



METHODS AND INITIAL RESULTS OF THE TAGGING DATA ANALYSIS OF YELLOWFIN TUNA IN THE INDIAN OCEAN

TOBIAS K. MILDENBERGER , DAN FU , AND ANDERS NIELSEN

Discussions (for information and open discussion):

The paper is to inform meeting members of the progress of the IOTC-DTU project on the analysis of the yellowfin tuna tagging data. It covers the movement model and the developments needed to fit it to the Indian Ocean data, as set out in Sections 2 and 3, and the tagging, environmental and fishing effort data, as set out in Section 4. It also presents initial estimates of movement from all usable recaptures, as set out in Section 5. The results show that the effect of ocean currents and the overall rate of dispersal are robust. The estimated habitat preferences are not robust yet, because recapture positions largely reflect where purse-seine effort is, as set out in Sections 5.3 and 5.6.

Feedback (to collect input and suggestions):

The paper is to request review and comments on the remaining steps of Part I, as set out in Section 6. These steps are:

- including fishing effort in the likelihood;
- how to pool fleets and gears into a small number of groups with separate catchabilities, on which guidance from the Working Party is specifically requested;
- how to treat recaptures with unknown school type or with several candidate purse-seine sets;
- which values to use for tag reporting rates, tag shedding and tagging mortality, and which sensitivity runs to make, including one with the RTTP-IO tags only, matching the 2024 assessment.



Abstract

This report presents progress on the collaborative IOTC-DTU project to estimate movement, mortality and abundance of yellowfin tuna in the Indian Ocean from tagging data. The project responds to the recommendation of the 2023 independent review of the yellowfin assessment to analyse the tagging data outside the assessment with a fine-scale spatio-temporal model. Movement is described by an advection-diffusion model with three components: diffusion, taxis towards preferred habitat defined by environmental fields, and advection by ocean currents. The model is fitted with a continuous-time Markov chain on a projected grid and is implemented in the R package *admove*. Several developments were needed to fit the Indian Ocean data. The main ones are uniformisation for the matrix exponential, shared calculations for tags released together, ocean currents as an advection field, and spline-based preference functions. Eight model configurations were fitted to 3389 recaptured tags. Most of these tags were released off Tanzania and Kenya and recaptured by purse seiners. The effect of currents is the most robust result: about a third of the current velocity in the upper 75 m is transferred to the fish. The estimated habitat preferences are not robust. They change when diffusion depends on temperature and when the grid is refined, and a temperature field mirrored east-west fits better than the real one. The recaptures stop where purse-seine effort stops, so the preference functions largely describe the distribution of fishing effort. The next step is to include fishing effort, tag reporting rates and tag loss in the likelihood, before the Part I report at the end of 2026.

Methods and initial results of the tagging data analysis of Yellowfin tuna in the Indian Ocean

Tobias K. Mildenerberger¹, Dan Fu², and Anders Nielsen¹

¹Technical University of Denmark, National Institute of Aquatic Resources (DTU Aqua), Denmark

²Indian Ocean Tuna Commission (IOTC) Secretariat, Victoria, Seychelles

1 Background

This document reports progress on the collaborative project between the Indian Ocean Tuna Commission (IOTC) and the Technical University of Denmark (DTU) to estimate movement rates, mortality rates and abundance indices for yellowfin tuna in the Indian Ocean from tagging data. The project is structured in two parts: Part I develops a movement model that estimates fine-scale movement patterns as a function of environmental fields and habitat preference functions, with delivery at the end of 2026. Part II builds on the resulting movement estimates to develop a spatiotemporal population model for mortality rates and relative, and possibly absolute, biomass indices, with delivery at the end of 2027. The work builds directly on the methodology developed for skipjack tuna in the eastern Pacific Ocean (EPO) (Mildenerberger et al., 2022, 2023, 2024, 2025), and recently reviewed for the western and central Pacific Ocean (WCPO) (Mildenerberger et al., 2026b).

The project responds to a recommendation of the independent review of the IOTC yellowfin assessment (Maunder et al., 2023). The review found the inclusion of the tagging data in the assessment problematic: the releases are spatially restricted, so the tagged fish are not initially mixed with the population, it is unclear after how many quarters, if at all, they become mixed, and a mixing period of several quarters greatly reduces the information in the tags. It also noted that the assessment does not fit the tagging data well. Among its priority research areas, the review therefore recommended to evaluate the tagging data outside the assessment with a fine-scale spatio-temporal model that addresses the lack of mixing, which may estimate movement, natural mortality, biomass or fishing mortality for use in the assessment. The 2024 assessment (Urtizberea et al., 2024) confirms the lack of fit: the model underestimates the purse-seine tag returns of the main recovery period 2007–2009, most strongly in the quarters immediately after the four-quarter mixing period of the large 2006 releases, which “*suggests that even the four-quarter mixing period may not be sufficient to allow for adequate dispersal of tagged fish in the population*” (p. 34; Urtizberea et al., 2024). Following the review, the 2024 assessment added a reference model in which the tagging data are down-weighted by 90%.

This report describes the movement model as it is applied to the Indian Ocean data (Section 2), the developments that were needed to fit these data (Section 3) and the data themselves (Section 4), presents initial results from a set of model configurations fitted to all usable yellowfin recaptures (Section 5), and sets out the remaining steps of Part I until the final report at the end of the year (Section 6). The results are initial: they show which conclusions are already robust and which depend on work that is still to be done, in particular on accounting for fishing effort.

2 Methodology

2.1 The movement model

Movement is described by an advection–diffusion process (Sibert et al., 1999; Thorson et al., 2021) with three components. *Diffusion* represents undirected movement and is the component that release–recapture pairs are most informative about. *Taxis* represents directed movement up the gradient of a habitat preference function, which is built as a sum of smooth functions of environmental fields (for example temperature or mixed layer depth). *Advection* represents passive transport by ocean currents. With $\mathbf{x} = (x, y)$ the position and $u(\mathbf{x}, t)$ the probability density of the position of a tagged fish at time t ,

$$\frac{\partial u}{\partial t} = \nabla^2 (D u) - \nabla \cdot (\mathbf{v} u), \quad \mathbf{v} = \kappa \nabla h_T + \gamma \mathbf{c}, \quad (1)$$

where $D(\mathbf{x}, t)$ is the diffusion coefficient (distance² per time), $h_T(\mathbf{x}, t)$ the habitat preference, κ a fixed scale that carries the units of the taxis term, $\mathbf{c}(\mathbf{x}, t)$ the current velocity and γ an entrainment coefficient.

2.2 Habitat preference functions

The preference function and the logarithm of the diffusion coefficient are sums of one-dimensional smooth functions of the covariate fields $c_i(\mathbf{x}, t)$,

$$h_T(\mathbf{x}, t) = \sum_i s_{T,i}(c_i(\mathbf{x}, t)), \quad \log D(\mathbf{x}, t) = \sum_i s_{D,i}(c_i(\mathbf{x}, t)). \quad (2)$$

Each smooth is a natural cubic spline through a small number of knots, and the estimated parameters are the values of the function at the knots. A fitted preference curve can therefore be read directly off the parameter estimates, with its uncertainty from the delta method. A smooth with one knot is a constant and with two knots a straight line, so constant, linear and non-linear responses are obtained from the same construction. Taxis enters the model only through the gradient of h_T , so the level of the preference function is not identifiable and the first knot value is fixed at zero; only differences in preference between covariate values are estimated.

