Environment
www.csiro.au

Exploring the ABC approach for IOTC Albacore OM conditioning

Rich Hillary
CSIRO Environment
and
Iago Mosqueira
Wageningen Marine Research, The Netherlands

CSIRO Environment
Battery Point, Hobart 7000, Tasmania, Australia.

Copyright and disclaimer

Important disclaimer

Contents

1   Background

For the current suite of IOTC MSE work, the general approach to conditioning the required set of Operating Models (OMs) has been to use the species-specific stock assessment model structure as the basis for the OMs. A grid of model runs, formulated using a set of alternative assumptions and inputs, is constructed based on the base case assessment model. In [1] an alternate, complementary approach was outlined where, instead of the assessment being the basis for conditioning, a suite of possible prior states of historical dynamics and current status are defined. The available, but mostly the more contemporary, data are included within an estimation scheme built on emerging Approximate Bayesian Computation (ABC) and Synthetic Likelihood (SL) concepts [2, 3]. The aim is to generate a distribution of current abundance, mortality and status that is consistent with both the available data and the suite of possible prior states of nature defined beforehand. This can then be used to initialise the OMs used to project the stock into the future and test the candidate MPs.

A stock assessment, in this context, can be viewed as our attempt to do both of these things at once. Ideally, this is arguably a sensible option; however, it is not always successful. The ongoing struggles with the Yellowfin tuna stock assessment, and the conditioning of OMs based upon it, outline this problem: what if you cannot adequately reconcile the data, assessment model structures, and the resultant estimates of current status and future projected dynamics? In [1] we proposed an alternate approach arguing that using a stable, agreed and robust stock assessment was a natural first option, but that the ABC approach was a potentially viable - and scientifically pragmatic - alternative approach if the assessment route was unsuccessful.

In this paper we parameterise a real world example using Indian Ocean Albacore tuna that mirrors (biologically and structurally) the most recent stock assessment, utilises length composition and longline CPUE data, and is able to explore a wide range of stock status prior hypotheses, many of them built on information from the results of the stock assessment.

2   Methods

ABC [2] and SL [3] methods can be used to define an *approximate* distribution for the parameters $\boldsymbol{\theta}$ we are interested in; subsequently, we obtain an approximate distribution for all the variables that depend on those parameters. Where they differ from more classical frequentist or Bayesian methods is how the data, D, are included. Classical methods posit a likelihood for the data, given the parameters: $\ell\left(D \mid \boldsymbol{\theta}\right)$; for a Bayesian analsysis we then define a prior distribution, $\pi(\boldsymbol{\theta})$ to then obtain the posterior distribution of the parameters given the data:

$$\pi\left(\boldsymbol{\theta} \mid D\right) = \frac{\ell\left(D \mid \boldsymbol{\theta}\right)\pi(\boldsymbol{\theta})}{\pi(D)}$$

The ABC approach relaxes the requirement for a specific likelihood (i.e. data generating probability model) to the idea of a *discrepancy* statistic that measures the difference between the observed data, and the model-derived process variables, X, that relate to it. The simplest example would be some distance metric $\rho(D, X)$ whereby we require that this distance between the observed data and our prediction is less than some value $\delta > 0$ (i.e. assumes uniform error on a radius δ). Values whereby $\rho(D, X) \geq \delta$ receive zero probability mass and, in a sampling scheme, would never be accepted. This simple approach will not necessarily work for certain types of data, especially the types we often have in fisheries contexts, but there are natural

generalisations of this simple discrepancy idea.

The other very useful thing we can easily embed within a sampling scheme is informative priors on various elements of the process variables, X, which effectively imply a prior on the key parameters, $\pi(\boldsymbol{\theta})$. This means we can also define informative priors for the various stock status variables (MSY ratios, SSB depletion *etc.*) and include the most relevant recent data, without having to fully model the historical dynamics and data and define appropriate likelihoods for these data. If we are departing from the stock assessment approach, this implies that we have not really succeeded in being able to robustly estimate the distribution of these key status variables. This approach takes a step back and instead tries to define scenarios (specifically distributional scenarios) for these variables at particular periods in time that are consistent with previous assessment experiences and the data. In terms of likely data sources, the two most obvious examples in tuna models are the CPUE indices of abundance, and the catch at length composition. For some examples we may also have mark-recapture data (e.g. tropical tuna) or other data sources. These can be naturally added to the algorithm, as long as the population and fishery model can generate observations for them and a distance metric can be computed.

