

An introduction to FLR fisheries simulation tools

Richard Hillary^a

Division of Biology, Imperial College, Ascot, SL5 7PL, UK

Received 27 August 2008; Accepted 8 April 2009

Abstract – Fisheries libraries in R (FLR) is an open-source software framework based in the R statistical software environment, with the main purpose of allowing for the efficient simulation and evaluation of fisheries management strategies. This paper details a generic (non-spatial) biological operating model used to simulate populations and their reaction to fishing, and an observation error model and methods for simulating structured sampling/observation error and bias in key fisheries assessment indices. It also contains the details of how the candidate harvest control rules, based on survey-derived information, can be characterised and implemented, and how the processes of management strategy evaluation and stock assessment can be included in the simulations.

Key words: Management strategy evaluation / Fishery-independent data / Observation error model / Harvest control rules / FLR

1 Introduction

The concept of management strategy evaluation (MSE), where one assesses the robustness of a given management procedure to the uncertainties and potential biases within the management system of interest, is a powerful one and is also an issue of current research interest for a wide range of different fisheries (Butterworth et al. 1997; Butterworth and Punt 1999; Punt and Smith 1999; De Oliveira and Butterworth 2004; Campbell and Dowling, 2005). Some of the evaluated management procedures (De Oliveira and Butterworth 2004) have been based upon information from surveys, not from more common fishery-dependent/assessment-derived indices. In this paper, I present a set of tools for simulating fisheries management systems, with particular attention given to parameterising the biological operating model, the simulation of structured error and bias in fisheries survey information, and the construction of candidate harvest control rules using only survey-derived information. One can view the MSE system as seen in the flow diagram (Fig. 1):

- An operating model (OM) simulates the biological and fishery “reality”, so both the dynamics of the stock and the dynamic impacts of the fishery upon the stock.
- Through observations of the system (catch, surveys, environmental/ecosystem information) we assess the status of the stock (assessment, with or without an actual assessment model of some kind) and this status information is combined with our pre-defined management options (harvest control rule or HCR) to form the management procedure (MP).

- This then effects some management action (with or without some implementation error and/or bias) which closes the loop as the implemented management action becomes fishery impact action once more.

All of the previous MSE studies do have one thing in common: they are all done using custom-designed software tools in a variety of differing languages and environments. In this paper, I present some basic tools (methodological and software-based in nature) based in the FLR simulation environment (Kell et al. 2007; <http://flr-project.org/>) specifically for use when we are using survey-based information (both raw indices or with variables derived from survey-based assessment methods) to construct candidate management procedures.

Fisheries library in R (FLR) is a collection tools, written in the open-source R programming language, that are designed to facilitate simulations of fisheries, ecological systems, and management strategies, taking into account the uncertainties that arise because many fisheries can only be monitored with significant levels of sampling error. The FLR project began as the framework for the simulation work of a variety of EU fisheries research projects (EFIMAS 2007; COBECOS 2008) but has progressed to become a key ICES (International Council for the Exploration of the Sea) stock assessment tool as well as a key research tool in a variety of topics and fora (Anonymous 2008; Hoff and Frost 2008; Needle 2008; SC-CAMLR 2008; Hillary 2009; Nash et al. 2009).

FLR benefits from being an open-source framework because this allows users to adjust programs for their own specific purposes and to pass on these developments to others. FLR also makes full use of R graphics to display simulation results in ways that are both communicable to managers and

^a Corresponding author: r.hillary@imperial.ac.uk

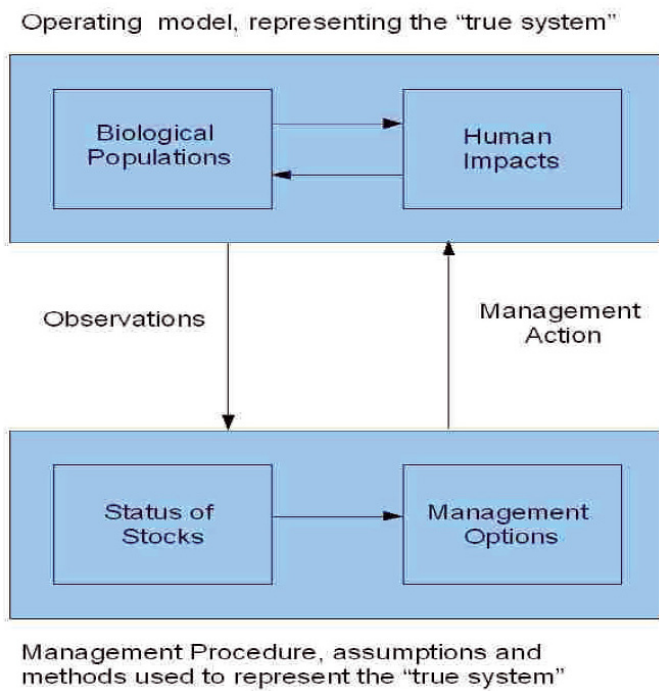


Fig. 1. Schematic flow diagram for the framework of management strategy evaluation.

stake-holders, yet technically informative for model construction and understanding. The present document provides:

- The mathematical specifications of the biological operating model, along with details of a worked example of how to parameterise it with standard ICES working group data;
- The mathematical specifications of the observation error model with details of how one can simulate various different surveys and, crucially, observation error models;
- How one can define candidate harvest control rules based upon survey-specific data or stock assessment information, to then be used to define survey-based management strategies;
- A review of the visualisation tools available with particular attention given to dealing with probabilistic information, and how to relate such information to relevant stakeholders.

