# A real-time data assimilative forecasting system for animal tracking ⋆

Marine Randon[a], Michael Dowd[b], Ruth Joy[a,c]

[a]*Department of Statistics and Actuarial Science, Simon Fraser University, 8888 University Drive, Burnaby, BC, V5A 1S6, Canada*

[b]*Department of Mathematics and Statistics, Dalhousie University, 6316 Coburg Road, PO Box 15000 Halifax, Nova Scotia, B3H 4R2, Canada*

[c]*School of Environmental Science, Simon Fraser University, 8888 University Drive, Burnaby, BC, V5A 1S6, Canada*



⋆Open Research: R code is available on figshare with DOI: `https://doi.org/10.6084/m9.figshare.17046026.v2`. The data associated with the code were simulated inside the R script.

*Email address:* `Corresponding author: marine_randon@sfu.ca` (Marine Randon)

*Preprint submitted to Ecology* *February 18, 2022*

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1002/ecy.3718

## Abstract

Monitoring technologies now provide real-time animal location information which opens the possibility of forecasting systems to fuse these data with movement models to predict future trajectories. State space modelling approaches are well established for retrospective location estimation and behavioural inference through state and parameter estimation. Here, we use a state space model within a comprehensive data assimilative framework for probabilistic animal movement forecasting. Real-time location information is combined with stochastic movement model predictions to provide forecasts of future animal locations and trajectories, as well as estimation of key behavioural parameters. Implementation uses ensemble-based sequential Monte Carlo methods (a particle filter). We first apply the framework to an idealized case using a non-dimensional animal movement model based on a continuous-time random walk process. A set of numerical forecasting experiments demonstrate the workflow and key features, such as the online estimation of behavioural parameters using state augmentation, the use of potential functions for habitat preference, as well as the role of observation error and sampling frequency on forecast skill. For a realistic demonstration, we adapt the framework to short-term forecasting of the endangered Southern Resident Killer Whale (SRKW) in the Salish Sea using visual sighting information wherein the potential function reflects historical habitat utilization of SRKW. We successfully estimate whale locations up to 2.5 hours in advance with a moderate prediction error ($< 5$ km), providing reasonable lead-in time to mitigate vessel-whale interactions. It is argued that this forecasting framework can be used to synthesize diverse data types, improve animal movement models and behavioural understanding, and has the potential to become an important new direction for movement ecology.


*Keywords:* animal movement, continuous time correlated random walk, data assimilation, ecological forecasting, particle filter, potential function, Southern Resident Killer Whale, state augmentation, state space models, trajectory prediction, whale collision avoidance

## 1. Introduction

Understanding ecological processes relies on our ability to make predictions and confront them with observations to refine hypotheses and theories. This is also the essence of the emerging field of ecological forecasting, which has arisen due to the many new data types becoming available. Ecological forecasting differs from standard statistical projection methods by its iterative nature and its reliance on dynamic models. The central idea is to generate forecasts of the future ecological state using dynamic models of ecological processes, compare the predictions to observations, and then refine hypotheses and models to improve predictive skill (Dietze et al., 2018). The focus on forecasting shifts the emphasis to the iterative refinement of ecological dynamic models, as well as to identifying key observational needs - thereby driving understanding and advancement of the ecological sciences.

Dynamical ecological forecasting is distinct from forecasting via statistical prediction. The former is based on using mechanistic or process-based models to project the ecological system forward in time, while the latter is based on using the established statistical models together with forecasts of their key environmental predictors. For instance, correlative species distributions project future animal distributions according to forecasts of their environmental drivers (e.g., Barlow and Torres (2021), Breece et al. (2021)). This type of model is increasingly used for managing human-wildlife conflicts in real-time with the purpose of limiting the probability of encounter. However, statistical predictions have a limited ability to incorporate ecological process and dynamics (Yates et al., 2018), and they rely upon existing conditions that may or may not hold in the future. Hence, transitioning from empirical to dynamical models may lead to better ecological forecasting (Payne et al., 2017). Dynamics-based ecological forecasting is reliant on accurate initial conditions and a useful mathematical description of processes that can project the ecological system state into the future. Contrary to correlation-based forecasts, dynamics-based forecasts are process-based and can actively learn from real-time observations when embedded in a data assimilative framework (Kitagawa, 1998). They can thus adapt to changing environmental conditions and structural ecological changes. Dynamics-based forecasting has not been extensively

applied in ecology (Payne et al., 2017, Dowd et al., 2014), especially in the field of animal movement, and this is the purpose of this paper.

New technologies for animal tracking (e.g., satellite tags, acoustic and electromagnetic detection) and communication networks (e.g., reporting apps) yield real-time information that has improved our understanding of movement ecology (Wall et al., 2014, Williams et al., 2020). Retrospective analyses of such tracking data has lead to the development of sophisticated fit-for-purpose statistical approaches, usually based on state space models (SSMs) (Hooten et al., 2017, Patterson et al., 2017). SSMs combine a statistical model of observations (i.e., a measurement model) with a model of the dynamic process (i.e., a movement model). The central goal is to estimate the system state (i.e., the unobserved animal locations) (Auger-Méthé et al., 2020), but they also can be used to determine system parameters linked to behavioural dynamics (Kitagawa, 1998, Dowd and Joy, 2011). SSMs can ingest and synthesize various sources of location information (e.g., tags, telemetry, visual or acoustic detections) (Patterson et al., 2017) and make use of increasingly sophisticated movement models (McClintock et al., 2017, Michelot et al., 2021). Consequently, SSMs are well adapted for ecological forecasting (Dowd et al., 2014, Dietze et al., 2018). Forecasting shifts the emphasis to predictive skill, which is distinct from retrospective model fitting that focuses on location in-filling and estimation of behavioural parameters and states. Forecasting strongly depends on having good movement models, which in turn requires understanding of ecological processes. A prediction system enhances this by putting emphasis on refining model structure, estimating their parameters, incorporating environmental features, and allowing them to adaptively learn from tracking data (Payne et al., 2017, Dietze et al., 2018).

Practical goals for studying the real-time location of animals and forecasting their future trajectories and locations include management and conservation objectives, especially for at-risk species. Forecasting systems may facilitate proactive management and increase the efficiency of mitigation measures by limiting the probability of human-wildlife conflicts in time and space (e.g., animal-vehicle or animal-vessel collisions, animal incursions into sensitive areas; Wall et al. (2014), Gervaise et al. (2021)). This study proposes a general framework for an animal forecasting

system that provides real-time fusion of location data with a movement model to yield probabilistic forecasts of animal location and key behavioural parameters. This iterative forecasting system uses state space models and ensemble methods. It follows the data assimilation cycle, alternating between a prediction step (i.e., forecast) using a process model, followed by an observation update (i.e., nowcast) using real-time observations (Dowd et al., 2014). An idealized non-dimensional example highlights the major features of the forecasting system including: use of a stochastic movement model, real-time data assimilation, state augmentation to estimate behavioural parameters, potential functions to incorporate habitat preference, and the evaluation of forecast skill. A realistic demonstration is then undertaken for short-term prediction for the endangered population of Southern Resident Killer Whales (SRKW), *Orcinus orca*, in the Salish Sea off southern British Columbia and northern Washington state, with the aim of mitigating disturbance from commercial shipping traffic (McWhinnie et al., 2021).

## 2. Methods

### *2.1. General framework*

#### *2.1.1. State Space Model*

SSMs are a general framework that couple a process model to a measurement model:

$$x_t = d(x_{t-1}, \theta_t, Z_t) + w_t \quad (1)$$

$$y_t = h(x_t) + \epsilon_t \quad (2)$$

where $x_t$ is the state of the system (e.g., the animal location), and $y_t$ represents the observations (e.g., error-prone location measurements) at time $t$. The process equation (1) represents the dynamics (e.g., an animal movement model) where $x_t$ depends on its value at the previous time, $x_{t-1}$, a set of parameters, $\theta_t$, and a set of covariates, $Z_t$. Note that parameters and covariates may or may not be time dependent. The stochastic error or forcing, $w_t$, is assumed additive, but could be multiplicative. The functional form of the model is embodied in the $d(\cdot)$ operator, and reflects the

time-dependent ecological dynamics. The measurement equation (2) relates the observations $y_t$ to the state $x_t$ through the measurement operator $h(\cdot)$. Direct observations of the state implies that $h(\cdot)$ is the identity operator. The observation error term is given by $\epsilon_t$. The goal of the basic state space model is online (real-time) estimate of the state, $x_t$, using observations, $y_t$, for $t = 1, \ldots, T$, with all other quantities known or specified. Parameters of the system can also be estimated online by the technique of state augmentation (see Section 2.1.4).