The population models and associated biological relationships will be very similar to the stock assessment, given they need to have the same structure as the existing OMs (which we don't propose need changing at this stage). The main technical challenge is constructing the discrepancy statistics for the data, and the sampling scheme to generate samples from our approximate distribution of the parameter and population dynamic variables. We have moved the technical details of how these are done to Appendix A to avoid an unnecessary amount of technical exposition in the main body of the paper: what we are trying to demonstrate is that we can get the variables we need for OM conditioning, using plausible status scenarios, key data sets and emerging powerful statistical sampling techniques.

3   Outlining Albacore example

The basis for the models and data explored in the Albacore example is contained in the most recent stock assessment [4]. In terms of model structure, the model is setup as follows:

- Annually structured but with four annual seasons and recruitment in a single pre-specified season.

- No explicit spatial structure but with the same areas-as-fleets approach as the stock assessment.

- Sexually-structured population dynamics driven by growth and selectivity-at-age (selectivity-at-length is not sexually structured by fishery).

- Time-frame for conditioning is $2000 - 2020$, so as to model all surviving cohorts.

- Model considers 4 long-line fleets, 1 "other" fleet and 1 purse seine fleet. The 16 seasonal long-line fleets are condensed into 4 fleets, each with 4 seasons of catch, effort and size data (e.g. LL1 is fisheries 1-4 in the assessment).

- A Beverton & Holt steepness-unfished recruitment stock-recruit relationship is used with lognormally distributed deviations constrained by a pre-specified $\sigma_r = 0.3$, as per the assessment [4].

The stock assessment model fits to the disaggregated length-frequency data and the seasonal CPUE data from the longline fleets. The initial approach taken in the ABC albacore modelling

was not to fit to disaggregated length frequency data, which is very variable and noisy from season to season and year to year. Instead, we aggregate the size data across years and seasons and fit to a mean size frequency data set per fishery. The reason behind this choice is to obtain a representative selectivity relationship for each fishery, but don't rely on these data to inform the model on size/age structure in the population over the time period. Additionally, it purposefully reduces the influence of the size data on the overall population abundance scale, which is a known issue already considered in the previous OM grid.

In terms of stock status prior information we use initial (female) SSB depletion and the final two years SSB at MSY ratios S_y/B_{msy}. At this stage we chose not to explore fishing mortality-related MSY ratios because of the difficulty with interpreting exactly how Stock Synthesis calculates these ratios in the seasonal modelling context. With respect to initial (year 2000) SSB depletion we have three options given its central importance for in particular this modelling approach. The first is to impose a single value; the second is to predefine an initial distribution and sample from it; the third is to define a prior for it and treat it as an estimated parameter. For this simple outlining example we use the first option: we fixed the year 2000 SSB depletion at 0.5 (very close to the apparent estimate from the assessment [4]). For the final two years of the SSB MSY ratios we define mean values of 2.25 and 2, respectively, and with standard deviations of 0.35 in both cases. Again, these are based roughly on those values listed in the assessment. We estimate R_0, annual recruitment deviations and double normal selectivity parameters for each of the five fisheries using the available LF data, with the *Other* fleet sharing the *PS* selectivity. This leads to a total of 36 estimated selectivity parameters. For the example, in this paper we have selected only the non-seasonal catchability parameter scenario, though we have also ran a seasonal one, and used the CPUE from LL1 (fleets 1–4 in the assessment) for all seasons and for all years in the model.

4 Results

Figure 4.1 outlines a selection of the possible population dynamic summaries - SSB depletion, mature biomass MSY ratio, recruitment, and selectivities - for the example ABC run. Figure 4.2 presents the fits to the seasonal LL1 CPUE data and the size frequency data for the five fisheries with estimated selectivity parameters.

5 Discussion

This paper outlines both the theoretical and practical ideas behind the suggestions originally raised in [1], with respect to methods for conditioning OMs that are not built upon the stock assessment model structures and output. The approach uses much, if not all, of the same demographic and life-history parameters the assessment does, and often key subsets of the same data if they are required to be simulated in the OMs. It also includes key prior information on current/recent/historical status variables (MSY ratios, depletion *etc.*) to help estimate the main abundance, biomass and mortality variables required in the OMs. There is already an IOTC precedent for these type of approach - the original skipjack tuna OM was conditioned in a similar fashion [7] using the methodology outlined in [8].