2 The biological operating model

The FISBOAT project has a specific biological operating model, which differs substantially from the population dynamics model applied in the many virtual population analysis (VPA)-type stock assessment models used in the ICES area. This is, in part, due to the nature of some of the case-studies in the project (some lack catch-at-age data) and also so as to be able to generate a more sensible model of the case-study stocks’ dynamics. In the wider sense, it represents a fairly simple single-species non-spatial but seasonal biological operating model. Here is the basis behind the biological OM:

- Time-steps are yearly and “seasonal” – multiple within-year periods allowed;

- Harvesting model is single-fleet: defined by a supplied selectivity; effected via harvest rates, defined by the ratio of catch to exploitable biomass;
- Three modes of stock-recruitment allowed: stochastic stock-recruitment relationship (Beverton-Holt, Ricker, Hockey stick); stochastic (geometric) mean recruitment; bootstrap option, given a recruitment series.

2.1 Population dynamics

For the initial numbers, the model assumes that the population is at exploited/unexploited population equilibrium (i.e. at a point at which the numbers-at-age will remain the same without a change in the system), where the initial equilibrium harvest rates are defined by $\bar{h}_{a,s}$ at age a , in season s . This allows us to run scenarios where the initial depletion (relative to unfished conditions) can be pre-specified via the equilibrium harvest rates. If recruitment begins at a_r , in season s_r , then, for $s = 1, \dots, s_r - 1$, $N_{1,a_r,s} = 0$, and $N_{1,a_r,s_r} = R_0$, where R_0 is a known parameter. For $s = s_r + 1, \dots, S$, where S is the number of seasons:

$$N_{1,a_r,s} = N_{1,a_r,s-1} (1 - \bar{h}_{1,a_r,s-1}) e^{-M_{1,a_r,s-1}}, \quad (1)$$

where $M_{y,a,s}$ is the natural mortality. For ages $a = a_r + 1, \dots, A^+ - 1$, where A^+ is the plus-group, and $s = 1$ we have that

$$N_{1,a,1} = N_{1,a-1,s} (1 - \bar{h}_{a-1,s}) e^{-M_{1,a-1,s}} \quad (2)$$

and for $a = A^+$

$$N_{1,a,1} = N_{1,a-1,s} (1 - \bar{h}_{a-1,s}) e^{-M_{1,a-1,s}} + N_{1,a,s} (1 - \bar{h}_{a,s}) e^{M_{1,a,s}}, \quad (3)$$

and finally, for $s = 2, \dots, S$ and $a = a_r + 1, \dots, A^+$

$$N_{1,a,s} = N_{1,a,s-1} (1 - \bar{h}_{a,s-1}) e^{-M_{1,a,s-1}}. \quad (4)$$

This takes care of the initial population dynamics, and, for the remaining years $y = 2, \dots, Y$ the dynamics are defined as follows:

For the seasons before recruitment, $N_{y,a_r,s} = 0$; we will explain the stock-recruit process in more detail later on; for the seasons following recruitment, the dynamics for $a = a_r$ are

$$N_{y,a_r,s} = N_{y,a_r,s-1} (1 - h_{y,a_r,s-1}) e^{-M_{y,a_r,s-1}}. \quad (5)$$

For season one, the dynamics for $a = a_r + 1, \dots, A^+ - 1$ are

$$N_{y,a,1} = N_{y-1,a-1,s} (1 - h_{y-1,a-1,s}) e^{-M_{y-1,a-1,s}} \quad (6)$$

and for the plus group:

$$N_{y,a,1} = N_{y-1,a-1,s} (1 - h_{y-1,a-1,s}) e^{-M_{y-1,a-1,s}} + N_{y-1,a,s} (1 - h_{y-1,a,s}) e^{-M_{y-1,a,s}}, \quad (7)$$

and for ages $a = a_r + 1, \dots, A^+$, in seasons $s = 2, \dots, S$, we have

$$N_{y,a,s} = N_{y,a,s-1} (1 - h_{y,a,s-1}) e^{-M_{y,a,s-1}}. \quad (8)$$

The harvest rate, $h_{y,a,s}$ is defined to be

$$h_{y,a,s} = \kappa_{y,a,s} \times H_{y,s} \quad (9)$$

where $\kappa_{y,a,s}$ is the selectivity function, and $H_{y,s}$ is the ratio of catch to total exploitable stock biomass. Note that the subscript y for κ allows for changing selectivity over time but it should only be estimated if there are good catch and survey data. Otherwise, changes of selectivity could be confounded with varying annual recruitments. The spawning-stock biomass (SSB) is defined in the standard manner, but with the inclusion of mortality before spawning in each of the seasons.