In the earlier EPO and WCPO analyses each preference function was a single global polynomial through the knot values. That construction oscillates between knots and diverges outside the knot range, and because taxis is driven by the gradient of the preference function both defects become spurious directed movement. With the natural cubic spline the influence of a knot decays with distance and the curve is linear beyond the outer knots. Linear extrapolation still matters for diffusion, which is modelled on the log scale: a slope continued into covariate values that no tag experienced gives very large or very small diffusion rates in a few cells. The knots are therefore placed within the range of covariate values at the release and recapture positions, and the covariate fields are limited to that range (Section 4.2).

2.3 Advection by ocean currents

Currents enter as a velocity field, not as a covariate with a smooth function. The model estimates one entrainment coefficient γ , the fraction of the current velocity that is transferred to the fish: $\gamma = 0$ means that currents have no effect on the displacement and $\gamma = 1$ that the fish drifts with the water. The eastward and northward current components, averaged over the upper 75 m (Section 4.2), are converted from m s^{-1} to km per month along the axes of the projected grid, so that γ is dimensionless.

2.4 Estimation engines

Two estimation engines are available in the framework and share the same parameters. Both are discretisations of Equation 1; they differ in whether space or time is discretised.

2.4.1 Kalman filter (KF)

The Kalman filter (Harvey, 1990) tracks the position of each tag in continuous space and discrete time. It was developed for skipjack tuna in the eastern Pacific Ocean; Mildenerger et al. (2023) describe the likelihood for archival and conventional tags and the differentiable local interpolation of the environmental fields that it requires, and Mildenerger et al. (2025) the residuals of the recapture positions for this engine. Each prediction step moves the mean by $\mathbf{v}\Delta t$ and adds $2D\Delta t$ to the variance, which is the Euler step of the movement model above. Movement rates are evaluated at the predicted position itself, so no spatial grid is needed and the resolution of the covariate fields is used in full. The cost is that the position distribution is Gaussian and the rate fields are linearised around its mean, so the coastline cannot be represented directly. This engine is well suited to archival tags, where frequent position estimates keep the predicted mean close to the fish. For the Indian Ocean yellowfin data it was not reliable: a conventional tag provides no observation between release and recapture, so the predicted mean moves freely for up to 104 months, and with most releases made close to the coast it reaches land, where the environmental fields are undefined. The results in this report are therefore obtained with the second engine.

2.4.2 Continuous-time Markov chain (CTMC)

The continuous-time Markov chain engine discretises space onto a grid of n cells and keeps time continuous (Thorson et al., 2021). It corresponds to the matrix exponential approach of the first implementation for skipjack tuna in the eastern Pacific Ocean (Mildenerger et al., 2022). The state of a tagged fish is the cell it occupies, and the model tracks the probability of each cell, a vector $\boldsymbol{\pi}_t$ of length n that starts with probability one in the release cell.

The movement rates between cells are collected in a generator matrix $\mathbf{Q}$. A fish in cell i can move to each of its four neighbours j , at distance Δ , with the rate

$$q_{ij} = \frac{D_i}{\Delta^2} + \frac{v_{ij}}{2\Delta}, \quad (3)$$

where D_i is the diffusion coefficient at the centre of cell i and v_{ij} the component of the velocity $\mathbf{v}$ (taxis plus advection) in the direction of neighbour j . The first term is the standard five-point discretisation of the diffusion term of Equation 1. The second term is a central difference: directed movement increases the rate towards the neighbour that $\mathbf{v}$ points at and decreases it by the same amount towards the opposite neighbour. The rates are smooth functions of the parameters and add no numerical diffusion, but they are only non-negative, and $\mathbf{Q}$ only a valid generator, where directed movement over one cell is small relative to diffusion, $|v_{ij}|\Delta/D_i < 2$. On the grid used here this condition is not met in 2–5% of the cell-months at the estimates. The diagonal $q_{ii} = -\sum_{j \neq i} q_{ij}$ makes each row sum to zero. Land cells are not part of the grid and no rate leads into them, so the coastline is a reflecting boundary by construction and probability is conserved. All other entries of $\mathbf{Q}$ are zero: the matrix is sparse, with at most five non-zero entries per row.

The rates depend on the environmental fields and therefore change from month to month, but are constant within a month. Over a step of length Δt within one month the distribution is propagated exactly by the matrix exponential,

$$\boldsymbol{\pi}_{t+\Delta t}^\top = \boldsymbol{\pi}_t^\top \exp(\mathbf{Q}\Delta t). \quad (4)$$

The steps are aligned to the months of the covariate fields, so that a tag at liberty for a year needs twelve full steps and a partial step at either end. The likelihood contribution of a recaptured tag is the probability that the model assigns to the cell of recapture at the time of recapture, and the negative log-likelihood is the sum over tags of the negative logarithms of these probabilities. Uncertainty in the recapture position, where it is known, can be included by spreading the observation over neighbouring cells.

Matrix exponential. The matrix exponential is the computational bottleneck. $\exp(\mathbf{Q}\Delta t)$ is a dense $n \times n$ matrix even though $\mathbf{Q}$ is sparse, and for the Indian Ocean data computing it exhausted the available memory already for a subset of 300 tags. The model therefore uses uniformisation (Grassmann, 1977; Jensen, 1953; Sherlock, 2021), implemented in RTMB (Kristensen, 2026), which computes the distribution after each step directly, as a sum of sparse matrix–vector products, without forming the matrix and with a guaranteed accuracy. The number of products grows with the largest rate out of any cell, and thus with the largest diffusion rate on the grid, with the length of the time step and with the inverse of the squared cell size: halving the cell size quadruples the products per step and also the number of cells. This is what currently limits the spatial resolution, and why a single cell with an extreme extrapolated diffusion rate can make a fit infeasible.

Release events. The distribution π_t depends on where and when a fish was released, but not on the individual tag. Tags released in the same cell at the same time therefore share one forward calculation, which is evaluated at each of their recapture times. Because tagging cruises released many fish at a time, this reduces the computation considerably without any approximation.

2.5 Prediction diagnostics

The fit is checked by how well the model predicts the recaptures. For each recaptured tag, the CTMC engine gives the predicted distribution over grid cells at the time of recapture, exactly as the likelihood evaluates it. Two measures are reported: the prediction error, the distance between the observed recapture position and the mean of this distribution, summarised as the root mean squared error (rmse); and the share of recaptures inside the 50% and 95% prediction regions of the distribution. If the predicted distributions are correct, these shares are 50% and 95%; larger shares mean that the predictions are too wide, smaller shares that they are too narrow.

2.6 R package `admove`

The framework is implemented in the `admove` R package (Mildenberger et al., 2026a). The likelihood is written in plain R and differentiated by RTMB (Kristensen, 2026), the R interface to TMB (Kristensen et al., 2016), which gives exact gradients for quasi-Newton minimisation and delta-method standard errors on derived quantities, including the estimated preference curves, without any compiled C++ code. The estimated fine-scale rates can be aggregated to transfer probabilities between any set of assessment regions and time steps, which decouples the movement grid from the spatial structure of the assessment.

3 Adapting the model to the Indian Ocean data

The model, the software and the data preparation were developed further to fit the Indian Ocean yellowfin data. Most of the following developments were needed before a first fit to all tags was possible.