The simple example presented in this paper for Indian Ocean Albacore was chosen more to outline (a) what data we can use, (b) what methods and variables we can use in the ABC algorithm, and (c) what kinds of outputs we can get that related directly to conditioning OMs in the MSE context. In terms of data we essentially use the same as the most recent assessment [4]: size
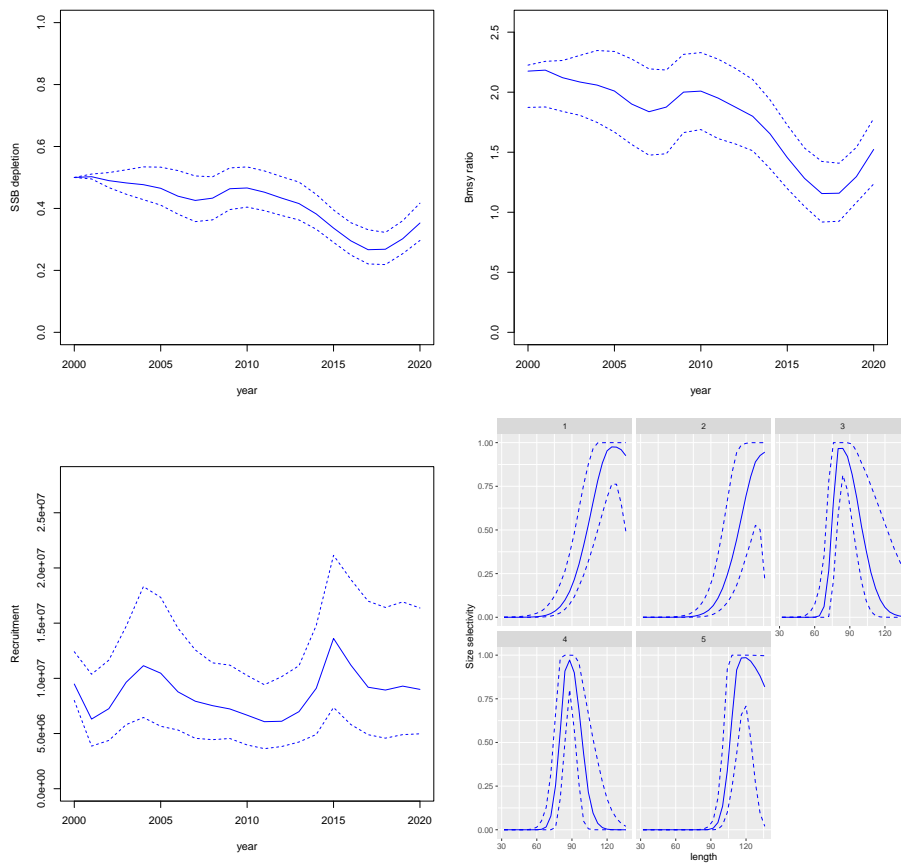
Figure 4.1: *SSB depletion (top left), SSB MSY biomass ratio (top right), overall recruitment (bottom left), and size selectivities (bottom right) posterior credible interval summaries (median full line and 95%ile dotted lines).*

frequency and longline CPUE data. In terms of variables we use predicted size frequency and exploitable biomass, as well as prior information that can be (but need not be) informed by the stock assessment e.g. MSY and SSB depletion ratios. In terms of the outputs we can reliably claim to be able to produce, this includes everything that the stock assessment-based OM conditioning approach can: abundance at age and SSB, harvest rates, MSY variables and catchability coefficients. The aim here is for the task force meeting to consider and discuss if the approach appears likely to be able to produce what we need for the OMs, and what additional features would be required to make that happen.
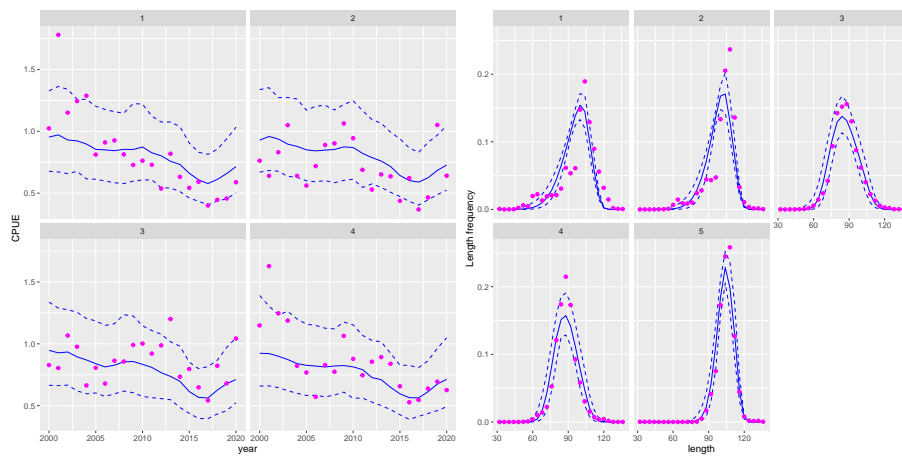
6   Acknowledgements