2.2 Stock-recruit behaviour

If a stock recruit relationship is to be used, then the recruits are defined as follows:

$$N_{y,a_r,s_r} = SRR(\theta, SSB_{y-a_r,s_{sb}}) e^{\xi_y} \quad (10)$$

In Eq. (10), $SRR(\theta, \bullet)$ is the particular stock-recruit function (Beverton-Holt, Ricker, hockey-stick), and θ are the parameters. The delay used to relate spawners to recruits is implicitly defined by the minimum age in the model, and s_{sb} is the spawning season. Note also, for reasons of common sense, that if $a_r = 0$, then recruitment cannot occur before spawning. The final term in Eq. (10) is the stochastic recruitment multiplier, defined as follows:

$$\xi_y = \rho \xi_{y-1} + \xi_y \quad (11)$$

where $\xi_y \sim N(0, (1 - \rho^2)\sigma_r^2)$ and ρ is the auto-correlation coefficient for the recruitment multipliers. No bias correction term is applied as the assumption here is that the median, not the mean, recruitment is the deterministic recruitment prediction.

If fixed geometric mean recruitment, \hat{R} , is requested, then

$$N_{y,a,s_r} = \hat{R} \times e^{\xi_y} \quad (12)$$

If a recruitment time-series is supplied, then the model simply resamples with replacement from this recruitment time-series to generate the model recruitments, but with no stochastic error term applied - with multiple simulations, the bootstrapping procedure will eventually realise the recruitment uncertainty and temporal correlation structure.

2.3 Parameterising the biological operating model

The process of estimating the parameters of our operating model is termed “conditioning” (Rademeyer et al. 2007). The conditioning of an operating model can be done directly from a stock assessment model with a similar population and fishery structure such as statistically integrated models like CASAL (Bull et al. 2005) or using the sorts of data that come from ICES stock assessment working groups: recruitment, spawning stock biomass, fishing mortality and so on. However, the idea of conditioning is not as restrictive as assessment usually is – for example we may wish to define an operating model

with a more complex spatial and trophic/multi-species structure yet lack the information to statistically sensibly estimate all the key parameters. Conditioning merely states that we ensure that our operating model is consistent with the available information, and that our assumptions are plausible. The conditioning process then provides us with the population and fishery model with which to begin our management simulations.

Using North Sea herring as an example case, I demonstrate how to use FLR tools to condition the operating model defined in Sects. 2.1 and 2.2. North Sea herring is assessed using the integrated catch-at-age, ICA (Patterson and Melvin 1995) tuned-VPA algorithm. For this stock there are data on historic maturity, stock and catch-weight-at-age as well as natural mortality so that leaves us with selectivity and stock-recruit relationship as the key factors that we must estimate to full condition the operating model.

2.3.1 Estimating the stock-recruit relationship

The main core of FLR is an R package called FLCORE (Kell et al. 2007), and within the core are many objects used to store data and methods to estimate relationships between data. For the purposes of estimating stock-recruit relationships, we have the FLSR class in FLCORE. This class allows us to estimate the parameters of a variety of stock-recruit relationships and also to choose between different stock-recruit models using maximum likelihood techniques. For the herring case, I consider the Ricker (Ricker 1958) and hockey-stick/piecewise linear stock-recruit models. Assuming a lognormal likelihood function both models were fitted to the herring stock-recruit data. Both models have the same number of parameters, thus, the model selection criteria was simple: the model which minimises the Akaike information criterion (AIC) was the best. For the Ricker model, the value was 74.65 and for the hockey-stick model the value was 72.08 so, henceforth, the hockey-stick model parameters, residual autocorrelation (ρ) and corrected variance were used in the operating model. It should be noted that there is no solid biological justification, like that used to derive the Ricker model, for the hockey-stick model employed here. The choice of stock-recruit relationship is based purely on statistical grounds – the potential for model uncertainties such as this one should be evaluated within the realm of management strategy evaluation; see Figure 2 for the FLR summary plot obtained when fitting a stock-recruit relationship.

2.3.2 Estimating the selectivity curve

The ICA algorithm allows one to impose a set number of years before the final year for which a fixed (but not functional) selectivity curve can be estimated. For herring, this “separable period” is fixed to be five years and I use the final year fishing mortality-at-age (normalised to be between zero and one by dividing by the maximum fishing mortality F) to estimate selectivity-at-age for the future. Note this assumes full selection of at least one age-class in the population. After selecting

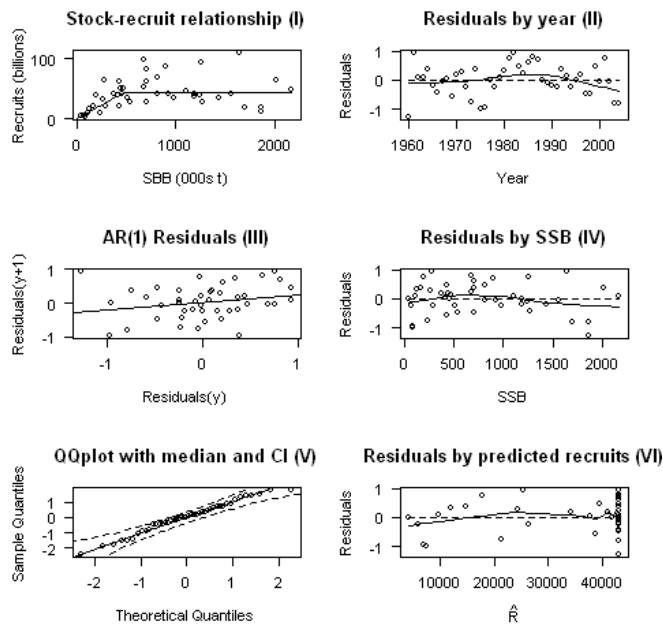


Fig. 2. FLR stock-recruit summary plot for hockey-stick model: (I) Fitted and observed relationship; (II) residuals-by-year; (III) auto-correlation trend (slope of line) in the residuals; (IV) residuals-by-SSB; (V) qq-plot of residuals and confidence intervals; (VI) residuals-by-predicted recruitment.