### *2.1.2. Data assimilation*

The target for our prediction system is to provide online estimates of the current location of the animal (a nowcast), and short term predictions of future locations (a forecast)[1]. Sequential state estimation follows the data assimilation (DA) cycle (Dowd et al., 2014). Figure 1 shows a schematic of this procedure. It describes the transition of the system from one time to the next (with the understanding that this is part of a continuously operating real-time sequential estimation). We assume that the probabilistic location nowcast is available at time $t-1$, and given by $[x_{t-1}|y_{1:t-1}]$ where $[\cdot]$ designates a probability density function and $y_{1:t-1}$ are the location observations from time 1 to time $t-1$ inclusive. A one-step ahead forecast is undertaken to transition the system from time $t-1$ to time $t$, yielding the animal location forecast $[x_t|y_{1:t-1}]$. This is done by applying the movement model (1) using the nowcast as the initial condition. Note that $n$-step ahead forecasts can also be produced to yield future predictions of animal locations on longer time horizons (Figure 1). Next, location observations, $y_t$, may become available at time $t$. If so, the assimilation step statistically blends location forecasts with the new observations yielding the nowcast at time $t$, or $[x_t|y_{1:t}]$. This probabilistic observational update is based on Bayesian principles treating the forecast as a prior and using the likelihood of the new observation. The procedure can continue indefinitely through time, cycling between movement forecasts and assimilation steps. It is initialized at time 0 with a location density $[x_0]$. In practice, prediction

[1] In this section, for clarity, we present state estimation but note that time-varying parameters may also be simultaneously estimated using an augmented state following the same procedure (see Section 2.1.4). As well, we suppress the explicit dependence on covariates and static parameters for notational simplicity.

and observation updates are carried out in an ensemble framework wherein samples (or particles) are used to represent the target nowcast and forecast densities. Specifically, forecasting (one-step or $n$-steps ahead) is based on ensemble prediction using (1), and assimilation is carried out with a particle filter (see Section 2.1.3).

### *2.1.3. Particle filter*

The particle filter is a sampling-based solution algorithm for sequential data assimilation. The data assimilation cycle is divided in two steps: (i) forecasting and (ii) observation update. Suppose we are at time $t-1$ and have a sample from the nowcast distribution, $[x_{t-1}|y_{1:t-1}]$. We designate this sample of size $N$ as $\{x^{(i)}_{t-1|t-1}\}_{i=1}^{N}$ where $i$ identifies a sample member, or particle. The standard particle filter algorithm (sequential importance resampling; Gordon et al. (1993)) proceeds as follows:

1. *Prediction*: Apply the movement model (1) for one step ahead prediction to each ensemble member of the nowcast $\{x^{(i)}_{t-1|t-1}\}_{i=1}^{N}$:

$$x^{(i)}_{t|t-1} = d(x^{(i)}_{t-1|t-1}, \theta_t, Z_t) + w^{(i)}_t, \qquad \text{for } i = 1, \ldots, N \tag{3}$$

   with $w^{(i)}_t$ an independent realization of the system noise. This yields the forecast ensemble $\{x^{(i)}_{t|t-1}\}_{i=1}^{N}$ which is a draw from $[x_t|y_{1:t-1}]$.

2. *Observation Update*: Carry out weighted resampling of the forecast ensemble $\{x^{(i)}_{t|t-1}\}_{i=1}^{N}$ using the observation $y_t$ at time $t$. The weights are based on the likelihood $[y_t|x_t]$ determined from (2) and computed as

$$W^{(i)}_t \propto p(y_t|x^{(i)}_{t|t-1}), \qquad \text{for } i = 1, \ldots, N. \tag{4}$$

   where $W^{(i)}_t$ is the weight given the $i$th particle. The weights are normalized so they sum to one. A weighted bootstrap (resampling with replacement) of $\{x^{(i)}_{t|t-1}\}_{i=1}^{N}$ is carried out to yield the nowcast ensemble $\{x^{(i)}_{t|t}\}_{i=1}^{N}$ at time $t$, which is a draw from $[x_t|y_{1:t}]$.

This single-state recursive transition of the system from time $t-1$ to $t$ is carried on sequentially through time by predicting forward with the model and assimilating new observations. An initial condition for the state, $[x_0]$, must be specified as an initial ensemble $\{x_0^{(i)}\}_{i=1}^{N}$.

### *2.1.4. State augmentation*

In the context of animal movement, estimation of key parameters may be important for representing underlying ecological processes and improving prediction. State augmentation appends such parameters to the original state vector so that the augmented state is $\tilde{x}_t = (x_t \ \theta_t)^T$, and therefore includes both the geographical location as well as parameters of interest. This allows for simultaneous estimates of states and time-varying parameters. Specifically, following Kitagawa (1998), the original process model (1) is transformed to the augmented one

$$\begin{pmatrix} x_t \\ \theta_t \end{pmatrix} = \begin{pmatrix} d(x_{t-1}, \theta_t, Z_t) \\ \theta_{t-1} \end{pmatrix} + \begin{pmatrix} w_t \\ \nu_t \end{pmatrix} \tag{5}$$

where the parameter $\theta_t$ varies as a random walk with a disturbance term $\nu_t$. The augmented measurement equation is a trivial alteration of (2) to reflect the fact that the state, but not the parameters, are observable. Most importantly, since the augmented state space model is the same general form as the usual state model (1)-(2), it can be estimated using standard sequential Monte Carlo methods, such as the particle filter. More sophisticated algorithms based on state augmentation are available if static parameter estimation is the goal (Ionides et al., 2011).

## *2.2. Idealized case*

The aim of the idealized simulation experiments are two-fold: (i) to demonstrate the general workflow and implementation of real-time data assimilation in the context of animal movement forecasting, and (ii) to provide a concrete illustration of key features such as state augmentation (for estimating behavioural parameters) and potential functions (for incorporating the environment through habitat preference), as well as forecast skill assessment.

### 2.2.1. Movement model

We choose a specific animal movement model corresponding to the general process model (1). This takes the form of a continuous-time correlated random walk, a reasonably sophisticated and well-used stochastic model in the random walk family (Johnson et al., 2008). This movement model also forms the basis for the application of Section 2.3, so additionally serves to introduce its major features. At its core is an Orstein-Ulhenbeck process for animal velocity (Russell et al., 2018):

$$dV_t = \frac{1}{\tau}(\mu_t - V_t)dt + \sigma dW_t \tag{6}$$

where $V_t$ is the velocity at time $t$, $\tau$ is a memory time scale parameter, $\mu_t$ is a time-dependent drift term, and $\sigma$ the scale factor for the Wiener process $W_t$. The application is two-dimensional and defined in the horizontal plane. For implementation, we numerically integrate (6) using the Euler-Maruyama approximation method (Kloeden and Platen, 2013) which yields the stochastic difference equation,

$$v_t = (1 - \phi_t)\mu_t + \phi_t v_{t-\Delta} + w_t \tag{7}$$

where $\Delta$ is the time step, $\phi_t = 1 - \Delta/\tau$ is a time-varying velocity persistence, and $w_t \sim \mathcal{N}(0, \sigma_w^2 I)$ is bivariate white noise forcing, with $\sigma_w = \sigma\Delta$ and $I$ the identity matrix. It is straightforward to use this discrete-time model to generate realizations $v_t$ from the probabilistic velocity process $V_t$ for any arbitrarily small $\Delta$. As a general rule, with shorter time steps, the approximation (7) of (6) is more accurate and the trajectory smoother and more continuous. To obtain the horizontal animal position $x_t$, we integrate the velocity $v_t$ by summing its increments,

$$x_t = x_{t-n\Delta} + \Delta \sum_{i=1}^{n} v_{t-i\Delta} \tag{8}$$

where $x_{t-n\Delta}$ is the animal position at a time $n\Delta$ before the present time $t$. The time history of $v_t$ is obtained from (7).