- **Projected coordinates.** The model treats space as a flat plane. In longitude and latitude, east–west distances are overstated away from the equator. Tags, grid and covariate fields are now all in an azimuthal equidistant projection in kilometres, centred on the median tag position (40.5°E, 6.1°S).
- **CTMC engine with uniformisation.** The Markov chain engine replaced the Kalman filter for these data, and the matrix exponential is computed by uniformisation (Section 2). The number of terms per step is capped, and the fit reports when the cap is reached at the estimates, because the likelihood is then truncated.
- **Tags grouped by release event, and monthly steps.** Tags released in the same cell at the same time share one calculation, and steps are aligned to the monthly covariate fields.
- **Ocean currents.** Currents are included as an advection field with an estimated entrainment coefficient, including the conversion of the current components to the axes of the projected grid.
- **Natural cubic splines.** The preference functions are natural cubic splines in place of the global polynomials of the earlier analyses, with knots placed within the covariate range experienced by the tags.
- **Covariate fields near land.** Cells whose centre lies on land have no covariate value, which initially removed the two coastal cells in which 83% of the recaptured tags were released. Covariate values are now extended a short distance from the sea onto land (Section 4.2).
- **Covariate range.** The covariate fields are limited to the range experienced by the tags, to avoid extreme extrapolated diffusion rates in cells that no tag visited.
- **Convergence checks.** Every fit is checked automatically for the status of the optimiser, the size of the final gradient, estimates at a bound, a positive definite Hessian and finite standard errors. Fits that fail are flagged and are not used for model comparison.
- **Diagnostics.** For every recaptured tag the model gives the predicted distribution of the recapture position, from which prediction errors and the coverage of prediction regions are calculated (Section 2.5). The model can also predict where fish released at a given place and time are expected to be over time, which is compared with the observed recaptures and with the distribution of fishing effort (Section 5).
- **Tag selection.** The selection of tags is recorded step by step from the database extract to the fit (Table 1), so that the effect of each filter on the released and the recaptured tags can be followed.
- **Fishing effort data.** The IOTC catch-and-effort data were explored, and the effort of the main purse-seine fleets was prepared by 1° cell and month (Section 4.3).

4 Data

4.1 Tagging data

The tagging data are an extract of the Indian Ocean tagging database, restricted to yellowfin tuna. They come mainly from the Regional Tuna Tagging Project – Indian Ocean (Hallier & Million, 2009), with additional small-scale tagging projects held in the same database. Table 1 follows the tags from the extract to the model, and Table 2 summarises the tags that are used. Figure 1 shows their release and recapture positions and the model domain, and Figure 2 the distribution of time at liberty.

The selection differs from that of the 2024 assessment (Urtizberea et al., 2024) in three respects. (i) The assessment uses the 54688 yellowfin releases of the main phase of the RTTP-IO (2005–2009), the same number as the RTTP-IO yellowfin releases in the extract, all of which date from 2005–2007. It excludes the small-scale programmes because they have not been analysed and their reporting rates are unknown. Of the 3389 tags in the model, 424 (13%) come from small-scale programmes; all tags released in 2003–2004 and in 2008–2009 belong to these programmes, and all RTTP-IO tags in the model were released in 2005–2007. A sensitivity run with the RTTP-IO tags only is planned (Section 6). (ii) The assessment uses 9916 recoveries up to the end of 2014; the model uses recaptures up to 2015, of which only three are from 2015. (iii) The assessment needs only the fishery, region and quarter of a recapture and can therefore use the 6531 recaptures that are known only from a set of candidate purse-seine sets, which the model currently excludes. This is the main reason why the assessment uses about three times as many recaptures.

Step	Removed	(recapt.)	Released	Recaptured
Yellowfin releases in the database extract			66543	10842
Missing release or recapture position	10	10	66533	10832
Recapture only from candidate purse-seine sets	6531	6531	60002	4301
Recapture without catch date or set information	494	494	59508	3807
Recaptured as another species	257	257	59251	3550
Outside the analysis period 2003–2015	3230	140	56021	3410
Recaptured at or before release	5	5	56016	3405
Implied speed above 200 km/day	2	2	56014	3403
Release or recapture position on land	7	6	56007	3397
Release or recapture outside the model domain	2851	8	53156	3389
Never recaptured	49767	0	3389	3389

Table 1: Yellowfin tags from the database extract to the model. Each row gives the tags removed at that step, how many of those were recaptured (in brackets), and the released and recaptured tags left after it. Tags that were never recaptured are not used by the movement likelihood. The analysis period is set by the coverage of the environmental fields.

Quantity	Value
Tags in the model (released and recaptured)	3389
Release events	703
Release years	2003–2009
Recapture years	2003–2015
Time at liberty (days), median	200
Time at liberty (days), mean	300
Time at liberty (days), range	1–3162
Tags at liberty < 30 days	554 (16.3%)
Displacement (km), median	851
Displacement (km), mean	910
Displacement (km), maximum	4463
Releases in the two main release cells	83.4%

Table 2: Summary of the yellowfin tags in the model.




Figure 1: Release and recapture positions in the projected coordinate system, with the model domain (blue cells, 500 km) and the tags that fall outside it.




Figure 2: Time at liberty of the 3389 tags in the model. Tags at liberty for more than 48 months are included in the last bar.

Several characteristics of these data determine what the model can resolve and what remains to be done.

- **Most releases are not used yet.** Of 66543 yellowfin releases, 10842 have a recapture record and 3389 enter the model. The movement likelihood used here is conditional on recapture, so releases that were never recaptured carry no information. They become informative once fishing effort and recapture probability are part of the model.

- **Most recaptures have no single recapture position.** 6531 recaptures are known only from a set of candidate purse-seine sets, with a median of six candidates per recapture. They are currently excluded. Including them requires the candidate positions and dates of each recapture and a likelihood that weights the candidates, for which fishing effort is the natural basis.
- **Releases are concentrated at the coast.** 83% of the tags in the model were released in two coastal cells off Tanzania and Kenya. The estimates are therefore dominated by how fish disperse from one area, and by where they can be recaptured.
- **Recaptures are mostly made by purse seine.** Purse seine accounts for 69% and pole and line for 19% of the recaptures in the model, so the recapture positions reflect where these fleets fish (Section 4.3).
- **Recapture date and species.** The recapture time is the reported catch date. The date on which a tag was returned is available for more tags, but it lags the catch by a median of 17 days and by up to several years, and would inflate the time at liberty and bias the diffusion rate downwards. Recaptures without a catch date are therefore not used. 257 tags released as yellowfin were recorded as another species at recapture, mostly skipjack. These are small fish, in the size range in which juvenile yellowfin and skipjack are hard to tell apart, and they are excluded.
- **Mostly juvenile fish.** The tags in the model were released on fish with a median fork length of 57 cm; 90% measured between 41 and 72 cm (range 23–129 cm; Figure 3). The lengths are recorded as reliable for 99% of these tags. Larger fish are rarer among the recaptured tags than among all releases, of which 5% were longer than 92 cm. Over all yellowfin releases in the extract, the share that was recaptured is about 17–18% for fish of 40–70 cm at release, 10–13% for 80–100 cm and 4% above 100 cm. The independent review asked for the cause of a higher return rate of larger fish to be determined, for example a higher tagging mortality of small fish (Maunder et al., 2023); in the extract the raw return rate is instead lower for larger fish, which may reflect the selectivity of the surface fisheries for small fish, and which measure the review referred to remains to be clarified. The fish grow during their time at liberty, but the estimates below describe the movement of juvenile and young adult yellowfin, not of large adults.
- **Short times at liberty.** 16% of the tags are recaptured within 30 days, less than the monthly time step of the model. A sensitivity run that excludes tags at liberty for less than one month or more than 24 months is included below. Overall, 30% of the tags are recaptured within three months, 47% within six months and 74% within a year. The 2024 assessment treats the first four quarters after release as a mixing period, during which the recaptures do not inform the fishing mortality of the population (Urtizberea et al., 2024). A movement model that follows each tag in space and time does not require mixing and uses these recaptures in full.
- **Double tags.** 737 of the tags in the model, and 2135 of all recaptured yellowfin in the extract, were released with two tags. 240 of the latter were recovered with only one tag. These data inform tag shedding directly (Section 6).
- **Tags are released in groups.** The 3389 tags come from 703 release events, and tags of the same event are often recaptured together. The model treats tags as independent, so the standard errors are likely to be too small.
- **No tracks from electronic tags.** The database records 233 releases of yellowfin with archival or pop-up tags, of which six were recovered. The extract holds one record per

