

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Karppinen, Santeri; Rajala, Tuomas; Mäntyniemi, Samu; Kojola, Ilpo; Vihola, Matti

Title: Identifying territories using presence-only citizen science data : An application to the Finnish wolf population

Year: 2022

Version: Published version

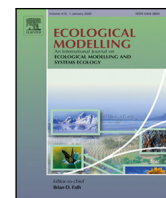
Copyright: © 2022 the Authors

Rights: CC BY 4.0

Rights url: <https://creativecommons.org/licenses/by/4.0/>

Please cite the original version:

Karppinen, S., Rajala, T., Mäntyniemi, S., Kojola, I., & Vihola, M. (2022). Identifying territories using presence-only citizen science data : An application to the Finnish wolf population. *Ecological Modelling*, 472, Article 110101. <https://doi.org/10.1016/j.ecolmodel.2022.110101>



Identifying territories using presence-only citizen science data: An application to the Finnish wolf population

Santeri Karppinen^{a,*}, Tuomas Rajala^b, Samu Mäntyniemi^b, Ilpo Kojola^c, Matti Vihola^a

^a Department of Mathematics and Statistics, University of Jyväskylä, P.O. Box 35, FI-40014, Jyväskylä, Finland

^b Natural Resources Institute Finland (Luke), Latokartanonkaari 9, FI-00790, Helsinki, Finland

^c Natural Resources Institute Finland (Luke), Ounasjoentie 6, FI-96200, Rovaniemi, Finland

ARTICLE INFO

Dataset link: <https://github.com/antiphon/tass-u-intensity>, <https://github.com/skarppinen/tassu-filtering>

Keywords:

Citizen science data
Bayesian statistics
Sequential Monte Carlo
Spatio-temporal model
Territory identification
Presence-only data

ABSTRACT

Citizens, community groups and local institutions participate in voluntary biological monitoring of population status and trends by providing species data e.g. for regulations and conservation. Sophisticated statistical methods are required to unlock the potential of such data in the assessment of wildlife populations.

We develop a statistical modelling framework for identifying territories based on presence-only citizen science data. The framework can be used to jointly estimate the number of active animal territories and their locations in time. Our approach is based on a data generating model which consists of a dynamic submodel for the appearance/removal of territories and an observation submodel that accounts for the varying observation intensity and links the data to the territories. We first estimate the observation intensity using past presence-only observations made by citizens, conditioning on previously known territories. We then infer the territories using a state-of-the-art sequential Monte Carlo method, which extends earlier approaches by allowing for spatial inhomogeneity in the observation process.

We verify our data generating model and inference method successfully in synthetic scenarios. We apply our framework for estimating the locations and number of wolf territories in March 2020 in Finland using one year of confirmed citizen-made wolf observations. The observation intensity is estimated using wolf observation data collected in 2011–2019, conditioning on official territory estimates and data from GPS-collared wolves.

Our experiments with synthetic data suggest that the estimation of territories can be feasible with presence-only data. Our location and territory count inferences for March 2020 based on past data are comparable to the official wolf population assessment of March 2020 by the Natural Resources Institute Finland. The results suggest that the framework can provide useful information for assessing populations of territorial animals. Furthermore, our methods and findings, such as the developed data generating model and the estimation of the spatio-temporal observation intensity can be relevant also beyond the strictly territorial setting.

1. Introduction

Volunteers contribute to many wildlife monitoring programs but standardised monitoring schemes are available for only a small number of taxa in a few countries (Gregory et al., 2005; Isaac, 2014). Citizens, community groups and local institutions participate in biological monitoring of population status and trends by providing species data e.g. for regulations and conservation (Conrad and Hichley, 2011; Lawrence, 2006). The involvement of citizens as data collectors has demonstrated its ability to gather massive amounts of data at a spatio-temporal scale unattainable by research teams and state authorities active in biodiversity monitoring (Silvertown, 2009). For instance, in many European countries, hunters are integrated as data-providers in

wildlife management structures that are intended to support sustainable harvest (Bragina et al., 2015; Cretois et al., 2020; Linnell et al., 2015).

Statistical developments in data integration as well as more rigorous protocols for data collection are needed to unlock further the potential that volunteers' data holds (Cretois et al., 2020; Isaac, 2014). The statistical interpretation of citizen-collected data faces problems less frequently encountered in traditional scientific research. For example, the spatio-temporal sampling effort of citizens is usually not known nor controllable. Sophisticated methods that model the data collection process offer the greatest potential to estimate e.g. timely trends (Isaac, 2014).

In this paper, we propose a statistical modelling framework that can be used to make inferences about animal populations with territorial

* Corresponding author.

E-mail address: santeri.j.karppinen@jyu.fi (S. Karppinen).

behaviour, using observations reported voluntarily by citizens. More specifically, our focus is on the following scenario:

- Citizens report presence-only observations of territorial animals. Each observation consists of a GPS coordinate and an approximate time stamp.
- We wish to estimate the number and locations of the animal territories within some area and time interval using the data collected by the citizens.
- Prior knowledge on the territorial behaviour of the species is assumed to exist in the form of a typical territory size and on the rate of appearance and disappearance of territories.

Because of the high spatio-temporal variability common in citizen science observation processes, modelling of the varying sampling effort, that is, the observation intensity, is a crucial first part of our framework. We take this variation into account by modelling the intensity based on past data. This yields intensity functions that capture the spatio-temporal variation in the observations reported by citizens. These functions are then fed into a data generating model consisting of two submodels that model the appearance and removal of territories, and the generation of citizen science observations from the territories, respectively. To estimate the number of territories and their locations, we use the latest available citizen science data and perform Bayesian inference for the data generating model.

The data generating model jointly approximates the evolution of the number of active territories and their locations in time, characterised by a sequence of posterior distributions conditioned on observation sets of increasing size. In the engineering literature, similar models are called ‘tracking models’ (cf. Goodman et al., 1997). Indeed, the inference algorithm we develop is a Rao-Blackwellised particle filter similar to those developed for tracking (Särkkä et al., 2007; Vihola, 2007). We further elaborate these methods by employing a state-of-the-art optimal resampling of Fearnhead and Clifford (2003), and further refine the inference algorithm so that it can incorporate the spatial inhomogeneity arising from our observation model.

Our data generating model is similar to dynamic occupancy models (Royle and Kéry, 2007) and open N-mixture models (Zhao et al., 2017) in the sense that it has a latent process model for the appearance and disappearance (“occupancy”) of animal territories, and a variable observation intensity. However, unlike in the work of Zhao et al. (2017), the principal objects of analysis in our model are animal territories rather than individual animals. In addition, our model does not assume a fixed set of potentially occupied sites but operates in continuous space, where territories are delineated without a pre-defined grid. Finally, our model is formulated in continuous-time, which allows the estimation of the state of the population at any time points within the interval of interest. For example, our model can be used to track the state of the population at daily or weekly time steps. In contrast, the methods of Royle and Kéry (2007) and Zhao et al. (2017) operate in discrete time, and are typically used for annual data with a considerably smaller number of time steps.

The motivation for the development of our modelling framework has been to aid in the task of assessing the Finnish wolf (*Canis lupus*) population, although the framework can be relevant for other territorial species as well. Currently, the Finnish wolf population is assessed annually in March by the Natural Resources Institute Finland (Luke). In the assessments, wolf observations provided by citizens from the beginning of August to the end of February are combined with non-invasive genetic samples, tracks of GPS collared wolves and records of known mortality (Kojola et al., 2018). The assessments are carried out in two phases. In the first phase a panel of experts conducts a systematic review of all the data and judges territory boundaries that are potentially occupied by wolf packs or pairs in March. In the second phase, a Bayesian state–space model is used to infer the number of wolves living in each territory by combining wolf observations, DNA-recaptures and known mortality (Heikkinen et al., 2020). In particular,

we envision that the developed framework can work as a useful tool in the first phase, providing a statistical look at the citizen science data and an aid in judging the territory boundaries.

We examine the performance of the developed particle filter with a sequence of simulation experiments, where we start from simple simulated conditions and work towards conditions that resemble more closely our concluding experiment, which is a realistic situation that could be faced in the assessment of the Finnish wolf population. Here, we use previous estimates of territory locations and citizen-provided observations to estimate the spatio-temporal variation of the conditional probability of wolf observations given known existence of wolf territories. Even though similar approaches have been used for species abundance estimations (e.g. Renner et al., 2015; Ver Hoef et al., 2021; Tang et al., 2021), the conditioning requirement provides a novel challenge. Using the results of said intensity modelling, we apply our data generating model to a real data set consisting of wolf observations made by Finnish citizens between April 2019 and March 2020. We estimate the number and locations of wolf territories and compare the result to the official estimates by the Natural Resources Institute Finland (Luke) which are based on the method discussed above.

The main contributions of this paper are as follows. First, we believe that the developed framework is of interest in assessing populations of territorial species using presence-only citizen science data. We focus on the application to wolves, but our methods are readily adaptable for other territorial species. Second, we believe that the observation intensity estimation is of its own independent interest, because it addresses the problem of estimating the conditional spatio-temporal intensity of presence-only citizen science observations. Third, from a methodological point of view, the developed data generating model and particle filter might be relevant also in the context of ‘general purpose’ target tracking (e.g. Vihola, 2007; Särkkä et al., 2007) applications where a spatially varying observation process is needed.

2. Materials and methods

The general modelling framework proposed in this paper can be summarised into four successive analysis steps numbered from one to four. The flowchart in Fig. 1 depicts their dependencies and relation with each other, highlighting the inputs, outputs and datasets associated with each step. The following subsections will explain how we apply the framework in the context of wolf territory estimation. Section 2.1 discusses the Datasets A, B and C. Section 2.2 then describes the data generating model, which motivates the intensity estimation consisting of steps 1 and 2, which in turn are discussed in Sections 2.3 and 2.4, respectively. Section 2.5 describes step 3 of the analysis, the statistical inference based on the data generating model, using a particle filter we have developed for the problem. Finally, in Section 2.6 we conclude with a description of step 4 where we extract the number and locations of the territories from the output of the particle filter.

2.1. Data

In this Section, we will discuss the Datasets A, B and C seen in the framework of Fig. 1. In summary, the Datasets A and B contain past data used in the construction of the data generating model, and Dataset C contains the latest data to be processed by the particle filter. Each datum in Datasets A and C is a spatio-temporal point, that is, it has the form (t, y) , where t is the time of observation, and y is a two-dimensional point on a domain we denote by D_y . The difference between Datasets A and C is that Dataset A is past data, and Dataset C corresponds to the latest data we wish to infer the territories with. In contrast, Dataset B is more heterogeneous and contains all additional data such as covariates and expert knowledge required in the construction of the data generating model.

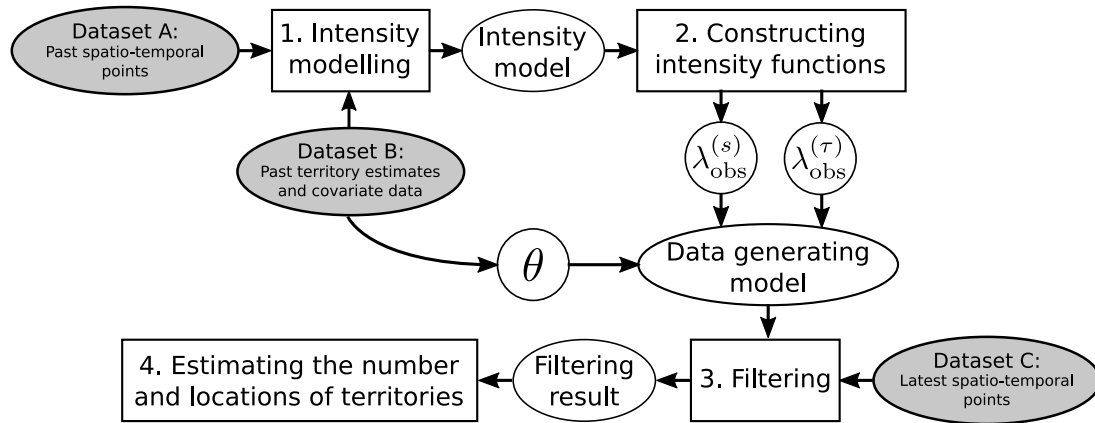


Fig. 1. A flowchart of the proposed general modelling framework. The rectangles depict the analysis steps 1–4. The elliptical shapes denote inputs and outputs of the analysis steps. Input datasets are marked with the gray fill. The symbols $\lambda_{\text{obs}}^{(s)}$ and $\lambda_{\text{obs}}^{(\tau)}$ denote the intensity functions and θ stands for the other parameters of the data generating model.