The time-dependent drift term $\mu_t$ in (7) is an externally imposed velocity perturbation, or forcing, due to exogenous environmental conditions. The drift term, $\mu_t$, is determined using a potential field approach (Brillinger et al., 2012). Its value depends on the current animal location such that the local gradient of the potential function influences the magnitude and directional tendency of the animal's movement (Russell et al., 2018). This potential function is a mixture of a Gaussian and a parabola (Figure 2a and e). It is isotropic, resembles a Gaussian near the origin, and decreases as a parabola far from the origin. The drift term, $\mu_t$, is proportional to its local gradient at the animal's current locations and thus movement is steered towards higher values of the potential function. Here, we interpret the potential function as a habitat preference, or resource selection.

The term $\phi_t$ in (7) is a behavioural parameter that encapsulates the tendency of an animal to move in the same direction, in other words, its auto-correlation properties (Russell et al., 2018). Note that $\phi_t$ in (7) also acts as a weighting factor, and so ranges from 0 to 1. As $\phi_t \to 0$ the drift term $\mu_t$ dominates with animal behaviour resembling foraging (tortuous paths); when $\phi_t \to 1$, the velocity process tends to a first order auto-regressive process and the behaviour resembles transiting (directed paths). This behavioural parameter has two notable features: (i) $\phi_t$ is a time-dependent parameter thereby allowing for continuum of behavioural states ranging from, say, foraging to transiting, and (ii) this parameter is estimated, along with the animal position, using a state augmented particle filter (Sections 2.1.4, 2.1.3). Hence, online parameter estimation uses information contained in the recent history of observed location data.

The idealized movement model is further transformed to be scale independent. To do this, the model (7)-(8) is rendered non-dimensional using the following quantities. The characteristic length scale is assumed to be $2L$, or two standard deviations of the Gaussian that (partly) defines the potential function. The velocity scale used is $\sigma_w$, or the standard deviation of the velocity forcing. These together imply a characteristic time scale of $2\sigma_w/L$ which can roughly be interpreted as the time it takes an animal to transverse the Gaussian part of the potential field. Using non-dimensional quantities makes the application scale-free, and thus applicable to organisms from viruses to whales.

### 2.2.2. *Numerical experiments*

In this section, idealized scenarios are presented to illustrate the implementation, features and properties of the forecasting system. The idea is to vary the accuracy (observation error) and availability (sampling rate) of the animal location data, as well as to illustrate the use of habitat preference through potential functions. These simulation experiments are based on realizations of the movement model, described in Section 2.2.1, which provides the known true positions $x_t$. The true time-varying velocity persistence behavioural parameter, $\phi_t$, followed a sinusoid (see Figure 2b, d, f, h). Synthetic observations $y_t$ are created by adding an observation error, $\epsilon_t \sim N(0, \sigma_\epsilon I)$, to the estimated true positions $x_t$ following (2) with $h(\cdot)$ being the identity operator. Simulated tracks are selected that show a clear overlap between the foraging behaviour (i.e. small velocity persistence) and the highest values of the potential field to mimic a habitat preference corresponding to a foraging or resting area.

Four simulation experiments, or scenarios, are considered: 1) low observation error, high sampling rate, drift; 2) low observation error, high sampling rate, no drift; 3) high observation error, low sampling rate, drift; and 4) high observation error, low sampling rate, no drift. Two realizations are used to generate the true positions: one for the drift case (i.e., using a potential function), and one for the no drift case. We run nowcast and forecast scenarios for each experiment and estimate the location state, along with time-varying parameter $\phi_t$. Our particle filter algorithm uses ensembles to yield probabilistic estimate of locations and time-varying parameters. For our idealized case, we use $N = 100$ particles, or ensemble members. A prototype R code (R version 4.0.3) of the forecasting system is provided in Supporting Information.

The central metric for nowcast and forecast skill is the Root Mean Square Error (RMSE). It measures the discrepancy between the observed and predicted animal position as $e_t = y_t - \hat{x}_t$, where $y_t$ represents the observed location at time $t$ and $\hat{x}_t$ is predicted position taken to be the median of the nowcast or forecast ensembles for time $t$. Then $RMSE = \sqrt{1/q \sum_{i=1}^{q} \|e_i\|^2}$, with $\|\cdot\|$ the vector norm and $q$ the number of observations. For each of the scenarios, we computed the RMSE for the nowcast location estimate, and for different forecast horizons, or $n$-step ahead

forecasts, with $n = 1 \ldots 30$ time units. Furthermore, to understand if forecast skill depends on the behavioural mode of the animal (i.e., foraging implied as $\phi_t \to 0$, or transiting as $\phi_t \to 1$), we assess the relationship between the behavioural parameter $\phi_t$, estimated via state augmentation, and the $n$-step ahead forecast RMSE using Simulation 1 (i.e., low observation error, high sampling rate drift).

### *2.3. Application: Southern Resident Killer Whales*

We illustrate and adapt the general framework to the specific problem of nowcasting and forecasting endangered SRKW pod trajectories in the Salish Sea. This population (clan) is composed of three stable matriarchal social groups termed J, K and L pod, each having a tendency to move as a coherent group. Hence, our application is designed to track pods, not individuals. We focus on J pod, the most observed pod in the Salish Sea during summer (Olson et al., 2018). The idealized movement model and data assimilation systems outlined in the previous section provides the basis for the SRKW application, and we outline the specific implementation details below.

#### *2.3.1. Movement model*

The whale movement model follows the model formulation described in 2.2.1, but is dimensional. For our particular application, we fix the model time step $\Delta = 5$ minutes to provide for accurate numerical implementation. This is also taken to be the time scale for the DA cycle meaning model output can match the times of available SRKW observations. The movement model is implemented for the spatial domain associated with the Salish Sea (a portion of which is shown in Figure 4). The main modification is to incorporate SRKW avoidance of land and shallow waters $< 5$ m. This is done in the movement model forecast step wherein ensemble members that fall on land or shallow water are removed.

The two main control parameters in the movement model, the persistence, $\phi_t$, and the drift term, $\mu_t$, allow whale trajectories to mimic different movement behaviours (e.g., transiting, resting, foraging, attraction to preferred habitat). To specify the value of $\phi_t$, we make it part of the online estimation procedure for whale location by using a state augmented particle filter (Sections 2.1.3,

2.1.4). The drift term, $\mu_t$, on the other hand, is designed to take account of SRKW historical habitat usage in the Salish Sea. Watson et al. (2019) develops the framework of a spatio-temporal point process model to create time indexed spatial whale intensity fields (maps) for each of pods J, K, and L. In this study, these whale intensity maps are created at monthly resolution and define the potential functions for each month of sightings data, e.g., $U_m(x)$, $m = 1, \ldots, M$. The drift term is then defined for month $m$ as the gradient of $U_m$ at the current whale location, i.e. $\mu_t = \nabla U_m$. This gradient acts as a force that determines the drift direction and magnitude, attracting trajectories towards areas of highest historical whale intensity. The drift term thus adds realism to simulated whale trajectories in the absence of direct location observations

### *2.3.2. Observations*

Observations of SRKW locations for this application are based on visual sighting data from the OrcaMaster database (Olson et al., 2018). These are available at irregular time intervals when SRKWs are present in the Salish and take the form of real-time opportunistic SRKW locations (to pod level, from a reporting app) during daylight hours. Like the idealized example, SRKW observations $y_t$ are assumed to have an additive error $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$. For simplicity, we fix $\sigma_\epsilon = 1$ km for SRKW observations. In reality, the error may be considerably more complex due to, for example, mismatches between the sighting and reporting times leading to location and timing errors, weather conditions, and observer effects. Another important issue is that SRKW pods may split or disperse meaning our visual detections may not necessarily reflect the core distribution of the pod. For our demonstration of real-time data assimilation and probabilistic prediction, we selected a single 5.5 hour track of J pod from 18 August 2016. The track consists of 14 observations between 10:34 and 16:00 (PDT), going southward (Figure 4a). Time intervals between observations are irregular and range from 5 to 90 min.