tag with the release and recapture only, so no positions between them are available. The framework uses such tracks directly, and they inform diffusion far more than a release–recapture pair does.

- **Tags before 2003.** The analysis period 2003–2015 excludes 140 recaptured tags released between 1990 and 1995. Extending the environmental fields backwards would recover them, at the cost of longer run times.




Figure 3: Fork length at release of the tags in the model (outlined bars) and of all yellowfin releases (grey bars), as densities.

4.2 Environmental covariates and model grid

Environmental fields are taken from the Copernicus Marine Service global ocean physics reanalysis (E.U. Copernicus Marine Service Information, 2024) as monthly means for 2003–2015. The fields used are the temperature averaged over the upper 75 m, the mixed layer depth, and the eastward and northward current components averaged over the upper 75 m. The depth averages weight the model layers by their thickness. The upper 75 m cover the layer in which yellowfin tuna spend most of their time: an archival-tagged yellowfin tuna in the western Indian Ocean spent 85% of its time above 75 m, the deepest thermocline it experienced (Dagorn et al., 2006), and the species is generally restricted to the mixed layer and the top of the thermocline (Brill, 1994). The fields were processed in four steps.

- **Projection.** The fields were averaged to 0.25° and then projected onto a regular grid of 110 km cells in the projected coordinate system by area-weighted averaging. Averaging cannot produce values outside the observed range, and a coastal cell receives the mean of its ocean part.
- **Extension onto land.** Missing values next to the sea were filled from the neighbouring sea values, by a distance-weighted mean, for three cells inland. This gives covariate values to model cells whose centre lies on land although a large part of the cell is sea. The filled values are extrapolated and have no meaning on land itself.
- **Limiting the range.** Temperature was limited to $23.5\text{--}30^\circ\text{C}$ and mixed layer depth to 10–50 m, the ranges experienced by the tags; values outside were set to the limit.
- **Currents.** The current components were averaged over the model cells and converted to km per month along the grid axes.

Monthly maps of the fields for 2007 are shown in Appendix A.

The model grid consists of 95 cells of 500 km (Figure 1). The cell size is a compromise: the cost of the CTMC engine increases steeply with the number of cells (Section 2), and the recaptures are too sparse over most of the domain to inform finer-scale structure. The sensitivity of the results to the cell size is shown in Section 5.4.

4.3 Fishing effort

Fishing effort is not yet part of the model, but it was prepared from the IOTC catch-and-effort data (IOTC Secretariat, 2023) for the analysis period and is used here to interpret the results.

Table 3 lists the gears that recaptured the tags in the model. Purse seine accounts for 69% of the recaptures and pole and line for a further 19%; gillnet, handline, longline and troll together account for about 12%. This report further explores purse-seine effort only. The effort of the other gears is left for the final report (Section 6).

Gear	Recaptures	%
Purse seine	2341	69.1
Pole and line	637	18.8
Gillnet	152	4.5
Handline	120	3.5
Longline	77	2.3
Troll	45	1.3
Unknown	14	0.4
Rod & reel	2	0.1
Other	1	0.0
Total	3389	100.0

Table 3: Recaptures of the tags in the model by gear.

The effort data are not a single homogeneous data set, and combining them is a task of its own. For 2003–2015 they hold about 225000 records from 6 fishery groups, 12 fisheries, 27 gears and 33 fleets, which combine with the school type to 111 strata. Three properties matter for their use in the tagging model.

- **Effort units.** Effort is reported in ten different units (for example fishing hours, hours, fishing days, days, sets, trips, hooks and boats), and the unit differs between fleets of the same gear. Effort in different units cannot be added, and each unit needs its own catchability. Some fleets report the same effort in several units: the main purse-seine fleets give it as fishing hours, hours and standardised hours, so one unit has to be selected to avoid counting the effort three times.
- **Spatial resolution.** Purse-seine and pole-and-line effort is reported almost entirely on a 1° grid, longline effort mostly on a 5° grid, and gillnet and line fisheries on a mixture of both. Part of the effort of every fishery group is reported for irregular areas without a position.
- **School type.** Purse-seine effort of the main fleets is not divided into sets on free schools and sets on floating objects, although their catch is. Different catchabilities for the two school types can therefore not be estimated from the effort data alone. The school type is also missing on the tag side: 1194 of the 2341 purse-seine recaptures in the model (51%) have no recorded school type, against 35% of the purse-seine returns used in the 2024 assessment. The assessment assigns these returns by size: fish larger than 80 cm at release to the free-school fishery, and smaller fish to the floating-object fishery if recaptured within 18 months and to the free-school fishery otherwise (Urtizberea et al., 2024).

Purse-seine effort further explored here. As most tags (69%; Table 3) are recaptured by purse seiners, we started further exploring this fishery. The data contain 14 purse-seine fleets. As a start, the six fleets that report effort in fishing hours on a 1° grid are considered: EU (Spain), EU (France), EU (France, Mayotte), Seychelles, I. R. Iran and the fleets grouped as not elsewhere included (NEI). Together they report 1.76 million fishing hours over the analysis period, of which Spain, Seychelles and France account for 93%. The effort of these fleets is summed by 1° cell and month, over school types (Figure 4). The effort is concentrated in the western equatorial Indian Ocean between the African coast and about 70°E , with a band extending eastwards along the equator, and its distribution changes over the year (Figure 5): it lies north of the equator off Somalia from August to October, moves south and east towards the end of the year, and extends into the Mozambique Channel from March to May.

The eight purse-seine fleets that are left out report effort in days, fishing days, sets or trips, on 1° or 5° cells or for irregular areas. The largest of them is the Thai fleet, whose 1.6 million fishing hours are all assigned to a single irregular area and cannot be placed on the grid.




Figure 4: Purse-seine effort (fishing hours) of the six fleets used here, summed over the analysis period by 1° cell. Colours are on a logarithmic scale.




Figure 5: Purse-seine effort (fishing hours per month) of the six fleets used here by calendar month, averaged over the analysis period. Colours are on a logarithmic scale.