2.1.1. Datasets A and C

We extract the spatio-temporal points in Datasets A and C from a digital large carnivore observation database named “Tassu” (meaning a “paw” in Finnish) (Kojola et al., 2018). The observations enter the database through a network of approximately 2000 large carnivore contact persons (LCCPs), who are nominated by management associations and educated by the Finnish Wildlife Agency and Luke in the biology, ecology and movement behaviour of wolves as well as footprint identification. There are, however, no formal exams used in the nomination process.

The LCCPs have their own local trusted network of people who report their observations of wolves to the LCCP. These networks consist mostly of hunters that are proficient in identifying wolves based on sightings, tracks, prey kills and camera-trap documents. In addition, it is in principle possible for any citizen to report their observations, since the contact details of the LCCPs are publicly available and known in local rural societies (Pellikka and Hiedanpää, 2017). However, the networks of the LCCP are particularly relevant for wolf sightings in snow-free conditions because such observations usually cannot be verified afterwards.

Wolf observations found to be valid by the LCCPs are saved into the Tassu database. Each saved datum includes information about the time and location of the observation, the type of observation (such as wolf track, sighting, droppings, game camera photograph, prey kill site or livestock predation) and the estimated count of wolves observed simultaneously. The count estimate is based on the judgement of the LCCP based on the information available. Since the observations saved to the database are subject to the confirmation of the LCCP (possibly days after the initial report), we consider the observation times to be accurate on a daily granularity. In total, Datasets A and C contain all observations from the Tassu database that reported two or more wolves between January 2011 and March 2020. Since the purpose of our framework is to infer the number and locations of wolf territories, we only focus on observations that report more than one wolf, since this indicates that the observed wolves form a wolf pack and very likely exhibit territorial behaviour. In contrast, observations of single wolves can originate from lone, vagrant wolves, that do not yet maintain a territory. Furthermore, for simplicity, we make no distinction for data points with different observation types; we regard each observation simply as a spatio-temporal point. We return to this matter in the discussion.

We split the data such that Dataset C contains the observations made between April 1st 2019 and March 31st 2020, and Dataset A the observations before this. The locations in Dataset C are illustrated in Fig. 2 (top left). The domain of the locations, D_y , is mainland Finland south of the reindeer husbandry region in the north. The wolf territories in the reindeer husbandry region are few and short-term owing to

lethal control that is justified by the prevention of damages to reindeer husbandry.

We organise the Datasets A and C according to so-called “wolf years”. A wolf year starts April 1st and ends in March 31st of the next year. Hence, Dataset C consists of the observations made during the wolf year 2019–2020. The organisation of the data to wolf years has two reasons. First, as described in the introduction, the annual Finnish wolf population assessments describe the state of the wolf population in March. Second, the data indicate that the highest observation intensity is reached during the winter season and declines towards the spring. Year-by-year changes in the observation activity are expected to occur between the winter seasons, rather than between calendar years. We will also use the term “wolf month” to refer to the months within a wolf year such that the first wolf month corresponds to April, the second to May, and so on.

2.1.2. Dataset B

Dataset B contains two kinds of information. Most importantly, Dataset B contains information about past known wolf territories until March 2019, that is, before the wolf year associated with Dataset C. In addition, Dataset B also contains covariates.

Dataset B is primarily used in the intensity modelling described in Section 2.3, but also for setting certain parameters in the data generating model. The details on how the data sources described below are used in the intensity modelling are given in Section 2.3. The relation of Dataset B to the parameters of the data generating model is discussed in the results of Section 3.3.

The information about past wolf territories was constructed from two sources of data, independently of Datasets A and C. We call the resulting territories ‘auxiliary territories’. The first source consists of the space–time trajectories of 34 GPS-collared wolves that were tracked between 2011 and 2019. The transmitters in the collars stored the wolf’s position at one- or four-hour intervals, depending on the season. The capture, handling and immobilisation protocols of these wolves are described in Kojola et al. (2016).

We assumed that each collared wolf was part of a wolf territory. The trajectories contain outlier recordings, such as test measurements at a lab or ‘glitch’ jumps of hundreds of kilometres occurring due to device malfunction or other reason. Some trajectories also cover two clearly separate territories. We therefore preprocessed the data as follows. First, the recorded GPS trajectories were divided into separate, contiguous trajectories at temporal jumps of more than two weeks or spatial jumps of more than 100 km. Trajectories less than 24 h were rejected. Second, each contiguous trajectory was processed by assigning to each trajectory point a probability of being an outlier. The probability was given by the velocity density $v \sim \text{Exp}(28.8)$ with median at 20 km/h, multiplied by the function $w(d_c = x) = 1(x <$

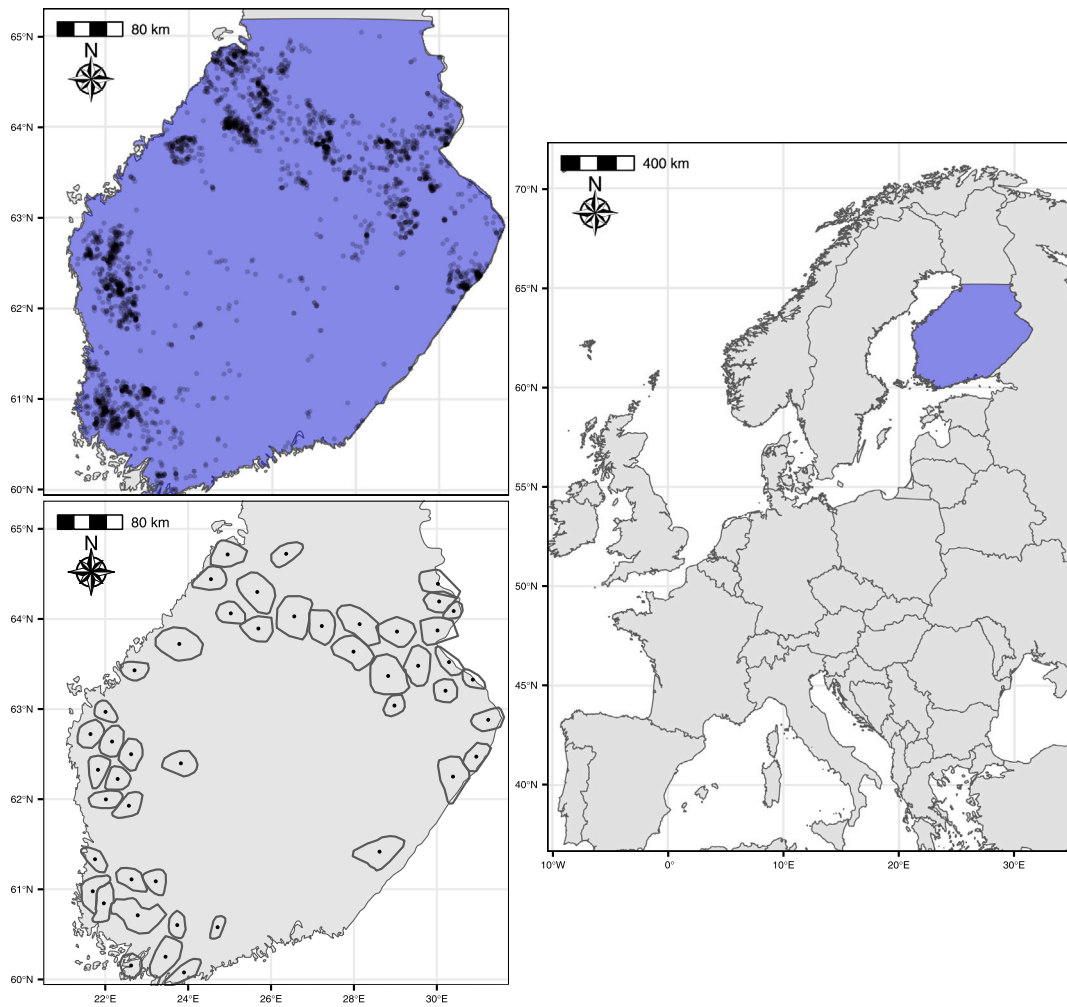


Fig. 2. Top left: The observed locations in Dataset C, that is, the locations of the Tassu observations (black points) that reported two or more wolves from April 2019 to March 2020. The blue overlay shows the domain of interest (the study area). Bottom left: Wolf territories found in the wolf population assessment of April 2019 by the Natural Resources Institute Finland. The point within each territory represents the centroid of the territory polygon. Right: The study area highlighted on a map of Europe. The distance scales are approximate due to coordinate transformations applied in drawing the maps.

$(5\sigma) + 1(x > (5\sigma)) \exp[-0.5(x - 5\sigma)^2/40000^2]$, where d_c denotes a point's distance from the trajectory's centre of mass and σ is the 90%-truncated standard deviation of the centre of mass-distances. Points with probability less than 50% were excluded from the trajectory. Finally, the first and second steps were repeated to account for significant gaps after the outlier detection second step. From each remaining trajectory, an auxiliary territory was constructed as a polytope in $D_y \times T_B$, where T_B denotes the time span 2003-03-04–2019-03-31, by taking the Cartesian product of the convex hull of the spatial locations and the time interval of the trajectory. In total, 59 auxiliary territories were constructed from the GPS trajectories, covering approximately 74,000 km² and with time spans that add up to approximately 36.5 years.

The second type of auxiliary territories were constructed based on expert knowledge using the official population assessments of Luke from 2017 onwards. The assessments include estimates of active wolf pack territories as polygons in D_y , during March of the corresponding years. Fig. 2 (bottom left) shows the active pack territory location estimates of experts in the assessment of March 2019. We assumed that these territories were active also during January and February. We then constructed polytopes in $D_y \times T_B$ as the Cartesian products of the polygons and January-March-intervals of each year between 2017–2019. The resulting 203 expert judgement auxiliary territories covered approximately 180,000 km² with time spans that add up to approximately 35.3 years.

The additional covariates in Dataset B consist of two datasets. The first of these is the CORINE land cover data for 2018 (Finnish Environment Institute SYKE, 2018). The dataset comes as a raster covering Finland and contains an approximate land use class (e.g. river, small road) for each of its 20 by 20 metre cells. The original 49 classes were first reclassified down to 8: Residential areas, other build areas, roads, cultivated fields, lakes and rivers, swamps and other wetlands, closed forests, and open forests. For each of the 8 classes, we aggregated their frequency in 1 km² cells, and to slightly reduce the amount of zeros, applied smoothing with a Gaussian blur with standard deviation 3 km. Each cell of the resulting 8-layer raster stack then contained a vector giving the (smoothed) frequencies of each land use class in (and near) the 1 km² cell. Since the resulting vector for each cell k , $[\text{Corine}_{k,1}, \dots, \text{Corine}_{k,8}]$, is (nearly) a simplex, we dropped the first class, residential areas, and kept the remaining 7 as frequencies. A log-ratio transformation, which is a popular approach in compositional data analysis, might have been more suitable here but was not done due to numerous zeros in all classes.