Figure 4.2: *LL1 seasonal CPUE (left) and size frequency (right) data posterior credible interval summaries (magenta circles are observed; median full line and 95%ile dotted lines predicted).*

References

[1] Hillary, R., Mosqueira, I., Preece, A., Rosa, D., Williams, A. and Edwards, C. (2021) Strategies for the conditioning of operating models for IOTC stocks, *IOTC–2021–WPM12–16*

[2] Sunnaker, M., Busetto, A., Numminen, E., Corander, J., Foll, M., *et al.*(2013). Approximate Bayesian Computation. *PLoS Comput. Biol.* **9**(1):e1002803. doi:10.1371/journal.pcbi.1002803.

[3] Wood, S. (2010) Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*. **466**: 1102–1107. doi: 10.1038/nature09319.

[4] Rice, J. (2022) Stock assessment of albacore tuna (*Thunnus alalunga*) in the Indian Ocean using Stock Synthesis. *IOTC–2022–WPTmT08–09*.

[5] Wilkinson, R. D. (2013) Approximate Bayesian Computation (ABC) gives exact results under the assumption of model error. *Stat. Appl. Genet. Mol. Biol.* **12**(2): 129–141.

[6] Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (1993) Handbook of Markov Chain Monte Carlo. *Chapman & Hall/CRC*.

[7] Bentley, N., and Adam. S. (2016) Management strategy evaluation for the Indian Ocean skipjack tuna fishery. Working paper IOTC-2016-WPM07-15 Rev1 presented to the 7 th IOTC Working Party on Methods, Mahe, Seychelles 11-13 November.

[8] Bentley, N. and Langley, A. (2011). Feasible stock trajectories: a flexible and efficient sequential estimator for use in fisheries management procedures. *CJFAS*. **69**(1): 161–177

[9] Kullback, S., and Leibler, R. A. (1951) On information and sufficiency. *Annal. Math. Stats.* **22**: 79–96.

Appendix A

Approximate McMC ABC algorithm

There are a wide variety of possible algorithms that can be used to generate a sample from the approximate posterior distribution [2, 3, 5]. Given the relative complexity of our likely suite of models, we consider that Algorithm D from [5] is the most applicable. It is basically an ABC-configured Metropolis-Hastings accept-reject algorithm as used in classic Bayesian McMC contexts. At time t, we have our joint parameter and process variable state $\Xi_t = \{\boldsymbol{\theta}_t, X_t\}$. We generate a proposal for a new parameter vector $\boldsymbol{\theta}'$ (and $X' = f(\boldsymbol{\theta}')$) from the pre-specified transition kernel $q(\boldsymbol{\theta}_t, \boldsymbol{\theta}')$. We define the following acceptance probability for $\Xi' = \{\boldsymbol{\theta}', X'\}$:

$$\alpha(\Xi_t, \Xi') = \min\left(1, \frac{\pi(D, X')q(\boldsymbol{\theta}', \boldsymbol{\theta}_t)\pi(\boldsymbol{\theta}')}{\pi(D, X_t)q(\boldsymbol{\theta}_t, \boldsymbol{\theta}')\pi(\boldsymbol{\theta}_t)}\right)$$

and generate a random variable $u \sim U[0, 1]$. If $\alpha(\Xi_t, \Xi') > u$ we accept the proposal and $\Xi_{t+1} = \Xi'$; if $\alpha(\Xi_t, \Xi') \leq u$ we reject the proposal and set $\Xi_{t+1} = \Xi_t$. By choosing a symmetric normally distributed transition kernel $q(\boldsymbol{\theta}_t, \boldsymbol{\theta}') = q(\boldsymbol{\theta}', \boldsymbol{\theta}_t)$ this term disappears from the acceptance rate calculations. The prior distribution, $\pi(\boldsymbol{\theta})$, is for the estimated parameters. The $\pi(D, X')$ term is our likelihood analogue or discrepancy function - it includes both the observed data (catch length composition, abundance indices) and our additional information on the key status variables contained in the suite of process variables, X (e.g. MSY and depletion ratios). The Markov chain transition kernels $q()$ are defined to use the random walk approach for sampling the posterior surface. To make the McMC algorithm more efficient we implement a Metropolis-within-Gibbs sampling approach [6], where parameters are grouped together depending on expected correlation. Each block is updated using the Metropolis-Hastings algorithm, conditional on the parameters not included in that block being fixed at their most recent value. After doing this for each block of parameters (the Gibbs sampling part of the algorithm) we have fully updated all the parameters of the model and repeat the same process many times. Random walk variances are adjusted to achieve acceptance rates around 40% - generally considered as optimally efficient [6]. A suitable burn in period is used to get the sampler moving on the surface before we decide to keep the samples, and these are then thinned to remove autocorrelation in the Markov chain. When this is done we have 1,000 samples from the approximate posterior distribution of interest and we use the Geweke statistic [6] to test for non-convergence of the Markov chains.