5 Initial results

5.1 Model configurations

In total, 8 configurations of the CTMC model were fitted to the 3389 tags (Table 4). The reference configuration (number 0) has a temperature preference with three knots (at 25.2, 26.5, 28.6°C), a constant diffusion rate and advection by currents. The other configurations change one element at a time: a diffusion rate that depends on temperature (1), mixed layer depth in place of temperature (6), a diffusion rate that depends on position (12), and no currents (14). Three configurations test the preference function itself: a null model without taxis (16), and the reference configuration with the temperature field mirrored north–south (17) or east–west (18). A mirrored field keeps the values and the spatial structure of the temperature field but places them elsewhere, so it has no relation to the temperature a fish experiences. All configurations use monthly time steps and one set of parameters for the whole period.

8 of the 8 fits passed all convergence checks. Each fit ran on a single core of a computing server and took at most 1.4 hours and 54 GB of memory (Section 5.4).

No.	Taxis / diffusion	p	nll	Δ AIC	rmse	50%	95%
0	temp (3) / constant	4	10998.2	908	926	0.66	0.94
1	temp (3) / temp (2)	5	10822.5	559	928	0.65	0.95
6	mld (3) / constant	4	10924.0	759	916	0.66	0.93
12	temp (3) / x+y (3)	8	10540.3	0	911	0.55	0.92
14	temp (3) / constant, no currents	3	11418.5	1746	987	0.66	0.94
16	no taxis / constant	2	11134.8	1177	923	0.67	0.93
17	temp (3) / constant, placebo ns	4	11116.9	1145	924	0.67	0.93
18	temp (3) / constant, placebo ew	4	10955.9	823	928	0.66	0.94

Table 4: Model configurations. Taxis / diffusion: covariates of the preference function and of the diffusion rate, with the number of knots in brackets (temp = temperature over the upper 75 m, mld = mixed layer depth, x+y = position). All include currents unless stated. p : number of estimated parameters; nll: negative log-likelihood; Δ AIC: difference to the lowest AIC among the converged fits to the same tags; rmse: root mean squared error of the predicted recapture position (km); 50% and 95%: share of recaptures within the 50% and 95% prediction regions (Section 2.5).

5.2 Reference configuration

Figure 6 shows the reference fit. The estimated preference increases with temperature over the whole range experienced by the tags, more steeply below 27 °C than above. Compared with the null model without taxis, the temperature preference improves the negative log-likelihood by 137 units for two parameters. The diffusion rate is estimated at about 118000 km² per month, which corresponds to a standard deviation of the displacement of about 490 km per coordinate after one month. The entrainment coefficient is 0.31 (95% confidence interval 0.29–0.32): the displacement of the fish is consistent with about a third of the velocity of the current over the upper 75 m.




Figure 6: Reference configuration. A: estimated temperature preference with 95% confidence band (dots: knots). B: taxis (arrows and colour: direction and strength). C: advection by currents. D: diffusion rate (constant). E: parameter estimates with 95% confidence intervals.

5.3 What is robust and what is not

Further results of the reference configuration and of other configurations are shown in Appendix B.

Currents. The effect of currents is the most robust result. Without currents the negative log-likelihood is 420 units higher for one parameter fewer, the prediction error increases from 926 to 987 km, and the estimated diffusion rate increases slightly (from 118000 to 128000 km² per month), because displacement that the currents explain is otherwise attributed to random movement. The entrainment coefficient lies between 0.23 and 0.35 in all configurations with currents.

Temperature preference. The temperature preference is not robust. When the diffusion rate is allowed to depend on temperature (configuration 1), the fit improves by 176 log-likelihood units for one additional parameter, the preference at the warmest knot drops from 3.3 to 2.7, and the preference levels off above about 27.5 °C (Figure 7). At the same time the diffusion rate decreases by a factor of about 1.8 between 26 and 28 °C. Part of what the reference configuration attributes to a preference for warm water is therefore attributed to slower movement in warm water. In the movement model a lower diffusion rate in warm water keeps fish there longer, so they accumulate in warm water without a preference for it. The data do not distinguish between fish moving towards warm water and fish moving less in warm water: both keep recaptures in the warm western equatorial area.




Figure 7: Configuration 1: estimated temperature preference (left) and log-diffusion rate as a function of temperature (right), with 95% confidence bands. The marks at the bottom are the temperatures at the release and recapture positions.

Placebo fields. The placebo configurations show that the preference function does not need to describe a response to temperature to fit the data. With the temperature field mirrored north–south the fit is 119 units worse than with the real field, but with the field mirrored east–west it is 42 units *better* (Figure 8), and better than every configuration with a constant diffusion rate in which temperature is the only environmental field. A field that has no relation to the temperature the fish experience explains the recaptures better than the real temperature. The preference function is therefore fitting the spatial structure of the recapture positions, which any field with a suitable gradient can reproduce. On a finer grid, however, the real temperature field fits better than the mirrored one (Section 5.4).




Figure 8: Configuration 18: estimated preference for the temperature field mirrored east–west, with 95% confidence band. The marks at the bottom are the values of the mirrored field at the release and recapture positions.

Mixed layer depth. With mixed layer depth in place of temperature (configuration 6) the fit is 74 units better than with temperature. The preference is highest for the shallowest mixed

layers and decreases steadily with depth (Appendix B). On a finer grid the order is reversed (Section 5.4). Given the placebo results, this shows that mixed layer depth describes the recapture pattern, not that it drives the movement.

Position explains more than the environment. The lowest AIC belongs to a configuration in which the diffusion rate depends on position and not on an environmental field (configuration 12). Its diffusion rates vary by a factor of about 35 over the domain, from 10000 to 369000 km² per month, low in the south and high in the north-west, and the temperature preference all but disappears. This is not a plausible description of behaviour. It is a description of where recaptures occur.

Predictions. Better likelihood does not translate into better prediction of the recapture position: the prediction error is between 911 and 928 km for all configurations with currents. The predicted distributions are too wide, with 55–67% of the recaptures inside the 50% prediction regions.

5.4 Sensitivity to the grid size

All results above are for the grid of 500 km cells. To test how much they depend on the cell size, the configurations were refitted on a grid of 250 km cells (359 cells instead of 95), with the same tags and settings (Table 5). The negative log-likelihood is the probability of the recapture cell and therefore depends on the cell size; it can only be compared between configurations on the same grid. On the predictive density per km², the reference configuration fits about 460 log-likelihood units better on the finer grid.

Some results hold on both grids. The currents improve the fit by 420 and 266 units, with an entrainment coefficient of 0.31 and 0.35, and the preference function improves on the null model by 137 and 122 units, while the field mirrored north–south does not. The diffusion rate is lower on the finer grid (87000 against 118000 km² per month), because part of the movement that the coarse grid cannot resolve is attributed to diffusion. On both grids a diffusion rate that depends on temperature improves the fit (by 176 and 272 units), and the configuration with position-dependent diffusion has the lowest AIC (458 and 705 units better than the reference).

Other results change. On the finer grid the real temperature field fits better than the field mirrored east–west (by 16 units) and better than mixed layer depth (by 100 units), the reverse of the order on the coarse grid. The temperature preference no longer increases over the whole range: it is highest at the middle knot, around 26.5 °C (1.6, against 0.8 at the warmest knot; 2.0 and 3.3 on the coarse grid). With a diffusion rate that depends on temperature, the preference on the finer grid is weak and decreases above the middle knot (0.8 at the middle and –0.3 at the warmest knot). With position-dependent diffusion the entrainment coefficient drops to 0.16 on the finer grid, so its size also depends on how diffusion is modelled. The comparison between environmental fields and the shape of the preference function therefore depend on the grid, which is a further reason not to interpret them before fishing effort is included.