The CORINE road information capture larger streets and highways, and to describe accessible forest areas we computed an additional forest road frequency variable to Dataset B. This variable is derived from the national road and street database Digiroad (Finnish Transport Infrastructure Agency, 2021). From the database we extracted the polyline feature class '12', roads and paths traversable by offroad vehicle. We

binarised the polylines on to the 20 by 20 metre cells of the CORINE land cover data, and then computed the frequencies of those cells on a 1×1 km raster. We denote the value of this variable in the k th 1 km^2 cell by $forestreet_k$.

2.2. Data generating model

In this section, we discuss the data generating model we have developed for citizen science observations of a territorial species. The model we have developed is more general than the instance of it that we use for modelling the wolf data. Therefore, this section will also highlight certain modelling decisions we make in the present application. Furthermore, the data generating model we describe here is ‘ideal’ in the sense that it must be approximated further to be tractable for our inference method. We will discuss this in more detail in Section 2.5 that is devoted to the filtering algorithm. For the interested reader, the mathematical details of the general data generating model are given in Sections 1.1 and 1.2 of the supplementary material.

The data generating model consists of two submodels, the birth and death process and the observation model, which we will discuss in Sections 2.2.1 and 2.2.2, respectively. In summary, the birth and death process models how new territories emerge and disappear, and the observation model describes how each existing territory produces citizen science observations (spatio-temporal points as in Dataset C).

2.2.1. Birth and death process

Our model assumes that the territories of interest exist and emerge within a domain denoted by $D_\mu \subset \mathbb{R}^2$, with $D_\mu \subseteq D_y$. The location of the territory i is represented by its centroid, $\mu_i \in D_\mu$, which is assumed to be constant in time. New territories emerge within D_μ with the instantaneous birth intensity $\lambda_b(u)N_u + \lambda_{b0}$ where $\lambda_b(u)$ is the (known) birth intensity function and N_u stands for the number of existing territories at time u . The function $\lambda_b(u)$ can be interpreted as birth intensity per each existing territory. The baseline birth intensity parameter λ_{b0} , on the other hand, models additional birth intensity due to external factors such as inflow from outside D_μ . For modelling of the wolf data, we simplify $\lambda_b(u)$ to a constant, denoted by λ_b , and set $\lambda_{b0} = 0$. As a new territory emerges, its centroid follows the uniform distribution on D_μ .

Similarly, each existing territory disappears with the instantaneous death intensity $\lambda_d(u)$, that is, the total instantaneous death intensity induced by all territories equals $\lambda_d(u)N_u$. In case of the wolf data, we fix $\lambda_d(u)$ to a constant that we denote by λ_d . Therefore, the lifetime of a single territory follows an exponential distribution with mean λ_d^{-1} . Furthermore, in a similar fashion as was done in Vihola (2007), we ‘symmetrise’ the birth and death process by setting $\lambda_b = \lambda_d = \lambda_{bd}$, where λ_{bd} then remains the only birth/death intensity parameter. This minimises the bias in the birth and death process, and a priori leads to a constant conditional expectation for the number of territories in time.

The initial distribution of the model is a joint distribution of the number of territories and the locations of their centroids. The initial locations of the territory centroids can either be distributed uniformly on D_μ or subject to Gaussian error (truncated to D_μ) around some location estimate.

2.2.2. Observation model

The observation model, conditional on the territory locations and lifetimes generated by the birth and death process, describes how each territory with its centroid on D_μ produces citizen science observations.

The ‘baseline’, underlying model for an observation from a single territory i that exists at any given time point, is bivariate normal $N(\cdot; \mu_i, \Sigma_{\text{obs}})$, where Σ_{obs} describes the size and shape of the territories. We assume that the territories are roughly circular in shape by setting $\Sigma_{\text{obs}} = \sigma_{\text{obs}}^2 I$, where I denotes the 2×2 identity matrix and $\sigma_{\text{obs}} > 0$ is a standard deviation related to the territory size.

It is useful to interpret this territory model using the circular contours of the distribution $N(\mu_i, \sigma_{\text{obs}}^2 I)$. A circle of radius

$$\sqrt{\chi_{\alpha,2}^2} \sigma_{\text{obs}} \quad (1)$$

centred at μ_i is assumed to enclose the instantaneous location of the wolves belonging to the territory with probability α . Here, $\chi_{\alpha,2}^2$ corresponds to the $100\alpha\%$ quantile of the chi-squared distribution with two degrees of freedom.

Because of the temporal and spatial variability inherent to citizen science observation processes, the observation model modulates the number of observations produced from the territories based on a temporal intensity function $\lambda_{\text{obs}}^{(\tau)}$, defined on a time interval of interest $[0, T)$, and a spatial intensity function $\lambda_{\text{obs}}^{(s)}$ defined on the domain of the observed locations, D_y . These intensity functions are assumed spatio-temporally separable, since there is limited data for their estimation, which is further discussed in Section 2.3. The values of these functions are tied to the number of observations the territories produce in time and space. For constant functions $\lambda_{\text{obs}}^{(s)}(y) = l_x \in [0, \infty)$ and $\lambda_{\text{obs}}^{(\tau)}(u) = l_t \in [0, \infty)$, our model assumes that the expected number of observations that a single territory produces on D_y in a unit of time is approximately $\lambda_{\text{obs}} l_x l_t$, where λ_{obs} is a scalar multiplier for the intensity functions. The intensity functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ are assumed known (fixed), and we discuss their estimation in Sections 2.3–2.4.

In addition to the observations originating from the territories, the observation model also accommodates so called ‘clutter’ observations, that are understood as ‘erroneous’ observations not originating from actual territories. These observations are assumed to be distributed uniformly on D_y and their intensity is likewise modulated by $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ but multiplied by a different scalar parameter, λ_c . The relative values of the scalar multipliers λ_{obs} and λ_c can be used to model the rate of the total number of observations believed to originate from the territories.

Mathematically, conditional on the territory locations and lifetimes, our observation model defines a three-dimensional inhomogeneous Poisson process in time and space, whose intensity function is given in Equation (5) of the supplementary material. The data generating model parameters in Fig. 1 are given by $\theta = (\lambda_{\text{obs}}, \lambda_{\text{bd}}, \lambda_c, \sigma_{\text{obs}}, D_y, D_\mu)$ and the intensity functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$.

2.3. Observation intensity modelling

The following two sections discuss how we estimate the intensity functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ in the observation model of Section 2.2.2. This section focuses on step 1 of the modelling framework in Fig. 1, detailing the intensity model we fit to the past Datasets A and B discussed in Section 2.1. The primary data for this step are the spatio-temporal points in Dataset A discussed in Section 2.1.1. We assume that each of these observations originated from an active wolf territory. In this section, we denote by $\psi = \{[s_i; t_i]\}$ the spatio-temporal point pattern of the wolf observations in Dataset A with locations $s_i \in D_y$ and dates $t_i \in T_A = [2011-01-01, 2019-03-31]$.

We assume the arrival of Tassu reports ψ can be approximated by an inhomogeneous Poisson process (Illian et al., 2008). Note that reporting depends on two consecutive events: An observer is at a territory, and they make and report an observation. The data contains no information if an observer was on a territory but did not observe wolf activity. The observer and reporting intensities are therefore confounded. The situation is notably different from presence–absence citizen science data, such as for birding analysed with point processes by Tang et al. (2021), as in addition to absences not being measured, the observers cannot be assumed to have been actively looking for wolf activity in the first place. Our situation is more akin to presence-only analysis (Renner et al., 2015; Ver Hoef et al., 2021), with the nuance that instead of estimating species abundances the goal is to estimate the connection between a wolf territory’s presence and the emergence of the Tassu-reports. To account for the conditioning on a wolf territory’s presence

in this estimation, we constrain the Tassu-observations to the auxiliary territories in Dataset B, detailed in Section 2.1.2.

Given the observations on the auxiliary territories, we estimate the intensities by aggregating the observations and then using standard Poisson regression in the generalised additive models (GAM) framework. More specifically, we fit the model on observation units given by spatio-temporal grid cells $C_k = \{V_k \times t_k\}$ with temporal resolution $|t_k| = 1$ month and spatial resolution $|V_k| = 1 \text{ km} \times 1 \text{ km}$. From this discretisation of $D_y \times T_A$ we only consider the subset of cells that intersect the auxiliary territories, which is ≈ 11.9 million cells, and then count the Tassu reports in each such cell, denoted by $n_k = \#(\psi \cap C_k)$, resulting in 4816 non-empty cells. The auxiliary territory (polytope) which a cell C_k overlaps is denoted by A_k . For each cell C_k we let u_k and y_k represent its centroid in time and space, respectively, and finally define $\lambda_k^{(\tau)} := \int_{t_k} \lambda_{\text{obs}}^{(\tau)}(u) du$ and $\lambda_k^{(s)} := \int_{V_k} \lambda_{\text{obs}}^{(s)}(y) dy$.

Then the counts in the grid cells are modelled with a Poisson regression model of the following form

$$n_k \sim \text{Poisson} \left(\frac{\lambda_k^{(\tau)} \lambda_k^{(s)}}{|A_k|} \right) \quad (2)$$

$$\log(\lambda_k^{(\tau)} \lambda_k^{(s)}) = \text{month}(u_k) + \text{year}(u_k) + X(y_k)\beta + \text{smooth}(y_k)$$

$$X(y_k) = [1_{\text{expert}}(A_k), \text{Corine}_k, \text{forestroa}_k]^T.$$

The offset $|A_k|$ is the area of the territory that grid cell C_k belongs to, and accounts for an assumption of the instantaneous location of the wolves following a uniform distribution within A_k , i.e. a pack on a larger territory is harder to observe.

The ‘year’ and ‘month’ effect were included as factors to model seasonal effects and year to year differences. We used ‘wolf years’ as discussed in Section 2.1.1 for the factors ‘year’ and ‘month’. We did not include weather station information (e.g. snow depth), mainly because their monthly aggregates are highly correlated with month-effects but also in order to avoid spatio-temporal interaction terms to reduce model complexity given we have only $<1\%$ non-zero units.

The spatial effects are modelled with covariates X and a residual smooth term. The indicator $1_{\text{expert}}(A_k)$ was included to adjust for potential discrepancies between the different types of auxiliary territories (GPS tracks & expert estimates). The numerical covariates $\text{Corine}_k = [\text{Corine}_{k,2}, \dots, \text{Corine}_{k,8}]$ and forestroa_k were described in Section 2.1.2 and capture environmental variability.

The estimation was carried out using the statistical software R and the GAM function `mgcv::bam` with a smooth term ‘smooth(y_k)’ defined as a tensor product `te`-term in x and y coordinates of the cell centroid y_k . The smoothness penalty choice was left to the default which is generalised cross validation (Wood, 2017).

2.4. Computing intensity functions for the data generating model

Next, we discuss how we compute the intensity functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ based on the intensity model of Section 2.3 fit to the past Datasets A and B. This section focuses on step 2 in the modelling framework of Fig. 1. The aim is to obtain intensity functions for the data generating model that anticipate the spatio-temporal intensity of the observations in Dataset C. We model the intensity functions as piecewise constant such that $\lambda_{\text{obs}}^{(\tau)}$ takes on the value c_{t_i} during wolf month i , $i = 1, \dots, 12$, and $\lambda_{\text{obs}}^{(s)}$ takes on the value c_{V_j} in each cell $V_j \in D_y$. In summary, the c_{t_i} ’s and c_{V_j} ’s are computed by using quantities calculated from the predictions of the intensity model (2) in Eq. (4) below.

The computation proceeds as follows. First, assume that the intensity model (2) is fit using the Datasets A and B. Then, using the fitted model, we predict the spatial effect $\lambda_k^{(s)}$, excluding the term $1_{\text{expert}}(A_k)$ for all 1 km^2 grid cells $V_k \in D_y$. This grid is then smoothed using Gaussian blur, and we denote the value in the smoothed grid cell V_k by $\tilde{\lambda}_k^{(s)}$. More specifically, we use a ‘border preserving’ Gaussian blur that only smooths cells that are within D_y and normalises the blur weights in the smoothing window such that only non-zero intensity

values contribute to the smoothed grid. We use σ_{obs} as the standard deviation in the Gaussian blur and set the window size to the first integer larger than $2\sigma_{\text{obs}}$ (in kilometres). We report the σ_{obs} value used in this step together with the results of Section 3. The smoothing of the predicted grid is motivated by an assumption of smoothness that our inference method places on the spatial intensity function. We will discuss this (Assumption C) in more detail in Section 2.5.