### *2.3.3. Numerical experiments*

A primary goal of this particular application is to produce a short-term SRKW location forecast on the time-scale of hours to aid in mitigation of ship collision risk or acoustic disturbance. This

time-scale is considered useful by marine operations in that pilots have sufficient warning to alter the pathways and speed of incoming vessels. We carry out two experiments: (i) an assimilation experiment, and (ii) a forecast experiment. For (i), we use all location observations from 10:34 to 16:00 ($n = 14$) in the data assimilation cycle to sequentially provide online probabilistic location estimates. The forecast experiment (ii) aims at assessing the capability of the prediction system over time horizons of interest (a few hours). We assimilate the visual observations from 10:34 to 12:25, i.e. the first $n = 6$ observations. We then forecast the pod locations up to 3.5 hrs ahead starting from 12:30 (corresponding to the current time) out to 16:00 (i.e., future times). These are then compared to the observations from 12:55 to 16:00 that were artificially removed from the system (not assimilated) for the purposes of validation. We computed the direct position errors as the discrepancy between the observations and predicted locations $e_t = y_t - \hat{x}_t$. We represented the forecast probability density function using a kernel density estimate (KDE) of the ensemble.

## 3. Results

### *3.1. Idealized example*

#### *3.1.1. Simulation experiments*

The four simulation experiments that make use of data assimilation for nowcasting are presented in Figure 2, as detailed in Section 2.2.2. The true animal reference track starts at the top right of the 2D domain (the 'X' in Figures 2a, c, e and g). The animal first moves towards the origin, consistent with the potential field and its velocity persistence. When near the origin, this section of the track corresponds to low values of $\phi_t$ where the gradient of the potential field is small, and mimics a foraging behaviour (Figures 2b, d, f and h). Finally, the animal moves towards the bottom left of the plane as $\phi_t$ increases towards 1 and velocity persistence dominates, i.e. a transiting mode. At the end of the track, the animal then loops back towards the origin as $\phi_t$ decreases and habitat preference asserts itself.

The first two simulation experiments assimilate accurate and regular location data close to the true track (Figures 2a and c). These highly informative data lead to predicted nowcast animal

positions very close to the observations, and the spread of the ensemble is quite limited. No major difference is found in terms of predicted positions between the first two simulations and the predicted positions are very similar to the true track. The last two simulation experiments ingest less accurate and more irregular data spread around the true track (Figures 2e and g). These less informative data lead to predicted positions that deviate noticeably from the observations, especially around the origin. There are small differences in the two cases since the movement model only uses habitat preference in one case. Ensembles generally spread more widely around the median due to the less informative location data.

The behavioural parameter, velocity persistence $\phi_t$, is estimated along with the location via the state augmentation procedure (Figures 2b, d, f and h). In Simulation 1, accurate and regular location data as well as the use of habitat preference, allows reliable estimation of the temporal pattern of the true $\phi_t$ (Figure 2b). The ensemble spread is generally large and bigger when $\phi_t$ decreases. The time lag between the true and estimated $\phi_t$ is small which means that the data assimilation is able to quickly learn the proper value for the behavioural parameter from the accurate and frequently available location data. In Simulation 2, despite the location data being the same, the lag is much larger. This suggests that the use of habitat preference in Simulation 1 improves online behavioural parameter estimates (this makes sense since the data were generated assuming habitat preference). In Simulation 3 and 4, recovery of $\phi_t$ also indicates that the accuracy and regularity of location data, as well as the use of habitat preference, influence the quality of the online behavioural parameter that can be estimated.

### *3.1.2. Forecast skill*

For all simulation experiments, RMSE increases with forecast horizon, as expected (Figure 3a). Forecast models with regular, lower error observations (Simulations 1 and 2) perform better than irregular, higher error observations (Simulations 3 and 4) up to forecast time horizons of 12 time units. Therefore, accurate and regular location data provides better short-term forecast of animal locations. However, beyond a forecast horizon of 12 units, simulations with a potential field (intensity map) increase error forecasts rapidly and exceed the error forecasts of Simulations 2 and

4 without a potential field. The use of habitat preference in the movement model steers the animal towards the origin (towards the higher potential function intensity). However, the second part of the observed track corresponds to the animal moving away from the center (i.e. going against the potential function). This contradictory feature of this particular realization of the observations thus produces higher long-term forecast error. Finally, the forecast error is higher when the animal is in a transiting mode ($\phi_t \rightarrow 1$) (Figure 3b), particularly when the forecast time horizon increases.

### *3.2. Southern Resident Killer Whale application*

#### *3.2.1. Assimilation experiment*

Figure 4 shows the results for the online estimation of J pod location nowcasts. The visual observations occur at irregular time intervals over the course of the day and are, on occasion, clustered closely together in space and time, consistent with the observation error. The movement of J pod is generally to the south and veering eastward and covers about 50 km. There are variations in both speed and directional persistence, including a brief doubling back to the north off the west coast of San Juan Island between 12:25 and 12:55 (Figure 4a). The historical whale occupancy for August, which acts as the potential function, is highest off the southwest coast of San Juan Island indicating preferred habitat. Predicted whale locations conform well to the observed locations at the observation times, as expected given that these nowcast estimates are sequentially corrected to be near the visual sightings as they become available.

Figures 4b and c provide a detailed view of how the data assimilation cycle operates and performs by showing the individual components of the nowcast location state in terms of the easting (longitudinal) and northing (latitudinal) coordinates. The full ensemble that represents the estimated whale location is shown along with its median value, i.e. the most likely whale location. The prediction-correction aspect of the data assimilation procedure is evident. To specifically illustrate this key feature, consider how the system reacts to the observed temporary northward reversal in whale direction seen at 12:55 (G). Immediately prior to this, a forecast was made after assimilating the last observation at 12:25 (F). The increasing uncertainty, or spread, in the forecast ensemble is clear, reaching its maximum at 12:50, just prior to the next observation. With the data

assimilation system then receiving the observation at 12:55 (G), the particles nearest the observation are re-sampled leading to a reduction in the ensemble spread and a median estimate closer to the observation.

Overall, the movement is biased eastward and southward due to the drift term $\mu_t$ that remains mostly positive in both longitudinal and latitudinal directions (Appendix S1: Figure S1c). More precisely, from C to K, the movement is forced eastward with a higher drift term in the longitudinal direction (Appendix S1: Figure S1c). After K, the movement is slightly biased southward and eastward corresponding to the pod entering an area of high whale intensity southwest of San Juan Island (Figures 4a, Appendix S1: Figure S1a).

Along with the whale location, the velocity persistence parameter, $\phi_t$, was also estimated online using state augmentation. Its value remains $< 0.5$ for most of the time series (Appendix S1: Figure S1b) which indicates an intermediate behaviour state between transiting ($\phi_t$ tending towards 1) and resting ($\phi_t$ tending towards 0), consistent with the general features of observed trajectory showing systematic north-south movement with some reversals. Two periods do show $\phi_t > 0.5$ which may correspond to exploratory behaviours (e.g., the northward reversal from F to G, Figure 4c).

### *3.2.2. Forecast experiment*

The forecast experiment demonstrates short-term predictions on the time scale of hours. Here, the first six observations (A-F) until 12:25 are assimilated, with the associated whale location nowcasts mimicking the assimilation experiment (Figures 4a, 5a). After the final observation F is assimilated, the longitudinal drift decreases 68% whereas the latitudinal drift increases 77% until observation N (Appendix S1: Figure S1a, f). Both drift components remain positive throughout the time series, which indicates slow directional movement towards the south and east, with J pod predicted to remain southwest of San Juan Island. The directed movement is in general smaller than in the assimilation experiment, and largely due to the whale intensity field (Appendix S1: Figure S1f) since there are no observations being assimilated to draw the whales southward. Therefore, in the absence of new observations, the pod has weak movement directionality. This may contribute to an increase in the forecast error over the experiment's time horizon.