The finer grid is computationally much more demanding. The number of terms of the uniformisation grows with the largest exit rate, which scales with D/Δ^2 , and the number of cells grows with $1/\Delta^2$. On the finer grid the fits with a constant diffusion rate took up to 3.1 hours and 230–253 GB of memory, against at most 1.4 hours and 54 GB on the coarse grid. The two configurations with a diffusion rate that varies in space (1 and 12) needed 480–586 GB and up to 11 hours, because the highest diffusion rate in the domain sets the number of uniformisation terms for all cells.

No.	Configuration	Δnl		Memory (GB)		Run time (h)	
		500 km	250 km	500 km	250 km	500 km	250 km
0	reference	0	0	41	249	0.8	3.1
1	D temperature	−176	−272	54	480	1.2	6.1
6	mixed layer depth	−74	100	38	232	0.8	2.5
12	D position	−458	−705	54	586	1.4	10.8
14	no currents	420	266	41	253	0.7	3.0
16	no taxis (null)	137	122	36	232	0.6	1.8
17	placebo north–south	119	118	36	230	0.8	2.7
18	placebo east–west	−42	16	38	235	0.8	3.0

Table 5: Model configurations on grids of 500 and 250 km cells. Δnl : negative log-likelihood relative to the reference configuration on the same grid (negative: better than the reference). Memory: peak memory of the fit; run time: fit, standard errors and predictions on one core.

5.5 Recaptures and fishing effort

Figure 9 compares the observed recaptures of the fish released off Tanzania and Kenya with the distribution predicted by the reference configuration and with the distribution of purse-seine effort. In the first six months the fish are recaptured further from the release area than predicted (a median of about 1140 km against 770 km), along the purse-seine grounds to the north-east. After that the observed distance no longer increases (about 920 km at 12–24 months), while the model keeps spreading the fish over the basin (about 1800 km). The recaptures stop where the purse-seine effort stops. The same pattern is found in all configurations.




Figure 9: Recaptures (yellow squares) of tags released off Tanzania and Kenya (blue circles) by time at liberty, with the 50% and 90% regions of the distribution predicted by the reference configuration (dark and light teal) and purse-seine effort over the recapture period (grey).

These results indicate that the recapture positions are shaped by the distribution of fishing effort as much as by movement. A model that conditions on recapture and ignores effort reads the absence of recaptures outside the fishing grounds as absence of fish, and attributes it to whichever model component can reproduce it: a preference for the temperature or mixed layer depth of the fishing grounds, or low diffusion where the fleet does not operate. The same was found for skipjack in the EPO (Mildenberger et al., 2025), and the general problem that catches reflect where tunas are vulnerable to the gear as much as where they are has long been recognised (Brill, 1994). The estimated preference functions should therefore not be interpreted as habitat preference until fishing effort is part of the model. The estimates that do not depend on this are the overall rate of dispersal and the contribution of the currents.

These findings are consistent with what the assessment and the review report about the tagging data. The analyses of the mixing period found that the maximum displacement of the tags reached a plateau within a few weeks of release, and that the recoveries of the floating-object fishery were not mixed for at least six months (Kolody & Hoyle, 2013; Langley & Million, 2012; Urtizberea et al., 2024). Few tags were recovered outside the western equatorial region and none in region 3 of the assessment, and the review noted that the lack of recoveries in the east could be due either to low effort or to an absence of movement (Maunder et al., 2023). Both observations are what a model with fishing effort is needed to separate. Environmental indices have been tested in earlier assessments as drivers of the movement between assessment regions and were

found to have no influence (Fu et al., 2018; Langley, 2016), although the movement between the western and eastern equatorial regions has been related to the Indian Ocean Dipole (Langley et al., 2023). The review summarised some evidence that the movement of juveniles is related to ocean currents and that of adults to temperature (Maunder et al., 2023). For the mostly juvenile fish in the model, the currents are the most robust result and the temperature preference is not, in line with this.

All estimates refer to the size range of the tagged fish, mostly between 41 and 72 cm fork length at release (Figure 3). Movement is likely to change with size: larger fish swim faster, which affects the diffusion rate, are less likely to be carried by currents, which affects the entrainment coefficient, and tolerate a wider range of temperatures and depths, which affects the habitat preference; thermal and oxygen tolerances change with body size, so physiological results for small tunas cannot simply be extrapolated to larger fish (Brill, 1994). The estimates should therefore not be transferred to large adult yellowfin.

6 Final steps for the movement model

Part I ends with a report at the end of 2026 that contains the movement model, its evaluation and recommendations on further work with the tagging model. The remaining steps are listed in the order in which they are planned (Table 6).

Period	Task
October–November	Fishing effort and mortality in the CTMC likelihood; effort fields by fleet; tag loss and reporting rates
November	Finer grid; recaptures with candidate positions; sensitivity runs
November–December	Model evaluation: residuals, cross-validation, comparison of configurations
December	Part I report with recommendations

Table 6: Remaining steps of Part I.

- **Fishing effort and mortality.** The main remaining task is to include fishing effort in the likelihood. The CTMC engine extends naturally: capture and natural death are added as states that a fish can enter but not leave, with a capture rate in each cell and month that is proportional to the fishing effort there. The likelihood then no longer conditions on recapture, so that the position and time of a recapture are evaluated against where effort was, and the releases that were never recaptured contribute information. This was developed and applied in the EPO analyses (Mildenberger et al., 2023) and is currently being implemented and tested in *admove*. Because all releases enter this likelihood, recaptures with incomplete information cannot simply be dropped: a dropped recapture counts as a fish that was never caught, which biases fishing mortality downwards and, if the missing information is concentrated in some areas or periods, also the movement estimates. Such recaptures are kept and their likelihood is summed over the possibilities that are consistent with the record, for example over the candidate positions of a recapture from purse-seine sets, or over the fleets or school types that could have made a recapture whose fleet or school type is unknown. For the school type, the size-based assignment of the 2024 assessment (Section 4.3) is an alternative. With capture by each fleet as a separate state of the Markov chain, this sum is the probability of the union of these states. Until effort is in the likelihood, the interpretation of any spatially structured term in the model remains provisional. Including the releases that were never recaptured increases the computational demands. The cost does not grow with the number of tags, because all tags released in the same cell and month share one propagation of the distribution, but with the length of time over which each release has to be propagated. A recaptured tag only needs to be followed to its recapture, whereas a tag that was never recaptured remains recoverable until the end

of the effort data or some specified end (e.g. maximum age).