For the length composition data we took a nonparametric approach using the concept of the Kullback-Leibler divergence (KLD, [9]): this is a measure of the divergence between in this case a discrete probability distribution P_i, relative to a reference distribution Q_i. It is defined as follows:

$$D_{KL}(P \parallel Q) = \sum_i P_i \ln\left(\frac{P_i}{Q_i}\right) \geq 0$$

with the convention that

$$\lim_{x \to 0^+} x \ln(x) = 0.$$

The KLD serves as a potentially very useful option in the ABC sense for the following reasons:

1. It is nonparametric so the underlying generating distribution of the length data does not need to be assumed

2. It reduces to zero when $P_i \equiv Q_i$ and increases the further Q_i diverges from P_i - by defining P_i as our observed data and Q_i our model prediction it makes an obvious candidate as a discrepancy measure of lack of fit

3. The units are interpretable. For natural logarithms the units are called nats - 1 nat is basically a difference in probability of $1/e$. This means we can set tolerance levels for how much divergence we are willing to accept in our predicted data that have a grounding in information theory

In practice we use the *negative* KLD - it reaches a maximum at perfect prediction of the data and decreases as this gets progressively worse, much like a likelihood does. If we define $p_{f,l}$ as our fishery-specific observed (annually and seasonally aggregated) length data and $\widehat{p}_{f,l}$ as our predicted length composition this part of the discrepancy function can be defined as

$$\mathcal{D}_{LF} = -\sum_f \sum_l p_{f,l} \ln \left(\frac{p_{f,l}}{\widehat{p}_{f,l}} \right)$$

We also set an effective maximum KL value of 0.8, so that no proposals of new parameter vectors are accepted if the KL value is above this maximum. The rational for a value of 0.8 is that this would roughly align with the upper confidence level of the KL value if the length data were truly distributed around the observed distribution with a multinomial effective sample size of 20.

For the CPUE data we assume an effectively lognormal distribution for the discrepancy function. For seasonal catchability models we calculate the catchability coefficient as follows:

$$\ln q_s = \mathbb{E}^y \left[\ln \left(I_{y,s} / \widehat{X}_{y,s} \right) \right]$$

and $\widehat{X}_{y,s}$ is the seasonal exploitable biomass. For the scenario with a constant catchability across all seasons it is calculated as follows:

$$\ln q = \mathbb{E}^{y,s} \left[\ln \left(I_{y,s} / \widehat{X}_{y,s} \right) \right]$$

The standard deviation for the CPUE discrepancy is taken by fitting a LOESS smoother to the seasonal log-transformed observed CPUE data and calculating the standard deviation in the residuals. The rationale being we want an overall (i.e. both observation and process error) measure of likely variation in the observed CPUE, not an observation only estimate.

For the status variables (MSY and depletion ratios) we simply defined quadratic kernels for the log-transformed variables averaged over the relevant time frame:

$$K(x, y \mid \epsilon) = \frac{\| x - y \|^2}{\epsilon^2}$$

In this case the tolerance for each kernel ϵ can basically be interpreted as twice the standard deviation of a normal distribution. The summation of the CPUE and length composition discrepancies and the status prior kernels makes up the overall discrepancy function $\pi(D, X)$.

WE DO THE EXTRAORDINARY
EVERY DAY
We innovate for tomorrow and help
improve today for our customers, all
Australians and the world.
Our innovations contribute billions of
dollars to the Australian economy
every year. As the largest patent
holder in the nation, our vast wealth of
intellectual property has led to more
than 150 spin-off companies.
With more than 5,000 experts and a
burning desire to get things done, we
are Australias catalyst for innovation.
WE IMAGINE. WE COLLABORATE.
WE INNOVATE.