After computing the predicted and smoothed spatial effect, we also predict the temporal effect for the wolf year associated with Dataset C by setting

$$\tilde{\lambda}_i^{(\tau)} = \exp(\beta_{\text{year}}^* + \beta_{\text{month},i}),$$

where $\tilde{\lambda}_i^{(\tau)}$ is the predicted temporal effect for wolf month $i = 1, 2, \dots, 12$ during the wolf year of interest. Here, β_{year}^* corresponds to a predicted (wolf) year regression coefficient obtained by running a linear regression on the previous (wolf) years’ regression coefficients (available from fitting model (2)). The coefficients $\beta_{\text{month},i}$, on the other hand, correspond to the estimated (wolf) month coefficients from model (2). The linear regression for the year coefficients was carried out to take into account a slight increasing trend in the yearly regression coefficients of model (2).

To compute the distinct values c_{t_i} and c_{V_j} that the piecewise constant functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ take, we match the predicted intensities such that

$$\frac{\tilde{\lambda}_i^{(\tau)} \tilde{\lambda}_j^{(s)}}{|A_j|} = \int_{t_i} \int_{V_j} \lambda_{\text{obs}}^{(\tau)}(u) \lambda_{\text{obs}}^{(s)}(y) / |A_j| dy du, \quad (3)$$

where A_j is defined as in Section 2.3.

Evaluating the integral (3) and taking the logarithm yields

$$\log(\tilde{\lambda}_i^{(\tau)}) + \log(\tilde{\lambda}_j^{(s)}) = \log(|t_i|) + \log(|V_j|) + \log(c_{t_i}) + \log(c_{V_j}).$$

To obtain equality in this equation, c_{t_i} and c_{V_j} can be chosen such that

$$\begin{aligned} c_{t_i} &= \exp(\log(|t_i|^{-1}) + \log(\tilde{\lambda}_i^{(\tau)}) - K) \\ c_{V_j} &= \exp(\log(|V_j|^{-1}) + \log(\tilde{\lambda}_j^{(s)}) + K), \end{aligned} \quad (4)$$

where K is a constant that needs to be chosen. Our choice for K is

$$K = -\max_j \{\log(|V_j|^{-1}) + \log(\tilde{\lambda}_j^{(s)})\},$$

which scales $\lambda_{\text{obs}}^{(s)} \in (0, 1]$.

Finally, for the purposes of Sections 3.3 and 3.4 we note that when $\lambda_{\text{obs}}^{(s)}$ is computed as discussed above, one of the territory centroids seen in Fig. 2 is outside the domain of $\lambda_{\text{obs}}^{(s)}$. Therefore, in the case of the wolf territory estimation, we additionally widen the domain of $\lambda_{\text{obs}}^{(s)}$ by 27 kilometres at the boundaries by repeating the maximal intensity value found in the neighbouring cells of the borders.

2.5. Filtering algorithm and its constraints

This section discusses step 3 of the framework in Fig. 1, the statistical inference of the territories given Dataset C, assuming the data generating model of Section 2.2. This section gives an overview of the inference, and the mathematical details are given in Section 1.6 of the supplementary material. In summary, given Dataset C and the data generating model of Section 2.2 with fixed parameters and intensity functions, the inference procedure outputs a sequence of joint filtering distributions of the territory centroid locations and their number at chosen times $t_1 < t_2 < \dots < t_n$ that are also input to the procedure. The filtering distribution at time t is the distribution of the locations of the territory centroids and their number, conditional on the observations of Dataset C up to time t .

The inference method we use is somewhat involved and computationally intensive, and based on sequential Monte Carlo/particle filtering (cf. Doucet et al., 2000). More specifically, the method is a Rao-Blackwellised particle filter similar to the works of Särkkä et al. (2007),

Vihola (2007), but employing a state-of-the-art optimal resampling algorithm developed by Fearnhead and Clifford (2003).

A summary of the method's operation can be given as follows. Denote by $\tilde{y}_{1:K} = (\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_K)$ the K temporally ordered observations with observation times $\tilde{t}_1 < \tilde{t}_2 < \dots < \tilde{t}_K$. Here, the input timepoints $t_1 < t_2 < \dots < t_n$ are among the \tilde{t}_i 's and each \tilde{y}_i is a spatial location from Dataset C, or $\tilde{y}_i = \emptyset$ (see also discussion on Assumption A below). The method works by processing the observations $\tilde{y}_{1:K}$ sequentially, such that the processing of \tilde{y}_k yields the filtering distribution at time \tilde{t}_k . During the algorithm, each filtering distribution is characterised by a set of M weighted particles. More specifically, at time index $k - 1$, each of the M particles represents a hypothesis about the locations and number of territory centroids on D_μ , conditional on the observations $\tilde{y}_{1:k-1}$ seen so far (with $\tilde{y}_{1:0}$ understood as the empty set).

The observation \tilde{y}_k is processed such that first a set of possible territory birth, territory death and observation association outcomes is built, which consist of all one step 'futures' that could happen for any existing hypothesis at time index $k - 1$, conditional on \tilde{y}_k and the time passed since \tilde{y}_{k-1} . This set of outcomes is then probabilistically pruned using the optimal resampling algorithm of Fearnhead and Clifford (2003), which yields a set of M chosen outcomes and their normalised weights. Finally, based on the chosen outcomes and the previous hypotheses, a new set of M updated and weighted hypotheses (particles) is constructed by adding and deleting territories and 'associating' \tilde{y}_k (if $\tilde{y}_k \neq \emptyset$) to a territory using an approximate Kalman filter update. The filtering distribution at time index k is characterised by these M particles. The process then repeats for the observation \tilde{y}_{k+1} .

The ideal data generating model of Section 2.2 is time-discretised and approximated before the filter can be used. The approximations used can be justified by introducing a set of additional assumptions, some of which are primarily computational, and some of which help reduce the discrepancy between the approximate and the ideal model. The following list highlights these assumptions, which we will refer to as Assumptions A to D.

- (A) The observed data can be processed sequentially, one at a time. In other words, each datum has an associated time, and the observation times are strictly increasing.
- (B) Compared to the rate of observations arriving from the territories, birth and death events of territories are rare.
- (C) The spatial intensity function $\lambda_{\text{obs}}^{(s)}$ is 'smooth'/'slowly varying' with respect to the territory size parameter Σ_{obs} . This means that for any centroid $\mu \in D_\mu$, $\lambda_{\text{obs}}^{(s)}(\mu)$ is a good approximation for $\lambda_{\text{obs}}^{(s)}(x)$ in the region where the distribution $N(\mu, \Sigma_{\text{obs}})$ has most of its probability mass.
- (D) For most territories, the territory centroid μ_i is not close to the boundary of D_μ , in the sense that a region of high probability of $N(\mu_i, \Sigma_{\text{obs}})$ is contained within D_μ .

We conclude this section with a brief discussion on these assumptions. Assumption A is satisfied for many datasets that are collected in real time. However, as mentioned in Section 2.1, the observation times in Dataset C are pooled with a granularity of one day. In order to make Assumption A hold, we introduce a preprocessing step before the filtering that artificially disperses the daily pooled observations in time, generating a 'pseudotime' for each observation within the day that it occurred. This step introduces a bias, which is small, since the arrival intensity of the observations still remains practically the same as with the pooled data. The preprocessing of the data is related to the time discretisation of the model, which we make fine enough so that during filtering we may assume that practically at most one territory birth or death may occur during each time-discretised interval, and that the time-discretised model approximates the ideal model sufficiently well. We ensure this by introducing 'discretisation points' to the dataset that contain no spatial location (that is, $\tilde{y}_k = \emptyset$). For further discussion on these matters, see Sections 1.3–1.4 in the supplementary material.

Assumption B is necessary for the identification of the territories based on the data. In the present application, Assumption B holds since the births and deaths of wolf territories are relatively rare events on the daily timescale at which the observations in Dataset C arrive.

Assumption C is necessary for approximating certain intractable integrals that arise in the filtering algorithm and are related to the spatial intensity function $\lambda_{\text{obs}}^{(s)}$. In Section 2.4 we described a smoothing step for the predicted spatial grid, which was carried out in order to satisfy Assumption C.

Finally, Assumption D arises since our method does not involve explicit edge correction. The computations in the particle filter are approximate for territories and observations close to the boundary of the finite domain D_μ . This may entail some bias, which is small under Assumption D. We investigate empirically the bias caused by the edge effect in Section 3.2.

2.6. Extracting the number and locations of the territories from the output of the particle filter

This section focuses on the final step of the framework of Fig. 1 and describes how we extract the number and locations of the territories from the filtering result. Our estimate for the number of territories at time t is the marginal filtering distribution of the number of territories at time t , computed as follows. Denoting by n_i , $i = 1, 2, \dots, p$ the unique numbers of territories found among the M particles at time t , the marginal filtering distribution for the number of territories consists of the tuples $(n_1, \pi_1), (n_2, \pi_2), \dots, (n_p, \pi_p)$, where π_i is the sum of the normalised weights of the particles having exactly n_i territories. In Section 3.4, we will summarise these distributions by taking their mean, mode and standard deviation, and by computing probability intervals I_α , that is, intervals whose end points are given by the $(100 - \alpha)\%$ and $\alpha\%$ quantiles of the distribution, where α is a given percentage point.

A common way to visualise the probabilistic location information of an unknown number of objects is to plot the so-called probability hypothesis density (Goodman et al., 1997) (PHD) of the filtering distribution, which in the present context corresponds to the expected intensity of territory centroids. More specifically, the PHD is defined for the territory centroids at time t by

$$\text{PHD}(\mu) = \sum_{j=1}^M w^{(j)} \sum_{i \in I_t^{(j)}} f_{ji}(\mu).$$

Here, $w^{(j)}$ is the j th normalised particle weight at time t , $I_t^{(j)}$ enumerates the territories in particle j at time t , and f_{ji} is the i th density in particle j at time t . The densities f_{ji} , $i \in I_t^{(j)}$, each represent the knowledge about a particular territory centroid location within one of the particles (hypotheses). In the context of our model, these densities are either normal densities $N(\mu; m_{ji}, C_{ji})$, with known means m_{ji} and covariances C_{ji} computed by our particle filter, or uniform densities $U(\mu; D_\mu)$. The form of the density, uniform or normal, depends on whether an observation has been 'associated' with a particular territory centroid in the particle. For more details regarding this, see Section 1 of the supplementary material.

Our approach for visualising the territory locations differs slightly from 'standard' PHD, and is as follows. First, we compute the PHD, but with the f_{ji} 's 'at the observation level', meaning that for territories associated at least once, f_{ji} corresponds to the normal density $N(y; m_{ji}, C_{ji} + \Sigma_{\text{obs}})$. For territories never associated, we take $f_{ji}(y) = U(y; D_\mu)$. After computing the PHD in this manner, we furthermore truncate the PHD values from above to the density value $N(0; 0, \Sigma_{\text{obs}})$. This value corresponds to the maximal contribution to the PHD value from a single territory which is known to exist and whose location has been estimated with maximal precision. In Section 3.4 we will visualise the estimated territory locations on the map using this computation. The rationale for this procedure is clearer visualisation of the regions where the filtering algorithm places the territories.

In Section 3 we focus mostly on the estimation of the number of territories, because it is relevant from the point of view of assessing the reproductive capacity of a wolf population. The number of territories is also easy to work with from a model validation perspective, since it is straightforward to compare it numerically to ‘ground truths’ of simulation experiments and to results of other estimation methods. In contrast, the PHD is important for graphical validation, visualisation and interpretation of the filtering result.