a logistic function to model selectivity-at-age, $s(a)$:

$$s(a) = (1 + \exp(\omega(a - a_{50})))^{-1},$$

where a_{50} is the age at 50% selectivity and ω is the shape parameter, I then use the `optim()` function in *R* to estimate a_{50} and ω given the herring selectivity information. Figure 3 shows the fit to the observed selection pattern.

The final part of the conditioning of the operating model is to define the historical dynamics for herring, and the simple and obvious answer is to simply impose the ICES working group estimates of numbers-at-age, harvest rate-at-age (derived from F) and SSB upon the historical dynamics, and use the selectivity and stock-recruit parameters (as well as the future catch) to define future dynamics.

2.4 Converting from age to length

The operating model simulates population indices-at-age, but it is easy to imagine harvest control rules and assessment methods that deal with length data, not age data – especially in a survey context. To this end, a conversion method was written into the FLFisboat package, called FLal, as the conversion from length-to-age (required to simulate a length-based survey from an age-based operating model) is as trivial as one might think.

Age-to-length conversion is the easier of the two “directions” - for a given growth model, and whatever the estimation error model and variance, the expected length, \bar{l} , for a given age, a , and a given growth model, $G()$, is as follows:

$$\bar{l} = G(a). \tag{13}$$

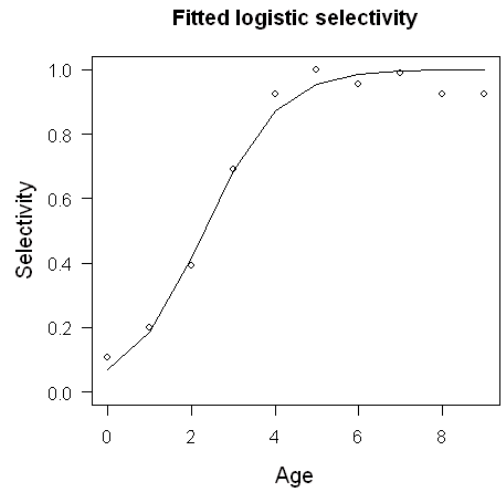


Fig. 3. Fitted (full line) versus stock assessment-derived “observed” selectivity (circles) for North Sea herring.

Given this expected length-at-age, we also must define a partition of lengths: l_1, \dots, l_N , for the length-based return object. Given this length partition, the conversion process simply finds the bin containing \bar{l} , for the given age, and the value of the data by age, year etc. is added to the length entry defined by the lower value of this length bin. If, for some reason, \bar{l} values fall outside the defined bins, then by default plus and minus groups are requested – a sensible choice of length bins, given the age range and growth model, should avoid the need for this option, and it can be switched off if required.

When converting from length to age, the fact that there is in fact a distribution of lengths – for any given age – must be taken into account, as I am, as yet, unaware of a perfectly estimated growth model. The methodology for this part of the conversion process is as follows:

We first assume that the inverse growth model, $G^{-1}()$, exists and can be calculated. For our given length, \hat{l} , what we do is to compute the expected age, given the inverse growth and error model:

$$\bar{a} = \int_L G^{-1}(l) \pi(l|\hat{l}) dl = E^{\pi(l|\hat{l})}[G^{-1}(l)] \tag{14}$$

and, for a normal error model with standard deviation σ , we have that

$$\pi(l|\hat{l}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(l - \hat{l})^2}{2\sigma^2}\right). \tag{15}$$

The integration step in Eq. (14) is numerically approximated. E means expectation (mean) over the distribution π . We allow for two error models: normal and lognormal. For the normal case, grid extrema are set-up at $\pm 2\sigma$ from \hat{l} (but if the lower is less than zero we set it to length zero) and an evenly spaced length grid of 100 points is defined between these two. For the lognormal case, we set these extrema at $\exp(\pm 2\sigma)$, so both encompass the 95 percentile of the error distribution being used. We then, after normalising, use this discrete approximation of $\pi(l|\hat{l})$ to compute the value of \bar{a} . If, for example, $\hat{l} > L_{\infty}$ for

the von Bertalanffy model, then \bar{a} is set at the maximum age, and, again, plus and minus groups act by default. The rounding up or down of the calculated value of \bar{a} is as one would expect, in that if $\bar{a} - \text{floor}(\bar{a}) < 0.5$, $\bar{a} = \text{floor}(\bar{a})$; if not, then $\bar{a} = \text{ceiling}(\bar{a})$.