Forecast errors are quantified with the direct position error (DPE) metric (Figure 5b). The main feature of note is the growing error as forecast time horizon increases. The median DPE remains less than 5 km for a forecast out to 2.5 hrs and exceeds 10 km for the final 3.5 hr forecast in comparison to observations at these time horizons. The other key metric of forecast skill is the forecast uncertainty, here quantified by the range associated with each DPE. This range is the spread of the forecast ensemble about the future observation, or the 90% outer credible interval (i.e., the most likely region for whales). This interval increases with the forecast time horizon, and shows a forecast uncertainty of about 10 km for 1 hr forecasts, 15 km for 1.5-2.5 hrs forecasts, and 30 km uncertainty for a 3.5 hr forecast. Figures 5c-f show the kernel-density estimated forecast probability density functions (PDF) for selected observations. In general, the forecast PDF increases its spatial extent with larger forecast time horizons. It is also clear that the forecast PDF has a shape that is distinctly non-Gaussian. For forecast horizons of 0.5 hr, 1.5 hrs, and 2.5 hrs, the forecast PDF overlaps with the corresponding observations (G,J,M), but for the 3.5 hr forecast it does not overlap with the final observation N.

## 4. Discussion

In this paper, we developed the statistical framework for a real-time forecasting system for animal movement for the purpose of advancing understanding of movement ecology through adaptive learning, and for providing a flexible framework that could potentially be operationalized to facilitate management and conservation of animal populations. A non-dimensional idealized case for generic animal movement demonstrated the key features of the system. We then presented a prototypical, but realistic, prediction system for endangered SRKWs.

Improving animal movement models is key for ecological forecasting, since forecast skill rests on the dynamic model's efficacy (Payne et al., 2017, Dietze et al., 2018). We used a stochastic movement model that is extremely flexible with respect to the animal trajectories it can generate, a salient feature of most stochastic movement models. Hence, we further constrained the movement model by integrating preferred habitat through a drift term computed as the local gradient of

a potential function. The behavioural persistence parameter $\phi_t$ acts as a weighting factor for the past velocity and the current drift component. Adaptive online learning of the persistence parameter was implemented as part of a forecasting system using state augmentation (Kitagawa, 1998). The dynamical interplay between the drift and velocity meant that high values of the persistence parameter ($\phi_t \rightarrow 1$) can reduce the influence of the drift term, which in the absence of location observations, can lead to weak directionality, and, in some cases, low predictive skill. The idealized example showed that the forecast error was typically higher when the persistence was larger, i.e., when animal was transiting. Low predictive skills also arose when the observations were inconsistent with the underlying preferred habitat. Improvement relies on better estimates of the positions and persistence parameter through the refinement of the movement model structure and the integration of high quality (small observation error and high sampling rate) location observations along with animal pathway information and environmental drivers.

Ecological understanding comes from a forecasting system as a byproduct of online learning, both in terms of adaptively estimating informative movement model parameters, and ultimately in the iterative refinement of the movement models themselves. Forecasting places the emphasis on optimizing predictive skill, rather than the traditional metrics for retrospective studies (goodness of fit, or cross-validation skill). Forecast skill is the key metric that underlies location estimates and parameter and model refinement. It depends on both accuracy of the initial condition (nowcast), and how well the movement model represents the actual ecological dynamics. We used a basic RMSE skill metric and the direct position error, but recognize that other skill metrics are possible (e.g., bias, mean absolute error, threat score, Brier skill score, ROC skill score; Hamill and Juras (2006)), as well as comparison to basic persistence forecasts. Many of these metrics would, however, need adaption for the probabilistic case.

Our SRKW demonstration system shows clearly how general ecological forecasting framework can be adapted to particular setting with aim of achieving conservation goals through the fusion of movement models and observations. We made use of one day of visual sightings compiled to pod level which would be, in practice, available at irregular intervals in real-time. These opportunistic

data are limited by viewing conditions (e.g., daylight, sea state), and in the near future we plan to integrate SRKW detections from passive acoustic monitoring. The statistical character of this data type is very different from the visual detection data. However, our system is flexible and can assimilate multiple complex data types through suitable specification of the measurement model. Future challenges to be addressed with respect to the SRKW measurement model includes: recording errors, the consequences of pod splitting, and detection false positives.

The use of habitat preference and the avoidance of shallow waters in our SRKW application allowed the model to bias the whale movement in the correct direction and provide moderate forecast error (5 km) up to 2.5 hrs. Considering that a container vessel transiting the Salish Sea moves at a median speed of 18 knots (Joy et al., 2019), the vessel would be 83 km from the median position of the whale pod, a distance well outside the envelope of error - hence our forecasting system provides reasonable lead-in time to mitigate vessel-whale interactions. With additional efforts to incorporate real-time location observations along with model refinements to improve directionality, this could become an operational tool for managing SRKW in the Salish Sea. Towards this end, future information that we are considering for our target SRKW system includes: pathway information derived from visual sightings (Olson et al., 2018), prey fields (Kent et al., 2020), and habitat use models (Abrahms et al., 2019). Another promising direction is to couple the system to other environmental or biological forecasts (Payne et al., 2017), such as, for the Salish Sea, an already existing oceanographic forecasting system. (e.g., Olson et al. (2020)). While our proof-of-concept system provides the basis for an SKRW forecasting system, operationalizing it is not trivial: quality controlled, real-time data feeds are required; ensemble data assimilation must be robust (e.g., particle collapse); and movement models need to be further refined to incorporate environmental information.

In summary, our forecasting system for animal tracking provides a synthesis tool for assimilation of the real-time location information into a movement model, and makes use of potential functions and online parameter learning. It can assimilate any direct (e.g., visual) or indirect (e.g., passive acoustic detections) location observations and can handle multiple location observations at the

same time (i.e., within the data assimilation window), or data available irregularly in time. We have extended the ensemble approach to a coherent group (whale pod) but it could integrate any level of aggregation, i.e., multiple individuals each represented by ensembles and interacting with one another (Russell et al., 2017). Our system is thus flexible enough to be adapted to any data types, movement models, animal species and environmental conditions. Importantly, the forecasting framework provides a step towards making more effective use of data streams on animal movement with a goal of forecasting and proactive management of animal populations, especially in the context of human-wildlife conflicts. Pursuing real-time animal movement prediction will drive advances in ecology by encouraging practitioners to confront their bio-logging data with model predictions and so drive the iterative refinement of movement models, ultimately leading to improved understanding of ecological processes. The approach could even be used for retrospective studies and so provides a complementary way to better interpret and understand the ecological implications of movement data. We anticipate that the continued improvement of such a system will provide for ecological hypotheses testing and the refinement of predictive movement models to drive future insights into animal ecology.

**Acknowledgements**

We thank The Whale Museum for providing sightings data, and A. Harris and E. Cummings for preparing the database. We thank J. Watson for sharing the SRKW intensity fields. We acknowledge H. Yurk of Fisheries & Oceans Canada (DFO) for championing our ideas. The project was funded by DFO as part of the Whale Detection and Collision Avoidance program.

**References**

Abrahms, B., H. Welch, S. Brodie, M. G. Jacox, E. A. Becker, S. J. Bograd, L. M. Irvine, D. M. Palacios, B. R. Mate, and E. L. Hazen (2019). Dynamic ensemble models to predict distributions and anthropogenic risk exposure for highly mobile species. *Diversity and Distributions 25*(8), 1182–1193.

Auger-Méthé, M., K. Newman, D. Cole, F. Empacher, R. Gryba, A. A. King, V. Leos-Barajas, J. M. Flemming, A. Nielsen, G. Petris, et al. (2020). An introduction to state-space modeling of ecological time series. arXiv preprint arXiv:2002.02001.

Barlow, D. R. and L. G. Torres (2021). Planning ahead: Dynamic models forecast blue whale distribution with applications for spatial management. *Journal of Applied Ecology 58*(11), 2493–2504.

Breece, M. W., M. J. Oliver, D. A. Fox, E. A. Hale, D. E. Haulsee, M. Shatley, S. J. Bograd, E. L. Hazen, and H. Welch (2021). A satellite-based mobile warning system to reduce interactions with an endangered species. *Ecological Applications 31*(6), e02358.

Brillinger, D. R., H. K. Preisler, A. A. Ager, and J. Kie (2012). The use of potential functions in modelling animal movement. In *Selected Works of David Brillinger*, pp. 385–409. Springer.

Dietze, M. C., A. Fox, L. M. Beck-Johnson, J. L. Betancourt, M. B. Hooten, C. S. Jarnevich, T. H. Keitt, M. A. Kenney, C. M. Laney, L. G. Larsen, et al. (2018). Iterative near-term ecological forecasting: Needs, opportunities, and challenges. *Proceedings of the National Academy of Sciences 115*(7), 1424–1432.