3. Results

We first discuss results regarding the observation intensity modelling in Section 3.1. We then move on to the experiments with the data generating model and the developed particle filter, starting with a synthetic scenario in Section 3.2, and then moving on to a semisynthetic scenario in Section 3.3 that resembles the situation with Dataset C but is still based on simulated observations. We conclude with the real data scenario based on Dataset C in Section 3.4.

3.1. Tassu observation arrival intensity estimation

Fig. 3 displays the estimated intensity functions $\lambda_{\text{obs}}^{(r)}$ and $\lambda_{\text{obs}}^{(s)}$ as well as the predicted spatial effect of the intensity model (2). Based on the figure, the temporal intensity function $\lambda_{\text{obs}}^{(r)}$ is seen to capture the rise in observation intensity in the winter time, and the spatial intensity function $\lambda_{\text{obs}}^{(s)}$ especially accounts for the high observation intensities in western Finland.

The overdispersion of the Poisson regression modelling of the observation intensities (Eq. (2)) fit was 1.83. After accounting for the overdispersion the yearly effects were not statistically significant (at 5% level), but monthly effects were clear: The fluctuations with respect to the baseline month of April ranged from a 78% (95% confidence interval [54, 91]) reduction in May to 371% ([206, 592]) increase in November, with smooth transitions in between.

The CORINE landcover variable effects were mostly increasing. When considering a 1% increase in the proportion of each cover class in turn, the estimated increase in intensity was: Cultivated fields 19% ([9, 29]); Closed forests 18% ([8, 28]); Open forests 16% ([7, 27]); Rivers and lakes 16% ([6, 26]); and Other wetlands 19% ([10, 30]). Effect of larger roads was not significant, but a 1% increase in forest roads increased the intensity by 5% ([1, 8]). The smooth component was significant, with a clear reduction effect in the central region and an increase in the west, south-west and south regions.

The explained deviance was 12.6%. For diagnostics we first checked the Pearson residuals aggregated at a month resolution, dropping the spatial dimension (see Figure 4 in the supplementary material). The cell counts showed slightly higher proportion of 0's than the model predictions, otherwise the overall quantiles were reasonably matched. There were no obvious patterns in time, apart from a potential positive trend during 2015. The residual variability was the same during 2011–2016 with only GPS tracking auxiliary territories and during 2017–2019 when both auxiliary territory-types were available. Pearson residuals exceeded +2 during five months (2014-10, 2017-12, 2018-01, 2018-09, 2019-02) with no clear pattern, with three over-estimates (predicted v observed counts: 39 v 14, 257 v 175, 543 v 475), and two under-estimates (6 v 21, 9 v 21).

We then studied the Pearson residuals in space without the time dimension. To visually check troubling areas we aggregated the observed and predicted counts to 10×10 km cells. Observed counts had again slightly larger amount of 0's, and also some higher-than-expected values, the latter mostly from the 2017–2019 period. No obvious spatial structure was visible in the Pearson residual map, with large residuals dotted around the domain (see Figure 4 in the supplementary material). We checked a version where the largest count in the temporal sum per cell was omitted (before aggregation in space). The residual sizes were greatly reduced (max.abs. from 12 to 3). This sensitivity suggests that

the observation counts are more concentrated in time and space than what we can capture with the model. Additionally, spatially contiguous regions of underestimation, particularly on the west coast, were revealed, indicating insufficient information in the spatial components of the model. A further check of before and after 2017 spatial sums revealed a tight cluster of unexpectedly high observation counts in the border region of eastern Kainuu.

3.2. Synthetic scenario and the edge effect

Our first territory estimation experiment is a purely synthetic, simple scenario. The purpose of the experiment is to ensure that the estimation algorithm works correctly. We also investigate explicitly the bias caused by the edge effect, as discussed in Section 2.5.

In this experiment we skip the intensity modelling discussed in Sections 2.3–2.4 and focus on the filtering of simulated datasets. We define the data generating model such that $D_\mu = [0, 100] \times [0, 100]$, $\lambda_{\text{obs}}^{(r)}(u) = 1$ and $\lambda_{\text{obs}}^{(s)}(y) = 1$ for all $u \in [0, 50]$ and $y \in D_y = \mathbb{R}^2$. For the remaining parameters we set $\lambda_{\text{bd}} = 0.0015$, $\lambda_c = 0$, $\lambda_{\text{obs}} = 1$. This configuration corresponds to a simple scenario for our particle filter, since there is no spatial inhomogeneity, and the largest approximation in the filtering arises from the finite domain.

Under these settings, for all combinations of the number of particles $M \in \{128, 256, 512, 1024, 2048\}$ and $\sigma_{\text{obs}} \in \{1, 2, 5, 10, 15\}$ we simulated 450 datasets as follows. First, we simulated the territory locations from the ideal birth and death model, and then conditional on the territories, simulated the observations from the ideal observation model, each time preprocessing the observations with the method discussed in Section 2.5. The initial distribution for the number of territories in the birth and death process was Poisson(20) truncated to the interval [10, 30]. For each sampled initial territory, we set the uniform distribution on D_μ as the initial distribution for the centroid of the territory.

We filtered each simulated dataset assuming that the initial distribution above was known, and computed the deviations $\hat{N}_t - N_t$, for $t = 1, \dots, 50$, where \hat{N}_t is the estimated mean number of territories at time t from the particle filter, and where N_t is the true number of territories for the dataset. We investigated the bias $\mathbb{E}[\hat{N}_t - N_t]$, which would be zero for an ideal Bayes estimator, by computing the empirical mean of the 450 deviations per t , M and σ_{obs} . Fig. 4 summarises the results of this experiment. With a small territory size compared to the size of the domain D_μ , we observe little or no bias in the estimation of the number of territories, given that a sufficient M is used in the filtering algorithm. For greater values of the territory size, the bias is more significant and appears to only slightly diminish with increasing M . This is expected, since with higher territory sizes, Assumption D of Section 2.5 is more likely to be violated, leading to observations outside or near the boundary of D_μ .

3.3. Semisynthetic scenario and feasibility of territory estimation

Next, we consider a more realistic ‘semisynthetic’ scenario with closer resemblance to the situation with Dataset C. In this scenario, the idea was to fix the territory centroids to realistic locations, use plausible parameter values and the intensity functions estimated as discussed in Sections 2.3–2.4. We then simulated data to see how well the particle filter can recover the true number of territories under a more realistic setting.

More specifically, we assumed that D_μ corresponds to the domain of the Tassu data (Fig. 2 top left), and the territory centroid locations were fixed to the centroids (Fig. 2 bottom left) of the territories found by Luke in the wolf population assessment of 2019. We further assumed that each of these 47 territory centroids existed for a period of one year, and that the territory count did not change.

We used the intensity functions in Figs. 3(c) and 3(b) as $\lambda_{\text{obs}}^{(r)}$ and $\lambda_{\text{obs}}^{(s)}$ in the data generating model. To set the value for the territory size parameter $\Sigma_{\text{obs}} = \sigma_{\text{obs}}^2 I$, we examined the shapes and sizes of the

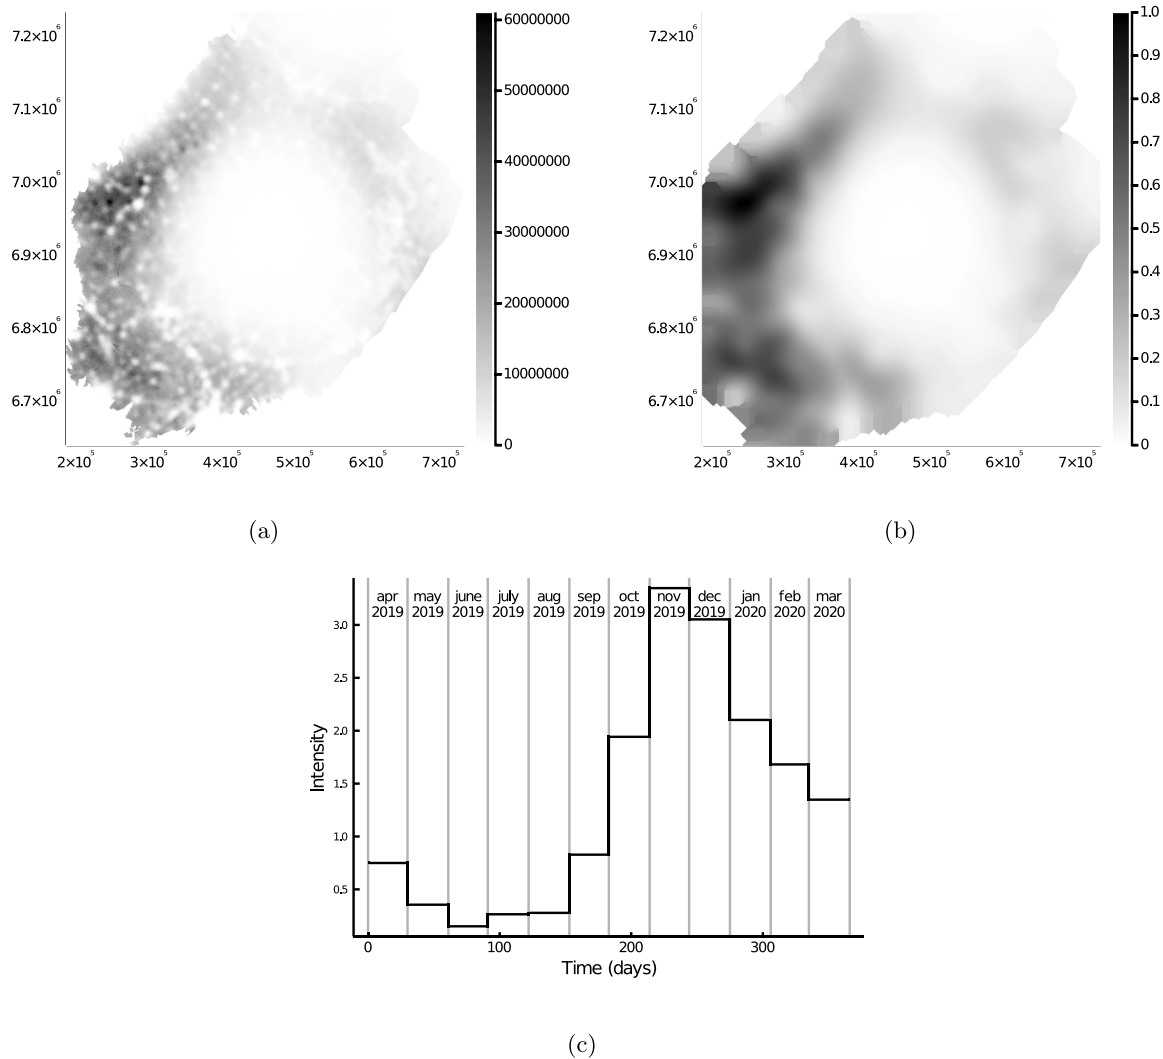


Fig. 3. Outputs of the intensity modelling of Sections 2.3–2.4. Plots (b) and (c) show the estimated spatial intensity function $\lambda_{\text{obs}}^{(s)}$ and temporal intensity function $\lambda_{\text{obs}}^{(t)}$, respectively. Plot (a) shows the predicted spatial effect $\lambda_k^{(s)}$ from the intensity model of Section 2.3 before applying the Gaussian blur as discussed in Section 2.4. The blur standard deviation was set to $\sigma_{\text{obs}} = 13376.67$ m and the domain of $\lambda_{\text{obs}}^{(s)}$ was additionally widened by 27 kilometres at the borders as discussed in Section 2.4.

territory (polygons) in Dataset B discussed in Section 2.1.2. As many of these territories were noncircular in shape (see Fig. 2 (bottom left) for some similar polygons) we estimated σ_{obs} as follows. First, we computed the 95% quantile, $d_{0.95}$, from the empirical distribution of the diameters of the territory polygons in Dataset B. Then, we used Eq. (1) with $\alpha = 0.95$ to compute the σ_{obs} value that corresponds to a radius of $d_{0.95}/2$, yielding the value $\sigma_{\text{obs}} = 13376.67$ m. Here, the ‘diameter of a polygon’ means the maximal length between any two points of the polygon. This procedure guarantees that typical territories ‘fit’ inside the 95% probability region of the territory model.