2.4.1 Calculating mean length in the population

Our operating model has numbers-at-age at its core and in the FLR environment this exists as an FLQuant (Kell et al. 2007) – the key object in the FLCore data object structure. The FLal method just requires the FLQuant with numbers-at-age, $N_{y,a}$, a vector of length classes, $\lambda = \{l_1, \dots, l_p\}$, and the relevant growth curve parameters and error model. The convert() method then creates an FLal object with the numbers-at-length FLQuant, $N_{y,l}$, now calculated. In terms of mathematical expressions the mean length, $\mu(l)$, is defined as

$$\mu(l) = \sum_{k=l}^{l_p} \left(k \times \frac{N_{y,j}}{\sum_{j=l_1}^{l_p} N_{y,j}} \right) \quad (16)$$

The expression in Eq. (16) can be achieved using the pay() method from FLEDA (exploratory data analysis) and the quantSums() method in FLCore: $\mu(l) = \text{quantSums}(\lambda * \text{pay}(N_{q,i}))$.

3 The observation error model

The second main part of the software framework was to develop an observation error model, which could simulate the known types of survey observations commonly available and used in most stock assessment methods (conventional or survey-based in nature). The observation error model takes data objects from the biological operating model, and together with information on catchability, observation error level and structure and so on and simulates observations that can then be used in the survey-based stock assessment methods and/or in developing survey-based harvest control rules.

The observation error package is also an FLR-based package, and is called FLOE. While the genesis of this package was largely driven by the FISBOAT project the simulation of abundance indices with different error and bias models is clearly of relevance to fisheries assessment and management as a whole, not just in the context of fishery information-independent management. The FLOE package is split into two main parts:

1. FLObsIndex: the part of the package which simulates standard survey-type observations, such as the International beam trawl survey (IBTS)-type trawl surveys, acoustic surveys and recruitment surveys, under a wide range of observation error, catchability and bias regimes.

2. FLprop: this part of the package can be used in conjunction with observations from the FLObsIndex simulations or as a stand-alone method. Its purpose is to simulate correlated error in simulations (ageing error or year effects, for example).

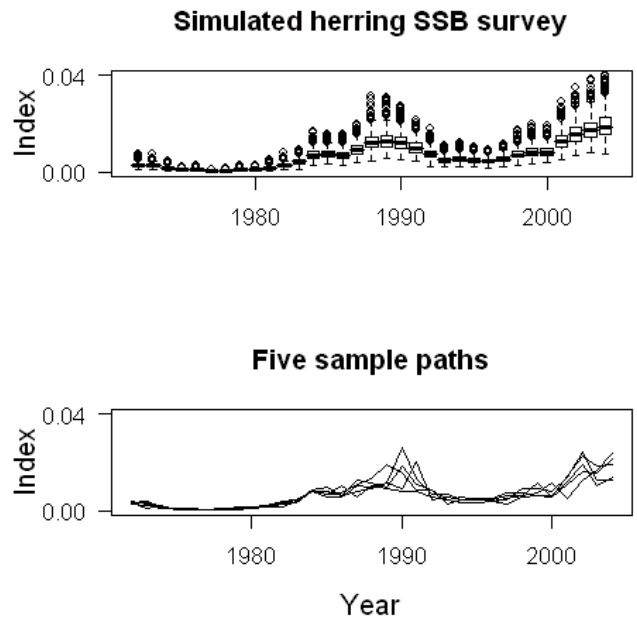


Fig. 4. Simulate North Sea herring egg survey (assuming an observation error coefficient of variation of 0.27 from observed variation): (top) a box and whisker plot of 1000 resamplings of the survey, given the observation error; (bottom) five indicative sample paths of these 1000 possible surveys indices.

3.1 The FLObsIndex suite of methods

The first thing to cover is the mathematical and statistical nature of the package. Speaking generically, we assume that any simulated index, \hat{I} , can be expressed in the following way:

$$\hat{I} = I^\beta \times q \times b (\times \text{ or } +) \varepsilon, \quad (17)$$

where I is the “true” population variable being observed; q is the catchability, which in truth could be a composite of a number of factors; b is the bias in the observations; and ε is the error term, which could be additive or multiplicative. The β parameter is the index-to-abundance power coefficient. The error term should be defined by the particular observation error model being applied and/or assumed. This basic structure is sufficient to generate a very large variety of observations – certainly enough to cover what is normally used as relative/absolute abundance tuning data. All the relevant parameters can vary by age, year and season, meaning that a very wide range of regimes can be simulated. Error structures permitted are normal (additive), log-normal and gamma (multiplicative), with a simple lag-1 correlation effect permitted in the normal and log-normal error regimes.

I used North Sea herring to parameterise the operating model, I now use herring to show how to simulate a survey of relative SSB using the FLR tools. For North Sea herring, a survey of relative egg production (MLAI) is done each year and is treated as a proxy for spawning stock biomass. To mimic this kind of observation, I set an arbitrary value of q of 10^{-5} (it is a relative index) and assume a lognormal error structure with a given coefficient of variation (CV), where the log-scale standard deviation is given by the following relationship: $\sigma_{\text{MLAI}} = \sqrt{\log(1 + CV^2)}$. Again to mimic the MLAI survey,

I simulate the survey from 1972 to the final year with a *CV* of 0.27 (which is a mean *CV* over the years of the survey). Figure 4 shows both a boxplot summary of the simulated SSB survey (given 500 Monte Carlo samples) and three sample paths of the SSB survey. The reason for displaying the information in this way is that people tend to look at a box plot and focus on the median trend and either the interquartile range or the extent of the whiskers but it is helpful to look at what actual survey indices might look like. The actual process of data collection, over time yields single paths and showing both summary information on the distribution of possible values with sample paths, is more informative than either in isolation.

