 3.0: current constraints and future priorities for development. *Remote Sensing in Ecology and Conservation*, 5(3):209–223, 2019.

[Helton and Davis, 2003] J. C. Helton and F. J. Davis. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. *Reliability Engineering and System Safety*, 81(1):23–69, July 2003.

[Kerry *et al.*, 2022] Rout George Kerry, Francis Jesmar Perez Montalbo, Rajeswari Das, Sushmita Patra, Gyana Prakash Mahapatra, Ganesh Kumar Maurya, Vinayak Nayak, Atala Bihari Jena, Kingsley Eghonghon Ukhurebor, Ram Chandra Jena, Sushanto Gouda, Sanatan Majhi, and Jyoti Ranjan Rout. An overview of remote monitoring methods in biodiversity conservation. *Environmental Science and Pollution Research International*, 29(53):80179–80221, 2022.

[Kleywegt *et al.*, 2002] Anton J. Kleywegt, Alexander Shapiro, and Tito Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002.

[Marino *et al.*, 2008] Simeone Marino, Ian B. Hogue, Christian J. Ray, and Denise E. Kirschner. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254(1):178–196, 2008.

[Oliver *et al.*, 2023] Ruth Y. Oliver, Fabiola Iannarilli, Jorge Ahumada, Eric Fegraus, Nicole Flores, Roland Kays, Tanya Birch, Ajay Ranipeta, Matthew S. Rogan, Yanina V. Sica, and Walter Jetz. Camera trapping expands the view into global biodiversity and its change. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 378(1881):20220232, July 2023.

[Royle and Young, 2008] J. Andrew Royle and Kevin V. Young. A hierarchical model for spatial capture-recapture data. *Ecology*, 89(8):2281–2289, August 2008.

[Royle *et al.*, 2013] J.A. Royle, R.B. Chandler, R. Sollmann, and B. Gardner. *Spatial Capture-Recapture*. Academic Press, 2013.

[Sexton *et al.*, 2013] Joseph O. Sexton, Xiao-Peng Song, Min Feng, Praveen Noojipady, Anupam Anand, Chengquan Huang, Do-Hyung Kim, Kathrine M. Collins, Saurabh Channan, Charlene DiMiceli, and John R. Townshend. Global, 30-m resolution continuous fields of tree cover: Landsat-based rescaling of modis vegetation continuous fields with lidar-based estimates of error. *International Journal of Digital Earth*, 6(5):427–448, September 2013.

[Shapiro, 2003] Alexander Shapiro. Monte carlo sampling methods. In *Handbooks in Operations Research and Management Science*, volume 10, pages 353–425. Elsevier, 2003.

[Steenweg *et al.*, 2017] Robin Steenweg, Mark Hebblewhite, Roland Kays, Jorge Ahumada, Jason T Fisher, Cole Burton, Susan E Townsend, Chris Carbone, J Marcus Rowcliffe, Jesse Whittington, Jedediah Brodie, J Andrew Royle, Adam Switalski, Anthony P Clevenger, Nicole Heim, and Lindsey N Rich. Scaling-up camera traps: monitoring the planet's biodiversity with networks of remote sensors. *Frontiers in Ecology and the Environment*, 15(1):26–34, 2017.

[Sun *et al.*, 2014] Catherine C. Sun, Angela K. Fuller, and J. Andrew Royle. Trap configuration and spacing influences parameter estimates in spatial capture-recapture models. *PLOS ONE*, 9(2):e88025, February 2014.

[Tanwar *et al.*, 2021] Kamakshi S. Tanwar, Ayan Sadhu, and Yadvendradev V. Jhala. Camera trap placement for evaluating species richness, abundance, and activity. *Scientific Reports*, 11(1):23050, November 2021.

[Verweij *et al.*, 2003] Bram Verweij, Shabbir Ahmed, Anton J. Kleywegt, George L. Nemhauser, and Alexander Shapiro. The sample average approximation method applied to stochastic routing problems: A computational study. *Computational Optimization and Applications*, 24:289–333, 2003.

[Wildlife Conservation Society, 2005] Wildlife Conservation Society. *Last of the Wild Project, Version 2, 2005 (LWP-2): Last of the Wild Dataset (Geographic)*, 2005. Archived by National Aeronautics and Space Administration, U.S. Government.

[Wilton *et al.*, 2014] Clay M. Wilton, Emily E. Puckett, Jeff Beringer, Beth Gardner, Lori S. Eggert, and Jerrold L. Belant. Trap array configuration influences estimates and precision of black bear density and abundance. *PLOS ONE*, 9(10):e111257, October 2014.

[Wolters, 2015] Mark A. Wolters. A genetic algorithm for selection of fixed-size subsets with application to design problems. *Journal of Statistical Software, Code Snippets*, 68(1):1–18, 2015.

[Wolters, 2018] Mark A. Wolters. *kofnGA: A Genetic Algorithm for Fixed-Size Subset Selection*, 2018. R package version 1.3.