- **Fleets and gears.** Each fleet or group of fleets in the model needs its own effort field and catchability. The work starts with the six purse-seine fleets described in Section 4.3, and is extended in three directions. (i) The remaining purse-seine fleets report effort in other units or at a coarser resolution; including them means combining effort units and spatial resolutions in one model, which needs testing. (ii) Sets on free schools and on floating objects are likely to differ in catchability, but the effort is not reported by school type, so a separation has to come from other information, for example the catch by school type. In addition, about half of the purse-seine recaptures have no recorded school type (Section 4.3). Fitted by school type, these recaptures would enter with the sum of the capture probabilities of both school types, as described above, rather than being dropped. This is valid if whether the school type was recorded does not depend on the school type itself; if, for example, recaptures from sets on floating objects were recorded more completely, the share of recorded school types would have to be modelled as well. (iii) The other gears recaptured about 30% of the tags (Table 3). Pole and line is the most important of them, and its effort raises issues of its own: the effort of the Maldives changes unit in 2013, from trips to fishing days, so the two periods cannot share one catchability, and the trips before 2013 are assigned to twelve fixed cells along the atolls, which appears to be an allocation and not reported positions. Gillnet, handline, longline and troll effort is reported in coarse or irregular areas and in mixed units. For each gear it has to be decided whether it is represented at its coarse resolution or whether the affected tags are left out. The IOTC data distinguish far more fleets than can be given separate catchability parameters, so fleets have to be pooled into a small number of groups. We would welcome guidance from the Working Party on this grouping.
- **Reporting rates and tag loss.** Reporting rate, tag shedding and tag-induced mortality scale the relationship between tags at liberty and tags recovered and are confounded with fishing mortality. The 2024 assessment (Urtizberea et al., 2024) uses the following values, which are the starting point here.
 - *Initial tag mortality:* 27.5% of the releases, a 20% effect of tagging covariates estimated for the RTTP-IO (Hoyle et al., 2015) added to a base rate of 7.5% adopted from the western and central Pacific. The review notes that this value, raised from 10% in earlier assessments, has a large effect on the estimated fishing mortality (Maunder et al., 2023).
 - *Chronic tag loss:* about 20% after 2000 days at liberty, entered as an annual rate of 0.03, from an update of Gaertner and Hallier (2015).
 - *Reporting rate, purse seine:* tag-seeding experiments gave reporting rates of 57% in 2005, 89% in 2006 and 94% in 2007, the increase being due to awareness campaigns (Hillary & Areso, 2008; Hillary et al., 2008). Tags recovered at sea are assumed to be fully reported; tags recovered from Seychelles landings are corrected by the seeding estimates and by the share of the EU purse-seine catch landed in the Seychelles (Fu, 2020; Kolody et al., 2011).
 - *Reporting rate, other gears:* there are no direct estimates. The assessment estimates them, and they are estimated to be very low for longline.

The review recommended that these rates be estimated from data collected during the tagging programme, such as double tagging, tag seeding or pen experiments (Maunder et al., 2023). The double-tagged recaptures in the extract (Section 4) allow tag shedding to be estimated within the model. The database also records where each tag was found; 92% of the recaptures in the model are recorded as found on the fishing boat, and how this

corresponds to the distinction between at-sea and landing recoveries used in the assessment remains to be checked. In the current likelihood, which conditions on the recapture and its time, a constant reporting rate, tag shedding and initial tag mortality cancel out. Reporting rates that differ between gears, areas or years do not: they change the spatial pattern of the reported recaptures, for example through the increase in purse-seine reporting from 2005 to 2007, or a lower reporting rate in the longline or pole-and-line fisheries. The rates will be fixed at the values above or estimated where the data allow it, and varied in sensitivity runs.

- **Recaptures with candidate positions.** The 6531 recaptures known only from candidate purse-seine sets would triple the number of recaptures in the model. With effort in the model, each candidate set can be weighted by its probability, and the model already allows several alternative recapture positions per tag. This requires the candidate positions and dates of each recapture, which are not part of the current extract.
- **Finer grid.** With 500 km cells the model cannot resolve structure at the scale of the fishing grounds, and some results change on a grid of 250 km (Section 5.4). The aim is to use the finest grid that is computationally feasible for the full data set with effort in the model, which depends on keeping memory use and the largest diffusion rate on the grid under control. A finer grid also reduces the share of cells in which the central-difference rates become negative.
- **Size-dependent movement.** The length at release, together with a growth curve, gives an approximate length of each fish at any time at liberty. The diffusion rate, the entrainment coefficient and the preference function can then be made to depend on length, for example through length as an additional covariate or separate parameters for length classes, as in the length-structured model developed for skipjack in the EPO (Mildenberger et al., 2023). With most tags released on juvenile fish, the data may support only a coarse size dependence; this is to be explored after the effort module.
- **Model evaluation.** The final evaluation will use the residual diagnostics and the comparison of predicted and observed recaptures shown here, extended by cross-validation in which release events are left out of the fit and predicted. Sensitivity runs will cover the time at liberty, the analysis period, the RTTP-IO tags only (the selection of the 2024 assessment), seasonal preference functions and the fixed tag loss and reporting rates.

Part II extends the framework to mortality rates and a relative biomass index during 2027, with a scientific manuscript and the Part II report at the end of that year.

7 Acknowledgments

This work was carried out under a Letter of Agreement between the Food and Agriculture Organization of the United Nations (FAO) and the Technical University of Denmark, funded by the European Union. The tag-recapture data used for the analysis were collected under the Regional Tuna Tagging Project – Indian Ocean, funded by the European Union and implemented by the Indian Ocean Tuna Commission.

Generative AI tools (Anthropic’s Claude; opus-5-5; accessed 2026) were used to assist with drafting and editing text in this document. All analytical choices, model specifications, diagnostic interpretations, and conclusions are those of the authors, who reviewed and verified all AI-assisted output. The methodology, model implementation, and the scientific conclusions are the responsibility of the authors.

References

- Brill, R. W. (1994). A review of temperature and oxygen tolerance studies of tunas pertinent to fisheries oceanography, movement models and stock assessments. *Fisheries Oceanography*, 3(3), 204–216.
- Dagorn, L., Holland, K. N., Hallier, J.-P., Taquet, M., Moreno, G., Sancho, G., Itano, D. G., Aumeeruddy, R., Girard, C., Million, J., & Fonteneau, A. (2006). Deep diving behavior observed in yellowfin tuna (*Thunnus albacares*). *Aquatic Living Resources*, 19(1), 85–88.
- E.U. Copernicus Marine Service Information. (2024). Global ocean physics reanalysis (GLO-RYS12V1, product GLOBAL_MULTIYEAR_PHY_001_030).
- Fu, D. (2020). *Tag data processing for IOTC tropical tuna assessments* (IOTC Document No. IOTC-2020-WPTT22(DP)-10). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Fu, D., Langley, A. D., Merino, G., & Urtizberea, A. (2018). *Preliminary Indian Ocean yellowfin tuna stock assessment 1950–2017 (Stock Synthesis)* (IOTC Document No. IOTC-2018-WPTT20-33). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Gaertner, D., & Hallier, J.-P. (2015). Tag shedding by tropical tunas in the Indian Ocean and other factors affecting the shedding rate. *Fisheries Research*, 163, 98–105.
- Grassmann, W. K. (1977). Transient solutions in Markovian queueing systems. *Computers & Operations Research*, 4(1), 47–53.
- Hallier, J.-P., & Million, J. (2009). *The contribution of the Regional Tuna Tagging Project - Indian Ocean to IOTC stock assessment* (IOTC Document No. IOTC-2009-WPTT-24). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Harvey, A. C. (1990). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Hillary, R. M., & Areso, J. J. (2008). *Reporting rate analyses for recaptures from Seychelles port for yellowfin, bigeye and skipjack tuna* (IOTC Document No. IOTC-2008-WPTT-18). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Hillary, R. M., Million, J., Anganuzzi, A., & Areso, J. J. (2008). *Tag shedding and reporting rate estimates for Indian Ocean tuna using double-tagging and tag-seeding experiments* (IOTC Document No. IOTC-2008-WPTDA-04). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Hoyle, S. D., Leroy, B. M., Nicol, S. J., & Hampton, W. J. (2015). Covariates of release mortality and tag loss in large-scale tuna tagging experiments. *Fisheries Research*, 163, 106–118.
- IOTC Secretariat. (2023). *Catch and effort data for IOTC fisheries* (IOTC Dataset No. IOTC-2023-WPTT25(AS)-DATA07-CEAll_Rev2). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Jensen, A. (1953). Markoff chains as an aid in the study of Markoff processes. *Scandinavian Actuarial Journal*, 1953(sup1), 87–91.
- Kolody, D., Herrera, M., & Million, J. (2011). *Indian Ocean skipjack tuna stock assessment 1950–2009 (Stock Synthesis)* (IOTC Document No. IOTC-2011-WPTT13-31). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Kolody, D., & Hoyle, S. D. (2013). *Evaluation of tag mixing assumptions for skipjack, yellowfin and bigeye tuna stock assessments in the Western Pacific and Indian Oceans* (WCPFC Document No. WCPFC-SC9-2013/SA-IP-11). Western and Central Pacific Fisheries Commission (WCPFC).
- Kristensen, K. (2026). *RTMB: R bindings for TMB* [R package version 1.7].
- Kristensen, K., Nielsen, A., Berg, C. W., Skaug, H., & Bell, B. M. (2016). TMB: Automatic differentiation and Laplace approximation. *Journal of Statistical Software*, 70(5), 1–21.
- Langley, A. D. (2016). *An update of the 2015 Indian Ocean yellowfin tuna stock assessment for 2016* (IOTC Document No. IOTC-2016-WPTT18-27). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.