Dowd, M., E. Jones, and J. Parslow (2014). A statistical overview and perspectives on data assimilation for marine biogeochemical models. *Environmetrics 25*(4), 203–213.

Dowd, M. and R. Joy (2011). Estimating behavioral parameters in animal movement models using a state-augmented particle filter. *Ecology 92*(3), 568–575.

Gervaise, C., Y. Simard, F. Aulanier, and N. Roy (2021). Optimizing passive acoustic systems for marine mammal detection and localization: Application to real-time monitoring north atlantic right whales in gulf of st. lawrence. *Applied Acoustics 178*, 107949.

Gordon, N. J., D. J. Salmond, and A. F. Smith (1993). Novel approach to nonlinear/non-gaussian

bayesian state estimation. In *IEE Proceedings F-radar and signal processing*, Volume 140 Issue 2, pp. 107–113. IET.

Hamill, T. M. and J. Juras (2006). Measuring forecast skill: Is it real skill or is it the varying climatology? *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography 132*(621C), 2905–2923.

Hooten, M. B., D. S. Johnson, B. T. McClintock, and J. M. Morales (2017). *Animal movement: statistical models for telemetry data*. CRC press.

Ionides, E. L., A. Bhadra, Y. Atchadé, A. King, et al. (2011). Iterated filtering. *The Annals of Statistics 39*(3), 1776–1802.

Johnson, D. S., J. M. London, M.-A. Lea, and J. W. Durban (2008). Continuous-time correlated random walk model for animal telemetry data. *Ecology 89*(5), 1208–1215.

Joy, R., D. Tollit, J. Wood, A. MacGillivray, Z. Li, K. Trounce, and O. Robinson (2019). Potential benefits of vessel slowdowns on endangered southern resident killer whales. *Frontiers in Marine Science 6*, 344.

Kent, C. S., P. Bouchet, R. Wellard, I. Parnum, L. Fouda, and C. Erbe (2020). Seasonal productivity drives aggregations of killer whales and other cetaceans over submarine canyons of the bremer sub-basin, south-western australia. *Australian Mammalogy 43*(2), 168–178.

Kitagawa, G. (1998). A self-organizing state-space model. *Journal of the American Statistical Association 93*(443), 1203–1215.

Kloeden, P. E. and E. Platen (2013, April). *Numerical Solution of Stochastic Differential Equations*. Springer Science & Business Media. Google-Books-ID: r9r6CAAAQBAJ.

McClintock, B. T., J. M. London, M. F. Cameron, and P. L. Boveng (2017). Bridging the gaps in animal movement: hidden behaviors and ecological relationships revealed by integrated data streams. *Ecosphere 8*(3), e01751.

McWhinnie, L. H., P. D. O'Hara, C. Hilliard, N. Le Baron, L. Smallshaw, R. Pelot, and R. Canessa (2021). Assessing vessel traffic in the Salish Sea using satellite AIS: An important contribution for planning, management and conservation in southern resident killer whale critical habitat. *Ocean & Coastal Management 200*, 105479.

Michelot, T., R. Glennie, C. Harris, and L. Thomas (2021). Varying-coefficient stochastic differential equations with applications in ecology. *Journal of Agricultural, Biological and Environmental Statistics 26*(3), 1–18.

Olson, E. M., S. E. Allen, V. Do, M. Dunphy, and D. Ianson (2020). Assessment of nutrient supply by a tidal jet in the northern Strait of Georgia based on a biogeochemical model. *Journal of Geophysical Research: Oceans 25*(8), e2019JC015766.

Olson, J. K., J. Wood, R. W. Osborne, L. Barrett-Lennard, and S. Larson (2018). Sightings of Southern Resident Killer Whales in the Ssalish Sea 1976–2014: the importance of a long-term opportunistic dataset. *Endangered Species Research 37*, 105–118.

Patterson, T. A., A. Parton, R. Langrock, P. G. Blackwell, L. Thomas, and R. King (2017). Statistical modelling of individual animal movement: an overview of key methods and a discussion of practical challenges. *AStA Advances in Statistical Analysis 101*(4), 399–438.

Payne, M. R., A. J. Hobday, B. R. MacKenzie, D. Tommasi, D. P. Dempsey, S. M. Fässler, A. C. Haynie, R. Ji, G. Liu, P. D. Lynch, et al. (2017). Lessons from the first generation of marine ecological forecast products. *Frontiers in Marine Science 4*, 289.

Russell, J. C., E. M. Hanks, M. Haran, D. Hughes, et al. (2018). A spatially varying stochastic differential equation model for animal movement. *Annals of Applied Statistics 12*(2), 1312–1331.

Russell, J. C., E. M. Hanks, A. P. Modlmeier, and D. P. Hughes (2017). Modeling collective animal movement through interactions in behavioral states. *Journal of Agricultural, Biological and Environmental Statistics 22*(3), 313–334.

Wall, J., G. Wittemyer, B. Klinkenberg, and I. Douglas-Hamilton (2014). Novel opportunities for wildlife conservation and research with real-time monitoring. *Ecological Applications 24*(4), 593–601.

Watson, J., R. Joy, D. Tollit, S. J. Thornton, and M. Auger-Méthé (2019). Estimating animal utilization distributions from multiple data types: A joint spatiotemporal point process framework. *The Annals of Applied Statistics 15*(4), 1872–1896.

Williams, H. J., L. A. Taylor, S. Benhamou, A. I. Bijleveld, T. A. Clay, S. de Grissac, U. Demšar, H. M. English, N. Franconi, A. Gómez-Laich, et al. (2020). Optimizing the use of biologgers for movement ecology research. *Journal of Animal Ecology 89*(1), 186–206.

Yates, K. L., P. J. Bouchet, M. J. Caley, K. Mengersen, C. F. Randin, S. Parnell, A. H. Fielding, A. J. Bamford, S. Ban, A. M. Barbosa, et al. (2018). Outstanding challenges in the transferability of ecological models. *Trends in ecology & evolution 33*(10), 790–802.

# List of figures

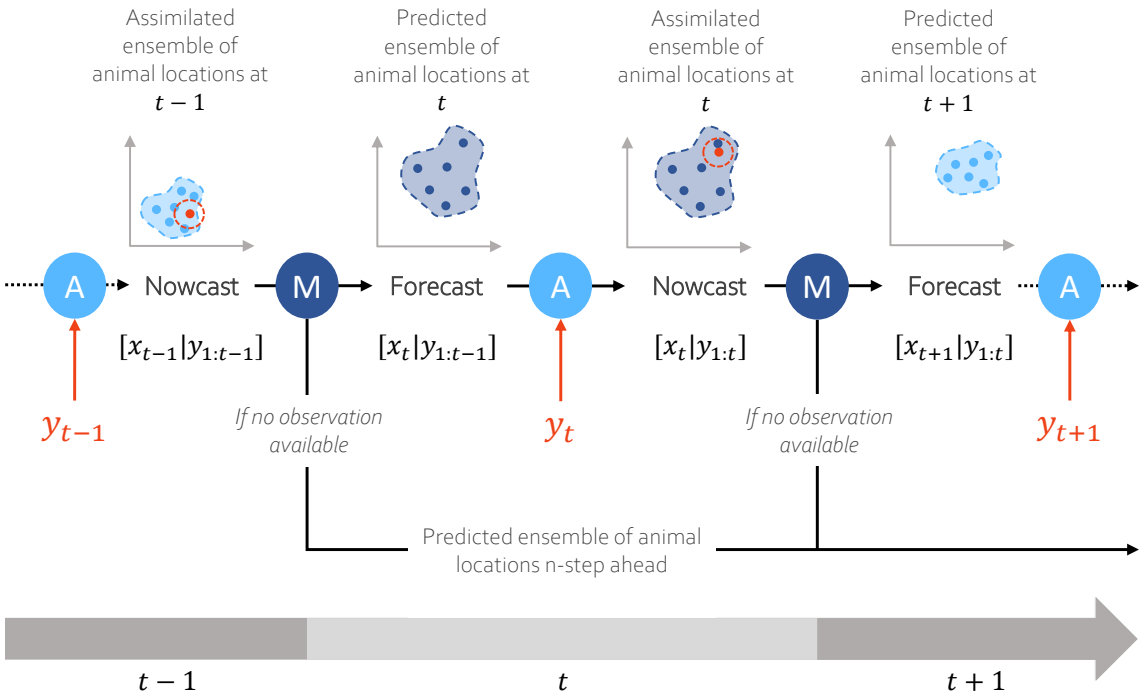


Figure 1: Schematic of the data assimilation cycle used for the animal prediction system. It shows a single stage transition of this probabilistic system from time $t-1$ to $t$, and how it toggles between movement model forecasts and particle filter based assimilation of incoming observations (see text for further details). Light and dark blue dots represent ensemble members (particles) at the nowcast and forecast steps, respectively. Red dots are the location observations and red circles correspond to measurement errors. Assimilation and forward model prediction are symbolized by A and M, respectively.