3.2 The FLprop suite of methods

This part of the package was designed to simulate the more complex forms of error seen in fisheries data-at-age/length. These forms of error normally correlate across age or length – the prime example of course is ageing error – and are rarely, if ever, properly accounted for in all forms of stock assessment, be they fisheries dependent or independent. The idea is that, given proportions data, p , at the FLQuant (the basic data array object in the FLCore package) resolution, if we assumed that these data are multinomially distributed, then after applying the logit transformation to the data:

$$\hat{p} = \ln\left(\frac{p}{1-p}\right), \tag{18}$$

the values \hat{p} will be normally distributed. We allow for the inclusion of structured noise (to simulate sampling/ageing error) but only along one dimension at a time. We use multi-variate normal noise to introduce the error, but only along either age, length or year dimensions. The noise can be structured with correlation at age or at year, for example, but not currently for both. The reason for this is that this type of noise represents a matrix, and matrix distributional theory is both extremely complicated and not as well developed as vector-based theory. Even allowing for this restriction, this should be more than enough freedom to explore many possibilities.

The delta method states that, given some differentiable transformation $\phi(p_{y,a})$, then the covariance matrix of this transformation, \sum^ϕ , can be expressed as follows:

$$\sum_{ij}^\phi = \frac{\partial\phi}{\partial p_i} \times \sum_{ij} \times \frac{\partial\phi}{\partial p_j} \tag{19}$$

where the partial derivatives of $\phi()$, with respect to the proportions, are evaluated at μ .

Given the logit transformation:

$$\text{logit}(p) = \ln\left(\frac{p}{1-p}\right), \tag{20}$$

we have that the required adjustment to the known covariance matrix, \sum , is as follows:

$$\sum_{ij}^{\text{logit}} = \frac{1}{p_i(1-p_i)} \times \sum_{ij} \times \frac{1}{p_j(1-p_j)} \tag{21}$$

where $p_i = \mu_i$.

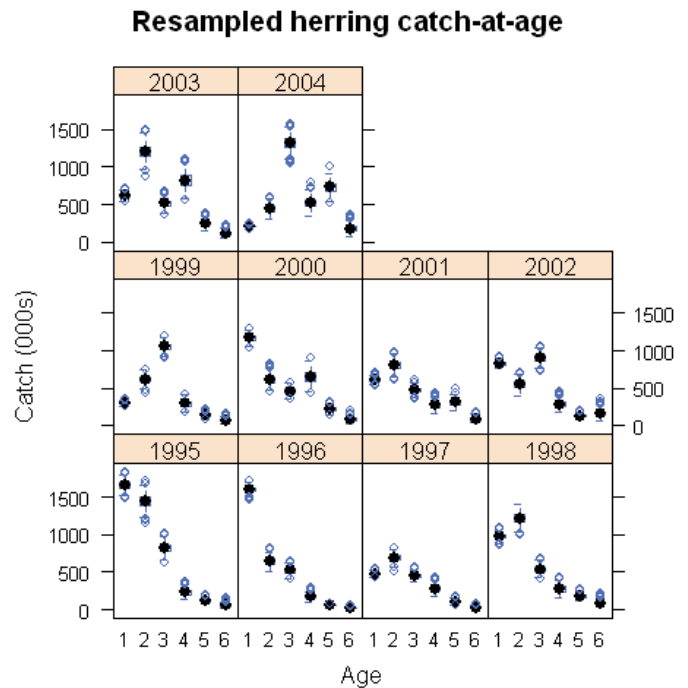


Fig. 5. Box and whisker plot of the resampled North Sea catch numbers-at-age, using the ageing error covariance matrix and the covariant error simulation method as detailed in Sect. 3.2. Catch is detailed in numbers and in thousands of fish.

The theory we use works as follows: we define a suitable adjusted covariance matrix, \sum^{logit} , which represents either the quant or the year dimension in the FLQuant of interest, and we generate the required number of simulations, assuming that

$$\hat{p}_N \sim \text{MVN}\left(\hat{p}_N, \sum^{\text{logit}}\right) \tag{22}$$

where N represents all the other quant dimensions not specified as the error covariate – in real terms, we add a multi-variate normal vector of mean zero and variance \sum^{logit} to every vector of age/length/year data in the FLQuant, depending on the covariate specified. MVN means the multivariate normal distribution.

Once this has been done, we then apply the reverse logit transformation:

$$p = \frac{e^{\hat{p}}}{1 + e^{\hat{p}}} \tag{23}$$

and re-normalise p to finally have our simulated proportions data. For the North Sea herring catch-at-age a covariance matrix for the proportions-at-age was obtained in the EVARES project (J. Simmonds, pers. comm.) and I use this to demonstrate the functionality of the FLprop class and methods. The FLprop methods `prop.sim()` and `regen()` simulate the correlated error in the catch proportions-at-age (defined by the transformed covariance matrix) and then use these resampled proportions to regenerate the observation error-corrected survey data, respectively. Figure 5 displays a box and whisker summary plot of the resampled catch-at-age matrix.