For the remaining filter parameters, we used the values $\lambda_{\text{obs}} = 1.0$, $\lambda_c = 0.475$ and $\lambda_{\text{bd}} = 0.0015$. The choice of the relative values of λ_{obs} and λ_c here corresponds to a situation where 1% of the total observation intensity is assumed to arise from clutter observations, when the spatial and temporal intensity functions are constant one, and there are 47 territories for a period of one year.

The choice of λ_{bd} corresponds to a mean territory lifetime of $\frac{1}{0.0015} \approx 667$ days, a little less than two years. This choice averages between the fact that in reality some wolf territories are short-lived, but some can exist for years. The chosen value also a priori predicts reasonable changes in the territory count over a period of one year, while maintaining a good agreement between the ideal and approximate birth and death models (see Figure 1 in the supplementary material).

With these settings, we simulated a total of 240 datasets for each particle count $M \in \{2^7, \dots, 2^{12}\}$, and applied the particle filter to each, estimating the mean number of territories at approximately weekly intervals. Each time, the filter was initialised with the initial number of territories following Poisson(47) truncated to [37, 57], with each territory centroid initially following the uniform distribution on D_μ . Fig. 5 shows the deviations computed by subtracting the true number of territories from the estimated mean territory count trajectories for each simulation and all particle counts. In addition, the average deviation and the average absolute deviation are shown. On average, the particle filter appears to recover the true number of territories quite accurately. With increasing numbers of particles, a slight underestimation of the true territory count is revealed. The average absolute deviation further indicates that the discrepancy from the true territory count is typically less than 3 given that a moderate amount of data has been processed.

3.4. Application to the Tassu dataset

Next, we applied the developed particle filter with 16 384 particles to Dataset C. As the initial distribution for the territory centroids, we used the centroids seen in Fig. 2 (bottom left), with Gaussian noise with covariance Σ_{obs} added to each. This way, the prior knowledge from the population assessment of March 2019 can be utilised in the filter.

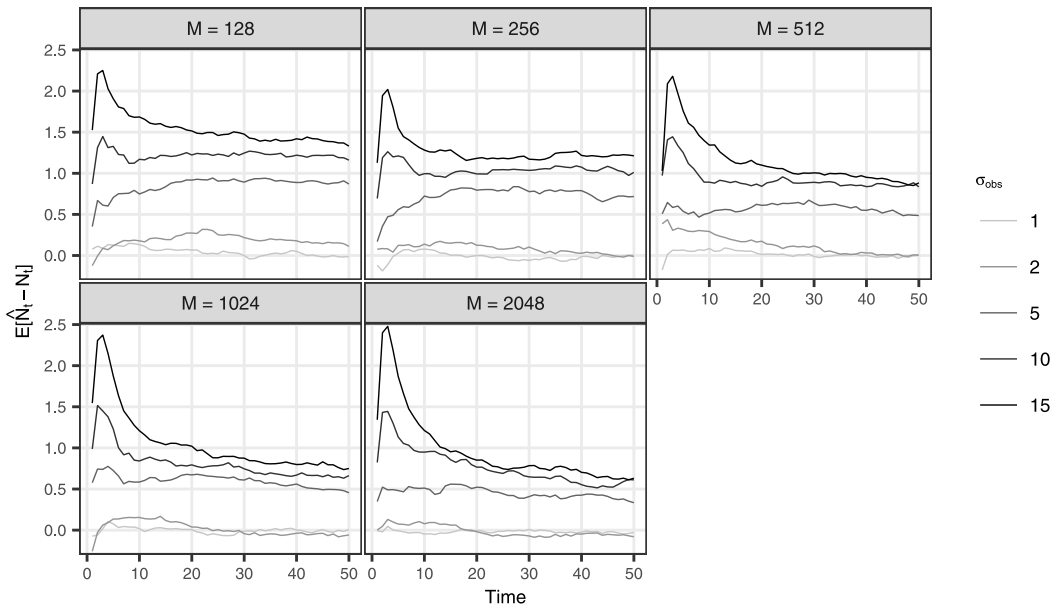


Fig. 4. The estimated bias based on 450 simulations for each M , σ_{obs} and τ in the simulation experiment described in the text.

— Average absolute deviation — Average deviation — Deviation in individual simulations

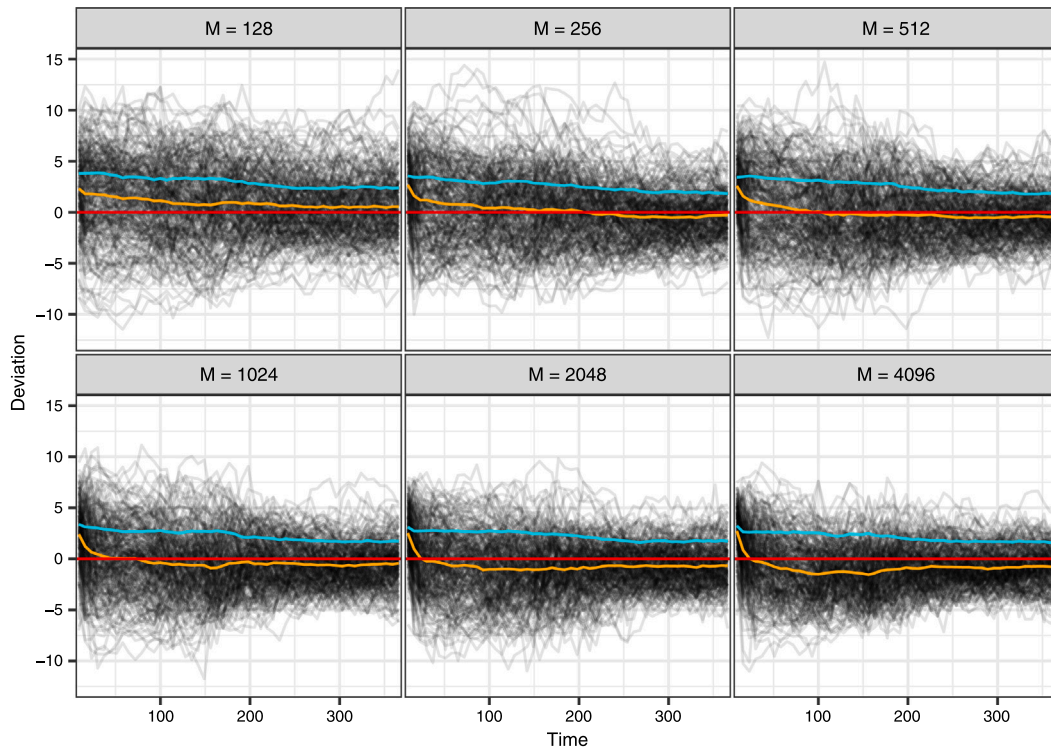


Fig. 5. The true number of territories subtracted from territory count estimates (on black) obtained by applying the particle filter to 240 datasets simulated conditional on the intensity functions in Fig. 3 and territory locations fixed to the centroids of the territory polygons in Fig. 2 (bottom left) for a period of one year. The orange and light blue lines represent the average deviation and the average absolute deviation between the true territory count and estimates, respectively.

We report results for four model variants, as follows. Models 1 and 2 correspond to the same model configuration we used in the semisynthetic experiment, but with Model 1 having $\lambda_c = 0$. Models 3 and 4 correspond to models 1 and 2, respectively, but with another set of intensity functions $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$, obtained by dropping the terms Corine_k , forestroad_k and $\text{smooth}(y_k)$ from the intensity model (2), and then estimating $\lambda_{\text{obs}}^{(\tau)}$ and $\lambda_{\text{obs}}^{(s)}$ as before, as described in Section 2.4. The

resulting spatial intensity function in models 3 and 4 is constant one in D_y . For a plot of $\lambda_{\text{obs}}^{(s)}$ and $\lambda_{\text{obs}}^{(\tau)}$ for models 3 and 4, see Figure 2 in the supplementary material.

Fig. 6 displays the estimated territory locations, and Table 1 shows summary statistics of the filtering distribution for the number of territories based on models 1–4 at the end of March 2020, after the

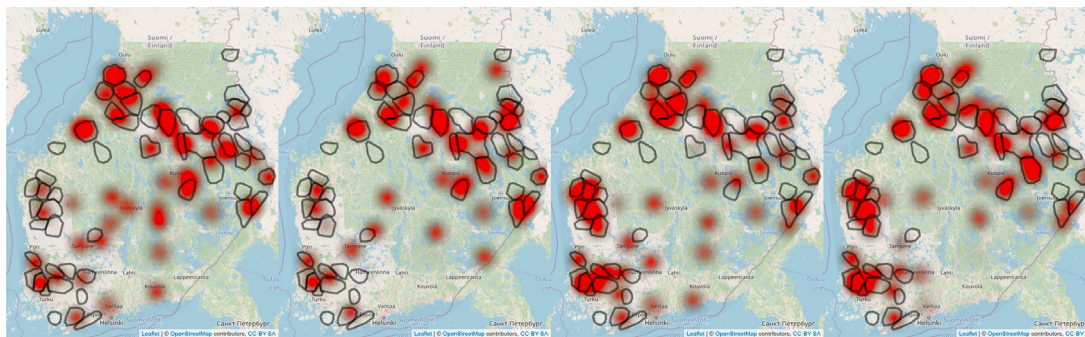


Fig. 6. The estimated territory locations (red) at the end of March 2020 and the territory polygons of the wolf territories in March 2020 (black) from the assessment by Luke. Higher intensities of the red colour depict higher plausibility of a territory location. The intensity of the red colour is computed such that the opaque red colour corresponds to the truncated PHD value discussed in Section 2.6. The individual maps show the results for models 1 through 4 from left to right.

Table 1

Mean, mode, standard deviation and probability intervals of the estimated distribution for the number of wolf territories in March 2020. The estimate 'Luke' corresponds to the estimate by the Natural Resources Institute Finland (Heikkinen et al., 2020). I_α denotes the $\alpha\%$ probability interval.

Estimate	Mean	Mode	St. Dev.	I_{90}	I_{95}	I_{99}
Luke	46.49	47	1.93	[43, 50]	[43, 50]	[41, 51]
Model 1	49.44	49	1.55	[47, 52]	[46, 52]	[46, 53]
Model 2	44.38	44	1.14	[43, 46]	[42, 47]	[42, 48]
Model 3	55.47	55	1.60	[53, 58]	[53, 59]	[52, 60]
Model 4	55.37	53	2.12	[52, 59]	[52, 59]	[51, 60]

full Dataset C has been filtered. The computations underlying these quantities were carried out as explained in Section 2.6.

The overlaid territory polygons in Fig. 6 correspond to the territories found by Luke in the wolf territory assessment of spring 2020. The plots show that each of the models 1–4 seem to find many of the wolf territories found by Luke. However, it appears that models 1 and 3 with $\lambda_c = 0$ place multiple 'extra' territories to central and southeast Finland, that are not found among the official territories. These territories are most likely not real wolf territories, since only a small number of observations have been reported from these areas; see Fig. 2 (top left). The territories most likely arise since the observation intensity is low (see Fig. 3(b)), and perhaps underestimated in these regions, causing the filter to attempt to explain these observations with additional territories. Based on the plots for model 2 and 4, setting $\lambda_c = 0.475$ seems to mitigate this issue a bit, as these observations are no longer interpreted as real observations from wolf territories. We also experimented with a higher value for λ_c , but this resulted in a territory count distribution not well aligned with the count estimates by Luke (see also discussion below). There is also no reason to believe that a substantial proportion of the data would not originate from the wolf territories.