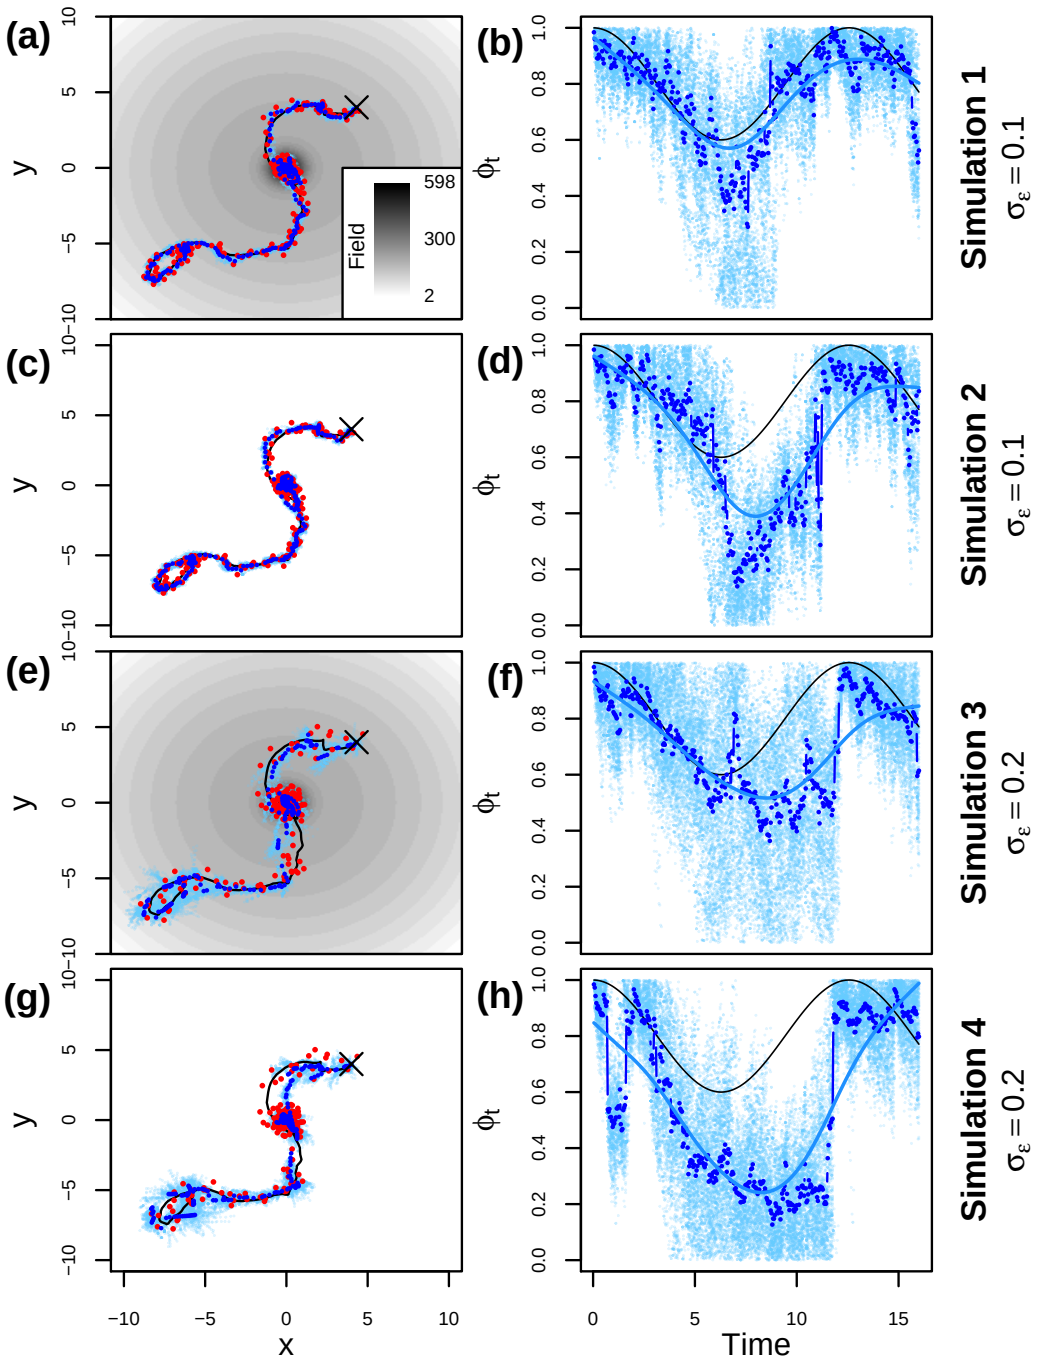


Figure 2: Idealized example. Left panels (a,c,e,g): Observed (red dots) and predicted animal locations (blue dots are the ensemble median; light blue dots show full ensemble). The true animal track is shown (black line) along with initial position (black cross). In panels (a) and (e), the grey scale represents potential field reflecting the animal's preferred habitat. Panels (a) and (c) present high quality location data ($\sigma_\epsilon = 0.1$, observations every time step) as opposed to panels (e) and (g) that present lower quality location data ($\sigma_\epsilon = 0.2$, observations every second time step). Right panels (b,d,f,h): Time-varying estimates of the velocity persistence parameter $\phi_t$. The black line show the true persistence velocity used for computing the true track (a sine wave), and estimation results (ensemble, light blue dots; ensemble median, blue dots; fitted smooth curve, blue lines). Note that spatial coordinates and time vectors are non-dimensional.

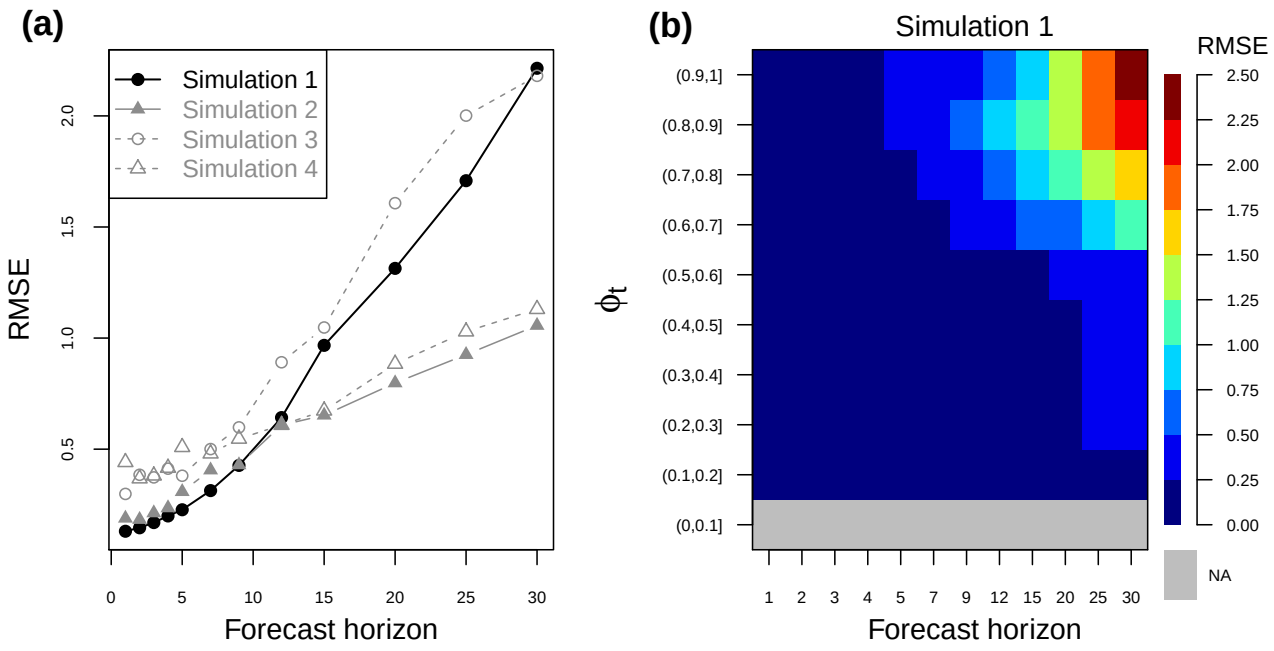


Figure 3: (a) $n$-step ahead error forecast (RMSE) of the idealized example. (b) Heatmap representing the relationship between the behavioural parameter $\phi_t$ and the $n$-step ahead error forecast (RMSE) for Simulation 1 (low observation error $\sigma_\epsilon = 0.1$, observations every time step, and drift term).