4 Assessment methods

FLR clearly offers a framework within which a range of assessment methods can be incorporated and, in the next section, we will see where these methods can fit into a management strategy evaluation process. Many of the standard assessment models such as an extended survivors analysis, XSA (Darby and Flatman 1994) and an integrated catch-at-age analysis, ICA (Patterson and Melvin 1995) have been translated into FLR – this means that the outputs from such assessments are in the same data structures as the operating models so using key information such as total mortality or trends in SSB/mean length and so on (in the survey-based paradigm) are readily useable in management simulations.

5 Designing and defining harvest control rules

A key facet of MSE work is the designing of suitable harvest control rules (HCR from now on) and from a FISBOAT perspective this means using survey-based data or assessment indices. There can be no one single package that encompasses all possible HCR – there is an infinite array of possibilities. To outline the mathematical considerations that are relevant, I provide two examples of harvest control rules that can be derived from either survey data directly or from survey-based assessment methods.

5.1 HCR using relative SSB trends

The example harvest control rule we apply is simple, based on year-to-year changes in abundance indices. The HCR essentially is a total allowable catch (TAC) adaptation scheme which changes the TAC from year to year, relative to the TAC from the year before and the status of the stock, given the information in the abundance indices of interest. The HCR can be parameterised as follows:

$$\Delta_{y-1} = \frac{I_{y-1}}{I_{y-2}} \quad (24)$$

$$\text{TAC}_{y+1} = \Delta_{y-1} \times \text{TAC}_y \quad (25)$$

where the delay effect of time $y - 2$ affecting the TAC in $y + 1$ is present because of the delay that is intrinsic to the observable effect of changing the TAC on the stock, for this particular management system. Here, I_y represents the trend in spawning stock biomass (relative or absolute) and can either be from a survey directly, or have been estimated by a survey-based assessment method. The basic idea is that the TAC will increase if the spawning stock index is considered to be increasing, and will be decreased if the spawning stock index is decreasing.

5.2 HCR using trends in total mortality

A secondary option is how to define a HCR based upon trends in total mortality, not fishing mortality, as this is a potential output of either an age-based survey or age-based survey

assessment method applied to an age-based survey. The idea was to have two precautionary reference total mortality values, one on the juvenile portion of a stock, and another on the adult/mature portion of a stock, and to adapt the TAC relative to the trends seen in total mortality seen on both these elements of the stock. The basis for this was to try and mirror aspects of the actual decision rule applied to North Sea herring (a case study in the project), but based on total, not fishing mortality levels on the juvenile and adult sections of the stock. Given our precautionary total mortality levels, Z_{pa}^{juv} and Z_{pa}^{adu} , the HCR is defined as follows:

$$\text{TAC}_{y+1} = \text{TAC}_y \times \min \left(\frac{Z_{pa}^{juv}}{Z_{y-1}^{juv}}, \frac{Z_{pa}^{adu}}{Z_{y-1}^{adu}} \right). \quad (26)$$

The HCR defined in Eq. (26) is designed to be a precautionary “traffic light” kind of HCR. We mean that there will only be an increase in TAC if both levels of Z are below their precautionary levels; if either is above or below then the TAC will be reduced. If both Z indices are below their precautionary levels, then the smallest increase in TAC will be allowed. If both are above their precautionary levels, then the largest decrease in TAC will be applied. The estimated values of Z can again be derived directly from surveys or can be estimated from the survey data using an age-based survey-specific assessment method.

There is a large amount of freedom in how we define a HCR. For the simple index-ratio HCR in Sect. 5.1 we might also include a longer set of historical index ratios in the TAC adaption scheme, where issues such as rebuilding or life-history can be incorporated within the HCR. For the Z -based rule in Sect. 5.2 one might wish to have more than a simple juvenile/adult structure in the HCR, or with differential weightings where trends in either the juvenile or adult populations had more influence on the TAC adaption. Of course one could also include combinations of different indices in the HCR – for example we update the TAC based on both the trend in juvenile mortality (relative to precautionary levels) and the trend in the SSB (as given by the index-ratio approach). The object oriented structure of FLR and the affiliated packages and methods mean that the practical implementation of these type of rules is straightforward once one is familiar with the FLQuant data structure – all of these rules involve manipulations of FLQuant objects holding the relevant survey/survey-derived information.

6 Visualisation tools

This final section relates to the visualisation and communication of the simulation results, to aid both the users and the intended stakeholder audiences to which these results will hopefully be presented.

6.1 Graphics within FLR

The default method of plotting FLR-based data objects is to use the lattice graphics package in R. All data objects in FLR

are based on the FLQuant data object, which is a six dimensional array that permits the storage and modelling of complex data objects. Figure 5 shows just some of the usefulness of lattice and the plotting of panel data and Monte Carlo information in FLR.