Another observation from Fig. 6 is that models 3 and 4, with the simplified intensity functions, seem to somewhat better capture the cluster of territories in west and southwest Finland, in comparison to models 1 and 2. This difference in the results occurs since the spatial intensity function for models 1 and 2 assumes that in these regions, the reporting intensity of the observations is higher than in other parts of Finland. This in turn results in less territories being needed to explain the observations arriving from these regions, under models 1 and 2.

Based on the territory count distributions summarised in Table 1, the territory counts for models 1 and 2 are best aligned with the estimate of Luke. In comparison, the territory count for models 3 and 4 is somewhat overestimated. Fig. 7 reports the sample standard deviations of the obtained mean territory counts at approximately weekly time points when we repeated the filtering of the Tassu data 195 times with the configuration of model 2 and different particle counts. The observation is that the variability in the estimated mean territory

counts is seen to diminish with increasing numbers of particles, but still remains noticeable even with 16 384 particles. We also experimented with different λ_{bd} , Σ_{obs} and intensity functions, but the phenomenon persisted. In contrast, when a dataset is simulated from the model, the results of this experiment are markedly different, as is seen from the second pane in the figure. Similarly, there is also some variability in the estimated territory locations. Figure 3 in the supplementary material shows the estimated territory locations after 10 independent runs of the filter to the Tassu data.

4. Discussion

We presented a statistical modelling framework for the analysis of citizen science data from territorial animals. At the core of our framework is the data generating model discussed in Section 2.2, that consists of a birth and death model giving rise to the animal territories, and an observation model that links the citizen science observations to the territories. In the developed data generating model, the high variability common to citizen science observation processes is modelled through a temporal and a spatial intensity function, which are assumed fixed and known. The Rao-Blackwellised particle filter described in Section 2.5, estimates the sequence of filtering distributions for the locations and number of territories that describe the knowledge of the animal territories in time as more data is brought in.

We found that the fitted intensity model and the estimated intensity functions were able to capture general trends in the arrival of the citizen science observations as is seen from the estimated intensity functions in Fig. 3. Clearly, the model captures the higher arrival intensity of observations in the winter, which mainly occurs because of snow that leaves wolf tracks visible for potentially long periods of time. The estimated spatial intensity function, on the other hand, captures the high intensity of observations in western Finland and the low intensity of observations in middle Finland compared to other regions.

The intensity model did, however, struggle to explain some characteristics adequately. For instance, sometimes the observations arrived in unexpected bursts, or were highly localised in space, and these features the model was not flexible enough to capture (cf. Figure 4 in the supplementary material). A potential remedy for this might be the addition of random effects e.g. an additional noise component to each spatial pixel, but it would be better to include interpretable rare-event overdispersion components based on the social analysis of the mechanisms for reporting the wolf observations. In fact, the outliers are worth a closer look to gain such insight. The model also struggled with an excess of 0-count cells due to how the conditioning on the auxiliary territories in Dataset B was constructed.

Another improvement for the intensity model could be a zero-inflation component with its own regression structure on, for example, environmental covariates. However, an even better option would be to forgo the aggregation and model the data as a spatio-temporal (marked)

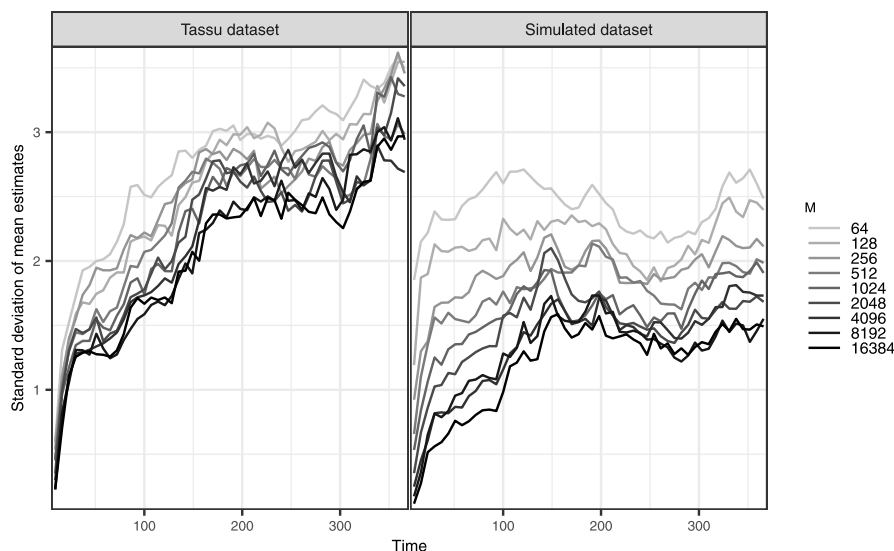


Fig. 7. Sample standard deviations over estimates of the mean territory count obtained by running the particle filter 195 times to Dataset C (Tassu dataset) and a single simulated dataset, for varying numbers of particles M .

point pattern (Sicacha-Parada et al., 2021; Tang et al., 2021). Modelled this way, events such as an observation being exactly on a road, or snow cover of previous days would not be averaged out, yet coarser resolutions for downstream analysis could still be easily computed. Other possible improvements include using the type of observation (sightings, paw prints, game-cameras etc.) in the intensity model and the improvement of the process of deriving the auxiliary wolf territories from the GPS trajectories. Finally, it might be beneficial to investigate the possibility to relax the space–time separability assumption, which would allow for the incorporation of weather data such as snow cover and/or allow for individual time trends for western and eastern Finland, for example. Such complex models could then be estimated from longer period or more frequent observation series, if available.

Our synthetic experiment in Section 3.2 investigated the mean territory count estimates of our particle filter under a simple data generating model, and found that our method works as expected when Assumption D of Section 2.5 holds. With high territory sizes, this assumption is violated, and there is some bias in the estimation of the mean territory count. This edge effect arises due to the finiteness of the domain D_μ . In particular, the violation of Assumption D reduces the accuracy of the approximation (25) (in the supplementary material), which likely causes the bias. In this experiment, territory parameters on a scale of 1%–2% of the width of the rectangular region led to small bias. Noting that the distance between the western and eastern borders of Finland is approximately 500 kilometres, we therefore expect that the bias caused by the edge effect should be small in the realistic experiments discussed in Sections 3.3–3.4.

The semisynthetic experiment of Section 3.3 showed that under realistic conditions resembling the situation with the Tassu dataset, the estimation of the territory count is possible. There was a slight underestimation of the true territory count with increasing numbers of particles, which we think occurs because of the discrepancy between the scenario and the ideal model, or the approximations used. This experiment was a proof of concept that showed that the territory estimation can be feasible even with presence-only data, when the model is correct.

While experimenting with different parameter values for the data generating model, we found that in general the model and especially the count estimation is most sensitive to the values of the territory size parameter Σ_{obs} and the intensity functions, and least sensitive to the choice of the constant and equal birth rate λ_{bd} . This is expected, since large territory sizes increase the probability of an observation being

associated with a territory that is far away, decreasing the relative probability of a new territory emerging. In a similar vein, the intensity functions are directly tied to the amount of territories needed to explain the number of observations arriving. With a fixed dataset, lowering the observation intensity results in more births, since a higher number of territories is then needed to explain the data.

The concluding analysis in Section 3.4 applied the developed data generating model and particle filter to analyse citizen science observations of wolves from April 2019 to March 2020. The results show similar patterns as the counts and locations reported by the Natural Resources Institute Finland (Luke) in the official assessment of March 2020. However, the results do not reach the accuracy of expert judgement. This cannot be expected, because the official assessment also incorporated additional information sources, such as DNA samples, GPS collared wolves and mortality records.

Based on Table 1, the inference of Model 2 incorporating the estimated spatial intensity function and clutter observations resulted in a slightly smaller estimated territory count and smaller variability than the official estimate by Luke. In general, such differences can occur because the estimates of Luke are based on a very different model and assumptions, and also take advantage of additional data. When using the citizen science data only, it is also possible that the particle filter might not find territories which rarely produce observations, leading into underestimation of the territory count. Furthermore, the uncertainty reported by the particle filter can be underestimated, because the results are based on a single particle filter run. If Monte Carlo variability is taken into account, the uncertainty is inflated (see also discussion below).

We noticed that with models using a constant one spatial intensity, the cluster of territories in the west and southwest Finland was captured better than with models that used the estimated intensity functions, as was seen from Fig. 6. This might suggest that the observation intensities in the more complex models are overestimated in these regions at the time of the observations arriving. Indeed, based on the Pearson residuals in Figure 4 of the supplementary material, the intensity model fit could not capture all of the variation in these regions. Despite this, it appears that using the spatially varying intensity improves and is likely a requirement for the accurate estimation of the territory count. This is supported by the territory count distribution in Table 1, which indicates that the territory count estimates of the more complex models were better aligned with the official territory count estimates of Luke.

In general, we found that estimating the territory count and locations of the territories reliably is challenging by only using the citizen science observations from the Tassu database. The challenges faced may be partly explained by unsatisfactory observation intensity modelling, but it appears that the inference algorithm is also struggling with real data. Fig. 7 (and 3 in the supplementary material) show repeated runs of the filter in our concluding analysis, indicating some Monte Carlo variability. We believe that the main challenge for the inference is the presence-only nature of the observations. These observations are quite informative about the births of new territories since a territory needs to exist so that an observation may occur. In contrast, the observations are not very informative about the deaths of the territories, because the information about a death of a territory is indirect and only mediated by the absence of observations arriving from a particular area. This difficulty with the data might explain the Monte Carlo variability in our particle filter, and further suggests that it could still benefit from further specialisation to the territory estimation task. For instance, this specialisation could come in the form of further heuristics that eliminate territories more efficiently, when no observations have been associated for a long time.

All in all, we think that the results obtained suggest that the developed modelling framework might be useful as an additional tool in the annual wolf population assessment, reducing the amount of subjectivity in the estimation process by providing a preliminary statistical interpretation of the citizen-made observations. Integration of the DNA samples, GPS collared wolves and other data with the results of the particle filter would still remain a task for the panel of experts. Our analysis with Datasets A, B and C of Section 2.1 showcased the intended use of the modelling framework in the context of the wolf population assessments. First, the intensity functions required by the data generating model are estimated based on the latest historical data. Then, the particle filter is initialised with the territory count and location distribution available from the latest population assessment. Finally, the yearly observations are analysed in a batch to obtain a model-based view on the status of the wolf population at the time of the next population assessment.

Filtering a year of data from April 2019 to March 2020 allowed both the comparison to the official wolf population assessment of 2020 and the use of prior information from the assessment of 2019. However, this yearly batch estimation is not the only way in which the developed data generating model and particle filter could be used. In fact, one of the motivations for the development of the data generating model and the particle filter was that the developed modelling framework could also be used in an online fashion. This way, smaller batches of new observations could be used to update the posterior distribution of the territory locations and track the population in finer timescales. In the context of the wolf territory estimation, the results from such an online estimation could provide dynamic feedback for the volunteers collecting the data, highlighting that their work is important. Such feedback might also be used to direct the effort of the volunteers to areas with the most uncertainty about the existence of territories.

We envision that our modelling scheme could also be a noteworthy tool for refining the population assessment of other large carnivores. For example, the female Eurasian lynx (*Lynx lynx*) are known to show territorial behaviour with cubs. This could be exploited in the estimation of lynx reproduction. It might also be possible to couple the developed framework with other developments for assessing animal populations, such as the spatial capture–recapture (SCR) model which estimates wolf density based on DNA samples (Bischof et al., 2020). For example, the two approaches might be used sequentially. The SCR model could be used first to estimate the spatial wolf density based on DNA samples, and then the density could provide another source of prior information about potential territory locations for our data generating model. On the other hand, in case that the DNA samples are collected by volunteers, the modelling of the sampling effort in the SCR model could be done by similar techniques as in this work.