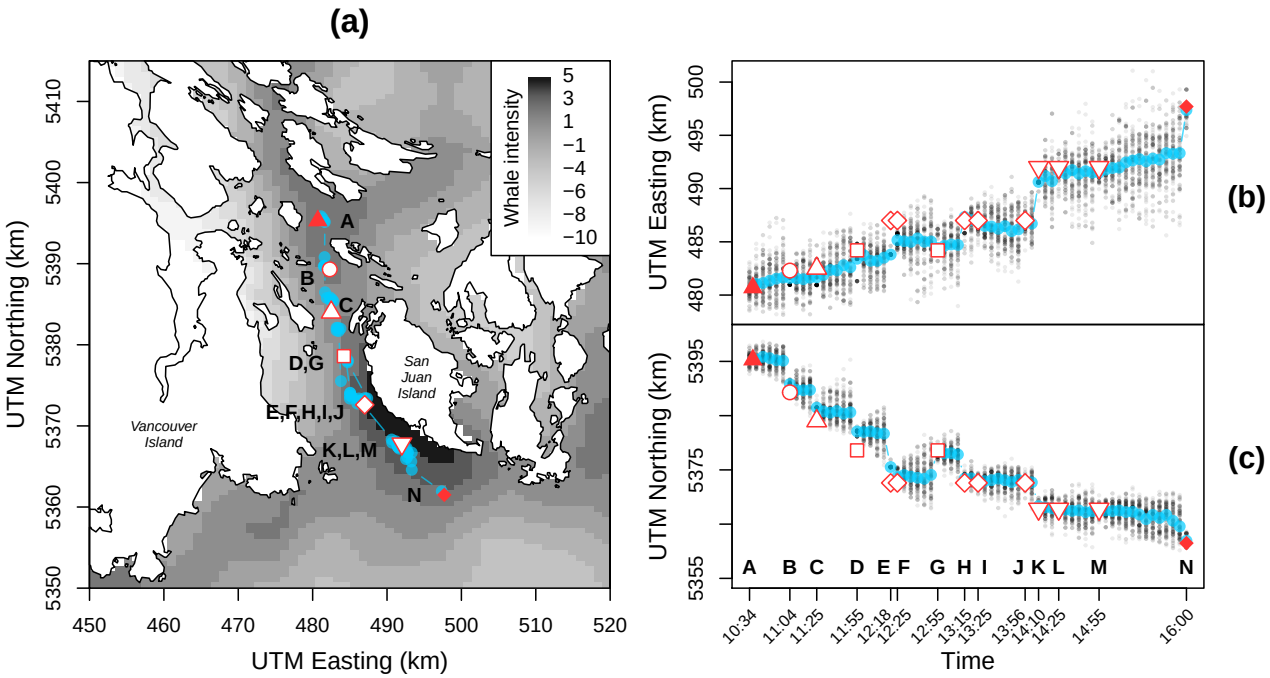


Figure 4: Assimilation experiment. Panel (a): Visual observations of SRKW J pod on 18 August 2016 (white symbols with red outline) and predicted whale locations (blue dots, ensemble median). Solid red symbols represent the starting (10:34) and ending (16:00) observations for the day. Letters A to N designate the chronology of these observations, with A being the first observation. The grey scale represents the whale intensity field of J pod in August, expressed in log scale (from Watson et al. (2019)). Panels (b) and (c): UTM easting and northing coordinates of whale locations including ensembles (grey dots) and their median (blue dots). Symbols denote the visual sightings location observations following Panel (a).

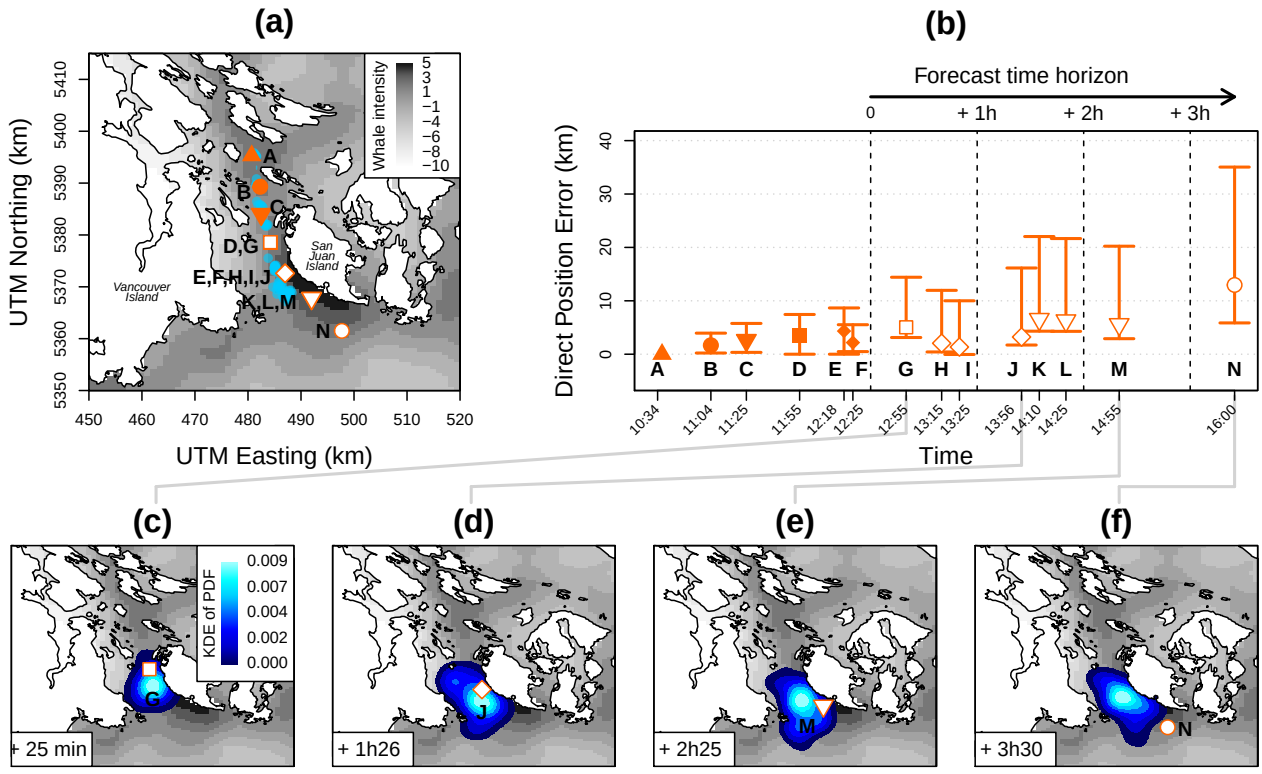


Figure 5: Forecast experiment. Panel (a): Follows Figure 4 except filled solid orange symbols represent first data point used for assimilation, and unfilled symbols represent location observations used for $n$-step ahead forecast validation (see text for details). Panel (b): Direct position error of the forecast (km) against time. Symbols show the discrepancy of the ensemble median and the observed location with the range being 5% and 95% percentiles of position error associated with the full ensemble. Panels (c) to (f): Kernel density estimates of the forecast probability density function of whale locations up to 3.5 hours ahead shown together with the future observation.