Being able to use such conditional plotting routines on panel data is extremely useful for exploratory data analysis (pre-model fitting) and also when looking for potential residual trends or when trying to communicate model results (post-model fitting). Plotting Monte Carlo information is also extremely important – especially when conducting management simulations with sometimes considerable observational and process error. There are numerous plotting and statistical functions within R (overloaded for FLR where feasible) that are very helpful with regard to relating Monte Carlo information to others graphically. In lattice the `bwplot()` method (used in Fig. 5) or the `boxplot()` method (Fig. 4) are ideal methods but, as previously stated, many seem to concentrate on the central tendency trend or one of the potential uncertainty ranges and it is incumbent on the user to make this information “real” also. I mean by showing some potential sample paths (be it in future SSB or recruitment or TAC, for example, Fig. 4) – by just showing the median or mean as our best prediction and not making it clear that the reality will be sometimes quite a lot more noisy we risk potentially misleading stakeholders into expecting such future stability.

6.2 Discussion

While the concept of using management strategy evaluation to assess the performance of management procedures is not new, even in a survey-based management paradigm, having a coherent set of freely available tools, such as the FLR platform, is clearly a step forward. In this paper I have tried to (a) outline some general mathematical and statistical tools – with particular attention to conditioning the operating model, simulating surveys and constructing survey-based HCR – that are useful for management simulations, and (b) show how they can be achieved in the FLR framework and how to relate sometimes complex, Monte Carlo information to stakeholders.

Acknowledgements. The author would like to thank an anonymous referee for their review which greatly improved the content and tone of the article.

References

- Anonymous, 2008, Report of the working party on tropical tunas. IOTC, Victoria, Mahe, Seychelles.
- Bull B., Francis R.I.C.C., Dunn A., McKenzie A., Gilbert D.J., Smith M.H., 2005, CASAL User Manual v2.07-2005/07/06. Nat. Inst. Water Atmosph. Res. NZ, Tech. Rep. 126.
- Butterworth D.S., Cochrane K.L., De Oliveira J.A.A., 1997, Management procedures: a better way to manage fisheries? The South African experience. In: Pikitch E.K., Huppert D.D., Sissenwine M.P. (Eds.) Global trends: fisheries management. Am. Fish. Soc. Symp. 20, pp. 83–90.
- Butterworth D.S., Punt A.E., 1999, Experiences in the evaluation and interpretation of management procedures. ICES J. Mar. Sci. 56, 985–998.
- Campbell R.A., Dowling N.A., 2005, Evaluating harvest strategies for a rapidly expanding fishery: the Australian broadbill swordfish fishery. In: Kruse G.H, Gallucci V.F., Hay D.E., Perry R.I., Peterman R.M., Shirley T.C., Spencer P.D., Wilson B., Woodby D. (Eds.) Fisheries assessment and management in data-limited situations. University of Alaska, Fairbanks. Alaska Sea Grant College Program Report, 05–02, pp. 509–532.
- COBECOS, 2008, Costs and benefits of fisheries control strategies. URL: <https://cobecos.jrc.ec.europa.eu/>
- Darby C.D., Flatman S., 1994, Lowestoft VPA Suite version 3.1, MAFF.
- De Oliveira J.A.A., Butterworth D.S., 2004, Developing and refining a joint management procedure for the multispecies South African pelagic fishery. ICES J. Mar. Sci. 61, 1432–1442.
- EFIMAS, 2007, Operational evaluation tools for fisheries management options. URL: <http://www.efimas.org/>.
- Hillary R.M., 2009, Assessment and tag program adaption methods for exploratory fisheries in the CCAMLR area: an example application for sub-area 58.4.3a. CCAMLR Sci. (in press)
- Hoff A., Frost H., 2008, Modelling economic response to harvest and effort control in North Sea cod fishery. Aquat. Living Resour. 21, 259–264.
- Kell L.T., Mosqueira I., Grosjean P., Fromentin J.M., Garcia D., Hillary R., Jardim E., Mardle S., Pastoors M.A., Poos J.J., Scott F., Scott R.D., 2007, FLR: an open-source framework for the evaluation and development of management strategies. ICES J. Mar. Sci. 64, 640–646.
- Nash R.D.M., Dickey-Collas M., Kell, L.T., 2009, Stock and recruitment in North Sea herring (*Clupea harengus*); compensation and depensation in population dynamics. Fish. Res. 95, 88–97.
- Needle C.L., 2008, Management strategy evaluation for North Sea haddock. Fish. Res. 94, 141–150.
- Patterson K.R., Melvin G.D., 1995, Integrated catch analysis version 1.2. - Scottish Fish. Res. Rep. N° 58.
- Punt A.E., Smith A.D.M., 1999, Harvest strategy evaluation for the eastern stock of gemfish (*Rexea solandri*). ICES J. Mar. Sci. 56, 860–875.
- Rademeyer R.A., Plagányi E.E., Butterworth D.S., 2007, Tips and tricks in designing management procedures. ICES J. Mar. Sci. 64, 618–625.
- Ricker W.E., 1958, Handbook of computation for biological statistics of fish populations. Bull. Fish. Resour. Board Can. 119.
- SC-CAMLR, 2008, Report of the working group on fish stock assessment. Report of the 26th Meeting of the Scientific Committee (SC-CAMLR-XXVI), Annex 5. CCAMLR, Hobart, Australia.