Besides the context of territorial animals, our methods might be relevant in target tracking where the modelling of the temporal and spatial variability of the observation process is required. The methods may also be regarded as a form of ‘dynamic clustering’. In different applications, the modular nature of the developed framework can be exploited to carry out the intensity modelling in a way that fits the application.

There are a number of ways how the developed data generating model and the inference algorithm could be improved in future works. The core of the filtering computation consists of evaluating the posterior probabilities for the birth, death and association variables (see Section 1.6 of the supplementary material). The main approximations made in the computation arise from the intractable integrals related to the spatial domains D_μ and D_y and the spatial intensity function $\lambda_{\text{obs}}^{(s)}$. We concentrated on the situation where $\lambda_{\text{obs}}^{(s)}$ may be assumed to be slowly varying by Assumption C, allowing for straightforward approximation of the intractable integrals in the probability computations and the measurement update. Introducing a numerical integration scheme could mitigate the bias from the approximations in the former. It might also be possible to incorporate a ‘no-overlap’ condition to the model that penalises large ‘overlaps’ of the territories.

The developed particle filtering approach assumes fixed parameter values. Even with the limited and noisy citizen science data, it might be possible to estimate some parameters of the data generating model, such as the intensity scaling parameter λ_{obs} and/or parameters related to the birth and death intensity functions λ_b and λ_d . This could result in a better fit of the data and model, and might result in improved location and count estimates. In theory, estimating the model parameters is possible using for example the particle marginal Metropolis–Hastings algorithm (Andrieu et al., 2010) similar to Kokkala and Särkkä (2015). Using these methods would however require the derivation of an unbiased estimate of the marginal likelihood of the observed data under the employed optimal resampling scheme of Fearnhead and Clifford (2003). For the current model and a moderate to large dataset, such estimation procedures are also computationally intensive and would likely require a tailored parallel implementation. The methods might also be difficult to tune in practice.

The data generating model could, in principle, readily incorporate moving territories as well, perhaps using an Ornstein–Uhlenbeck-type movement model as in the work of Johnson et al. (2008). We did not attempt this with the wolf territory estimation, because our main interest was in immobile territories. In addition, we suspect that the additional flexibility in the model allowing for movement would make the inference task substantially more difficult or even infeasible. Furthermore, we do not believe that in our dataset each territory is observed frequently enough to make the inference with an additional movement model practical.

The observation model could also be further refined. We did not separate between the different observation types, but some observation types can be more reliable than others and may be subject to different observation intensity (‘detectability’); consider for instance tracks vs. sightings. We chose to simplify since we do not believe that adequate intensity estimation is possible for the different observation types separately with our dataset. In another context, however, it would in principle be possible to modify the observation model to also accommodate different observation types. The intensity function of the observation model (see Equation (5) in the supplementary material) could be augmented with ‘independent data streams’ for the different observation types, each with their own spatio-temporal and clutter intensities. This could lead to interesting observation models, such as ones that designate higher clutter intensity for more uncertain observations types such as sightings.

Another means of model improvement are more fine-grained birth–death models. In ecological applications, for example, further information such as mating seasons or typical lifespans of the species could be encoded into the birth and death intensity functions λ_b and λ_d .

in the general model described in the supplementary material. However, these kinds of models would likely achieve their full potential when coupled with parameter estimation discussed above. In our wolf territory estimation, we did not investigate time-varying birth/death intensity functions and opted for a model where a territory can emerge at any time. Even though wolves only reproduce in the spring, new wolf territories can form at any time of the year when vagrant wolves pair up and establish new territories. Furthermore, the majority of the citizen science wolf observations are made in the winter time, and it is possible that the first observation from a territory formed in the spring comes later, in the winter.

CRedit authorship contribution statement

Santeri Karppinen: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration. **Tuomas Rajala:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration. **Samu Mäntyniemi:** Conceptualization, Writing – original draft, Resources, Writing – review & editing, Supervision, Project administration. **Ilpo Kojola:** Resources, Writing – original draft, Project administration. **Matti Vihola:** Conceptualization, Methodology, Project administration, Writing – original draft, Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data accessibility statement

The code used to reach the conclusions of the paper is available at the repositories

- <https://github.com/antiphon/tassu-intensity> (intensity analysis)
- <https://github.com/skarppinen/tassu-filtering> (particle filter implementation and experiments of Sections 3.2–3.4).

The repositories (and links therein) also contain simulation data and results related to the experiments. The data on the Finnish wolf population is not publicly available because it contains sensitive information on local wolf behaviour that could be exploited in poaching or disturbing wolves.

Acknowledgements

We thank Gabriel Kasmí, who was involved with early experiments related to this research. SK and MV were supported by the Academy of Finland research grant 315619. The authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ecolmodel.2022.110101>.

References

- Andrieu, C., Doucet, A., Holenstein, R., 2010. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 72 (3), 269–342.
- Bischof, R., Milleret, C., Dupont, P., Chipperfield, J., Tourani, M., Ordiz, A., de Valpine, P., Turek, D., Royle, J.A., Gimenez, O., Flagstad, Ø., Åkesson, M., Svensson, L., Brøseth, H., Kindberg, J., 2020. Estimating and forecasting spatial population dynamics of apex predators using transnational genetic monitoring. *Proc. Natl. Acad. Sci.* (ISSN: 0027-8424) 117 (48), 30531–30538. <http://dx.doi.org/10.1073/pnas.2011383117>, URL <https://www.pnas.org/content/117/48/30531>.
- Bragina, E.V., Ives, A., Pidgeon, A., Kuemmerle, T., Baskin, L., Gubar, Y., Piquer-Rodríguez, M., Keuler, N., Petrosyan, V., Radeloff, V., 2015. Rapid declines of large mammal populations after the collapse of the Soviet union. *Conserv. Biol.* 29 (3), 844–853.
- Conrad, C., Hichley, K., 2011. A review of citizen science and community-based environmental monitoring: issues and opportunities. *Environ. Monit. Assess.* 176, 273–291.
- Cretois, B., Linnell, J.D., Grainger, M., Nilsen, E.B., Rød, J.K., 2020. Hunters as citizen scientists: Contributions to biodiversity monitoring in Europe. *Global Ecol. Conserv.* 23, e01077.
- Doucet, A., Godsill, S., Andrieu, C., 2000. On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.* 10 (3), 197–208.
- Fearnhead, P., Clifford, P., 2003. On-line inference for hidden Markov models via particle filters. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 65 (4), 887–899.
- Finnish Environment Institute SYKE, 2018. CORINE Land Cover 2018. Data is downloaded from the data download service of SYKE on 18.3.2019 under the license CC 4.0 BY.
- Finnish Transport Infrastructure Agency, 2021. Digiroad. Data is downloaded from the download and viewing service of the Finnish Transport Infrastructure Agency on 19.01.2021 under the license CC 4.0 BY.
- Goodman, I.R., Mahler, R.P.S., Nguyen, H.T., 1997. *Mathematics of Data Fusion. Series B: Mathematical and Statistical Methods*, vol. 39, Kluwer Academic Publishers, AA Dordrecht, The Netherlands, ISBN: 0-7923-4674-2.
- Gregory, R.D., van Strien, A., Vorisek, P., Gmelig Meyling, A.W., Noble, D., Roy, D., 2005. Developing indicators for European birds. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 360, 269–288.
- Heikkinen, S., Kojola, I., Mäntyniemi, S., Holmala, K., Härkölä, A., 2020. Susikanta suomessa maaliskuussa 2020. Luonnonvarakeskus 37, URL <http://urn.fi/URN:ISBN:978-952-326-979-8>.
- Illian, J., Penttinen, A., Stoyan, H., Stoyan, D., 2008. *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley & Sons.
- Isaac, N., 2014. Statistics for citizen science: extracting signals from noisy ecological data. *Methods Ecol. Evol.* 5, 1052–1060.
- Johnson, D.S., London, J.M., Lea, M.-A., Durban, J.W., 2008. Continuous-time correlated random walk model for animal telemetry data. *Ecology* 89 (5), 1208–1215.
- Kojola, I., Hallikainen, V., Mikkola, K., Gurarie, E., Heikkinen, S., Kaartinen, S., Nikula, A., Nivala, V., 2016. Wolf visitations close to human residences in Finland: the role of age, residence density, and time of day. *Biol. Conserv.* 198, 9–14.
- Kojola, I., Heikkinen, S., Holmala, K., 2018. Balancing costs and confidence: volunteer-provided point observations, GPS telemetry and the genetic monitoring of Finland's wolves. *Mammal Res.* 63, 415–423. <http://dx.doi.org/10.1007/s13364-018-0371-3>.
- Kokkala, J., Särkkä, S., 2015. Combining particle MCMC with Rao-Blackwellized Monte Carlo data association for parameter estimation in multiple target tracking. *Digital Signal Process.* 47, 84–95.
- Lawrence, A., 2006. No personal motive? Volunteers, biodiversity, and false dichotomies of participation. *Ethics Place Environ.* 9, 279–298.
- Linnell, J.D., Kaczensky, P., Wotschikowsky, U., Lescureux, N., Boitani, L., 2015. Framing the relationship between people and nature in the context of European conservation. *Conserv. Biol.* 29 (4), 978–985.
- Pellikka, J., Hiedanpää, J., 2017. Looking for a common ground: useful knowledge and adaptation in wolf politics in southwestern Finland. *Wildlife Biol.* 2017 (4).
- Renner, I.W., Elith, J., Baddeley, A., Fithian, W., Hastie, T., Phillips, S.J., Popovic, G., Warton, D.I., 2015. Point process models for presence-only analysis. In: O'Hara, R.B. (Ed.), *Methods Ecol. Evol.* (ISSN: 2041210X) 6 (4), 366–379. <http://dx.doi.org/10.1111/2041-210X.12352>, URL <http://doi.wiley.com/10.1111/2041-210X.12352>.
- Royle, J., Kéry, M., 2007. A Bayesian state-space formulation of dynamic occupancy models. *Ecology* 88 (7), 1813–1823.
- Särkkä, S., Vehtari, A., Lampinen, J., 2007. Rao-Blackwellized particle filter for multiple target tracking. *Inf. Fusion* 8 (1), 2–15.
- Sicacha-Parada, J., Steinsland, I., Cretois, B., Borgelt, J., 2021. Accounting for spatial varying sampling effort due to accessibility in citizen science data: A case study of moose in Norway. *Spat. Statist.* (ISSN: 22116753) 42, 100446. <http://dx.doi.org/10.1016/j.spasta.2020.100446>, URL <https://linkinghub.elsevier.com/retrieve/pii/S2211675320300403>.
- Silvertown, J., 2009. A new dawn for citizen science. *Trends Ecol. Evol.* 24, 467–471.
- Tang, B., Clark, J.S., Gelfand, A.E., 2021. Modeling spatially biased citizen science effort through the eBird database. *Environ. Ecol. Stat.* <http://dx.doi.org/10.1007/s10651-021-00508-1>, (ISSN: 1352-8505, 1573-3009). URL <https://link.springer.com/10.1007/s10651-021-00508-1>.

- Ver Hoef, J.M., Johnson, D., Angliss, R., Higham, M., 2021. Species density models from opportunistic citizen science data. *Methods Ecol. Evol.* 2041–210X.13679. <http://dx.doi.org/10.1111/2041-210X.13679>, (ISSN: 2041-210X, 2041-210X). URL <https://onlinelibrary.wiley.com/doi/10.1111/2041-210X.13679>.
- Vihola, M., 2007. Rao-Blackwellised particle filtering in random set multitarget tracking. *IEEE Trans. Aerosp. Electron. Syst.* 43 (2), 689–705.
- Wood, S., 2017. *Generalized Additive Models: An Introduction with R*, second ed. Chapman and Hall/CRC.
- Zhao, Q., Royle, A., Boomer, G.S., 2017. Spatially explicit dynamic N-mixture models. *Popul. Ecol.* 59 (4), 293–300.