- Langley, A. D., Fu, D., & Maunder, M. N. (2023). *An investigation of the recruitment dynamics of Indian Ocean yellowfin tuna* (IOTC Document No. IOTC-2023-WPTT25-12). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Langley, A. D., & Million, J. (2012). *Determining an appropriate tag mixing period for the Indian Ocean yellowfin tuna stock assessment* (IOTC Document No. IOTC-2012-WPTT14-31). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Maunder, M. N., Mente-Vera, C. V., Langley, A. D., & Howell, D. (2023). *Independent review of recent IOTC yellowfin tuna assessment* (IOTC Document No. IOTC-2023-WPTT25-13_Rev2). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.
- Mildenberger, T. K., Maunder, M. N., & Nielsen, A. (2026a). *admove: Mechanistic advection-taxis-diffusion movement inference from tagging data in R*.
- Mildenberger, T. K., Nielsen, A., Magnusson, A., Scutt Phillips, J., & Hamer, P. (2026b). *Skipjack tagging data preliminary analysis, review and recommendations for further development and application* (WCPFC Scientific Committee Document No. WCPFC-SC22-2026/SA-WP-03). Western and Central Pacific Fisheries Commission (WCPFC). Apia, Samoa.
- Mildenberger, T. K., Nielsen, A., & Maunder, M. N. (2022). *Spatiotemporal tagging model for skipjack tuna in the EPO* (IATTC Document No. SAC-13-08). Inter-American Tropical Tuna Commission (IATTC).
- Mildenberger, T. K., Nielsen, A., & Maunder, M. N. (2023). *Length-structured spatiotemporal tagging model for skipjack tuna in the EPO* (IATTC Document No. SAC-14-INF-E). Inter-American Tropical Tuna Commission (IATTC). La Jolla, California.
- Mildenberger, T. K., Nielsen, A., & Maunder, M. N. (2024). *A spatiotemporal petersen-type model for skipjack tuna in the EPO* (IATTC Document No. SAC-15-INF-G). Inter-American Tropical Tuna Commission (IATTC). La Jolla, California.
- Mildenberger, T. K., Nielsen, A., & Maunder, M. N. (2025). *Ongoing development of the spatiotemporal tagging model* (IATTC Document No. SAC-16-INF-D). Inter-American Tropical Tuna Commission (IATTC). La Jolla, California.
- Sherlock, C. (2021). Direct statistical inference for finite Markov jump processes via the matrix exponential. *Computational Statistics*, 36(4), 2863–2887.
- Sibert, J. R., Hampton, J., Fournier, D. A., & Bills, P. J. (1999). An advection-diffusion-reaction model for the estimation of fish movement parameters from tagging data, with application to skipjack tuna (*Katsuwonus pelamis*). *Canadian Journal of Fisheries and Aquatic Sciences*, 56(6), 925–938.
- Thorson, J. T., Barbeaux, S. J., Goethel, D. R., Kearney, K. A., Laman, E. A., Nielsen, J. K., Siskey, M. R., Siwicke, K., & Thompson, G. G. (2021). Estimating fine-scale movement rates and habitat preferences using multiple data sources. *Fish and Fisheries*, 22(6), 1359–1376.
- Urtizberea, A., Correa, G. M., Langley, A. D., Merino, G., Fu, D., Chassot, E., & Adam, S. (2024). *Stock assessment of yellowfin tuna in the Indian Ocean for 2024* (IOTC Document No. IOTC-2024-WPTT26-11). Indian Ocean Tuna Commission (IOTC). Victoria, Seychelles.

A Environmental fields

The figures in this appendix show the environmental fields for 2007, a year with many releases and recaptures. Temperature, mixed layer depth and currents are shown as the model uses them: projected, extended onto land and limited to the range experienced by the tags (Section 4.2); the limits explain the uniform areas in the south of the domain. Sea surface temperature is not used in the model and is shown for comparison, as projected but not extended or limited.




Figure 10: Sea surface temperature (°C) by month in 2007.




Figure 11: Temperature averaged over the upper 75 m ($^{\circ}\text{C}$) by month in 2007, limited to 23.5–30 $^{\circ}\text{C}$.




Figure 12: Mixed layer depth (m) by month in 2007, limited to 10–50 m.




Figure 13: Currents averaged over the upper 75 m by month in 2007, converted to km per month along the axes of the projected grid. Arrows: direction; colour: speed.

B Further results

Reference configuration by month. Figures 14 and 15 show the directed movement of the reference configuration (configuration 0) by month in 2007: the taxis implied by the temperature preference, and the advection by the currents, that is the current velocity multiplied by the estimated entrainment coefficient. The diffusion rate of the reference configuration is constant and is not shown by month.




Figure 14: Configuration 0: taxis by month in 2007 (km per month). Arrows: direction; colour: speed.




Figure 15: Configuration 0: advection by currents by month in 2007 (km per month). Arrows: direction; colour: speed.

Mixed layer depth. With mixed layer depth as the covariate of the preference function (configuration 6), the preference is highest for the shallowest mixed layers and decreases steadily for deeper mixed layers (Figure 16).




Figure 16: Configuration 6: preference as a function of mixed layer depth, with 95% confidence band. The marks at the bottom are the values at the release and recapture positions.

Position-dependent diffusion. Figure 17 shows the diffusion rate of configuration 12, which has the lowest AIC of all configurations. The diffusion rate is highest in the north-west, off Somalia and in the Arabian Sea, and lowest in the south, where there are few recaptures. With position-dependent diffusion the temperature preference nearly disappears (0.0 against 3.3 at the warmest knot).




Figure 17: Configuration 12: diffusion rate (km^2 per month) by month in 2007. Colour and circle size: diffusion rate. The rate depends on position only and is the same in every month.