



Not to be cited without prior reference to the author

C.M. 1991/D:19
Sess.U

Progress and Problems in Fish Abundance Estimation

Experience with the Adaptive Framework as a Calibration Tool
for
Finfish Stock Assessments in CAFSAC

Stratis Gavaris
Marine Fish Division
Department of Fisheries and Oceans
Biological Station
St. Andrews, New Brunswick E0G 2X0

ABSTRACT

The adaptive framework has been adopted by CAFSAC as a standard calibration tool for conducting analytical assessments. ADAPT, the acronym for this approach, is not a rigid model, rather it provides a common statistical estimation framework based on least squares theory while allowing sufficient flexibility in the definition of model relationships and structures to accommodate the diverse situations encountered. Since the approach computes the sequential population analysis while estimating the survivors of each yearclass and the calibration constants for the indices simultaneously, it falls into the class of methods referred to as "integrated methods" by ICES. The diagnostics available from a least squares analysis have proven useful in determining how much information content is encapsulated in the data and thereby indicating which population parameters can be estimated reliably. Finally, the adaptive framework can be advantageously employed to examine the implications of alternative assumptions in the model relationships or alternative treatments of the three classes of information, the biological sampling of the commercial catch, the landings and the abundance indices. A common, objective estimation method providing statistical diagnostics is essential for this type of data analytic approach to fish stock assessments.

INTRODUCTION

Until the late 1980s the Canadian Atlantic Fisheries Scientific Advisory Committee (CAFSAC) employed 'ad hoc tuning' methods to conduct analytical assessments of the finfish resources of eastern Canada. Much time was spent discussing the criteria to be used in the selection of the 'best' solution. Debates revolved around issues such as the inclusion of the terminal year observation in the catchability regression or whether residuals should be minimized along cohorts. In traditional data analysis using statistical techniques the problem of selecting the 'best' solution is transformed to one of considering properties (eg. precision, bias) of alternate estimators (eg. least squares, maximum likelihood). There is a large body of knowledge in the statistical literature regarding the properties of many classes of estimators. The complexity of the stock assessment problem mitigated against straight forward application of standard statistical techniques which would then permit us to draw from that broader experience regarding how to select the 'best' solution. Specifically, there are two characteristics of the problem which create difficulties. The first is the problem of simultaneously estimating the population abundance and the catchability coefficients for the fleets or research surveys being used to monitor trends. The second is the need to account for the propagation of impact along cohorts which results in linking the catchability relationships of all age groups being estimated. Typically 'ad hoc tuning' methods circumvented these difficulties by dealing with the issues sequentially and going through several iterations. The methods were 'ad

hoc' in two respects, the order of the sequence and the definition of 'best' solution. Further, the population models which were assumed were vaguely described. Despite all these problems, 'ad hoc tuning' methods were capable of providing usefull results when applied judiciously. However the process of arriving at a consensus on 'best' solution was cumbersome and very time consuming and the rationalization of the results arising from an 'ad hoc' process were becoming increasingly difficult to defend to representatives of client groups.

Through the mid/late 1970s and the 1980s several investigations focused on the problem of expressing the stock assessment problem in terms of a standard statistical estimation problem. This approach promised to resolve the major issues; selection of the 'best solution' would be a matter of exploring statistical properties of competing estimators, population and catchability coefficients could be estimated simultaneously and all year classes being estimated could be handled simultaneously thereby accounting for impact along cohorts. A shared feature amongst most of this work was the acceptance of common statistical estimators as being adequate while endeavouring to integrate the fisheries models with the estimation techniques. While these techniques were viewed as advancing the science, few were implemented for stocks being considered by CAFSAC because of the perceived difficulty in tayloring them to specific stock situations.

During the late 1980s the adaptive framework (ADAPT) was introduced in CAFSAC. The adaptive framework also endeavoured to express the stock assessment problem in terms of a standard statistical estimation problem but it incorporated two distinguishing features. The first, a philosophical difference, was to emphasize the distinction between the estimation technique and the definition of the fisheries models. Thus, several alternate fisheries models could be compared within the same estimation framework. A corollary to this was that assessment committees or working groups could retain features of models which were particularly suited to one stock but which were not generally applicable. The second distinguishing feature, a practical concern, was to choose an estimation implementation which minimized the fisheries biologist's involvement with the statistical algorithms. Eliminating the need to modify statistical software while providing flexibility to define fisheries models were the key to presenting an attractive option within CAFSAC.


The adoption of the adaptive framework had the effect of redirecting the focus of CAFSAC Subcommittees discussions away from how to select the 'best' solution to consideration of the adequacy of relationships employed and the validity of the assumptions being made. This was facilitated by the flexibility which the adaptive framework provided in model definition. Initially, this flexibility was a detraction as it demanded that the user be more knowledgable to implement the technique than would be required with a 'fixed model' approach. I contend though, that collective experience gained within CAFSAC through the experimentation and exploration of fisheries models with the adaptive framework has contributed greatly to our understanding of the information value of fisheries related data. Though introduced as an assessment method, the adaptive framework is proving a useful research tool.

My objective is to demonstrate the utility of the adaptive framework as a tool to explore the information content of fisheries related data with respect to estimating abundance. I will first outline the fish abundance estimation problem and describe how the adaptive framework addresses it. This is followed by a selection of some typical issues which can be studied and have received some consideration by CAFSAC, demonstrating how the adaptive framework serves as a tool for data analysis.


THE FISH ABUNDANCE ESTIMATION PROBLEM AND ADAPT

For our purpose it is useful to view statistical estimation as the identification of model parameter values which produce predicted values of variables that best match the sample observations. Now, pose the fish abundance estimation problem in this way. Observations have been made on several variables which reflect the state of a stock, typically the commercial fishery catch at age and abundance indices from either research surveys or commercial fishery

catch rates. Which parameters, that is which population abundance, result in predicted values of the observed state variables which are in best agreement with those observations. The vehicle for going from the parameters to the predicted state values is some imputed model.



We select those parameters which give the best agreement between the observed and predicted state variables for the specified model. Stated in this way, the problem is suitable for standard statistical practices. Capitalizing on this feature and replacing the 'compare' with an accepted objective statistical tool, the least squares estimation method (LS), we obtain the adaptive framework. The nature of the problem necessitates the use of non-linear least squares (NLLS) techniques. The software generally employed by CAFSAC implements the well known Marquardt algorithm.



As stated previously, one objective was to eliminate the requirement for fisheries biologists to become involved with the statistical algorithms. Non-linear techniques require derivatives of the model equations with respect to each parameter. Though numerical methods for obtaining derivatives require more computation and result in some loss of information, by nature of being approximations, the simplification afforded compels their use. The software employed by CAFSAC implements traditional finite difference techniques to obtain numerical approximations of derivatives.

As described here, the adaptive framework, given the acronym ADAPT, provides a consistent statistical framework for the estimation of fish abundance without dictating a specific model. This approach espouses a data analysis philosophy of employing statistically based estimation practices while maintaining flexibility in model specification.

STUDIES ON SPECIFIC ISSUES

The following studies are not intended to provide definitive solutions for the issues posed, indeed the answers may differ from stock to stock as situations vary. Further, these issues do not necessarily represent those requiring priority attention. Priorities will also vary from stock to stock. They are presented here to demonstrate how ADAPT can help to increase our understanding.

These studies were done with data for haddock in unit areas 5Zj and 5Zm on Georges Bank from Gavaris and Van Eeckhaute(1991). Landings data and information on the age composition of the commercial catch were available for 1969 to 1990. Three age structured survey abundance indices were available, two conducted by the USA since 1963 and 1968 in the autumn and spring respectively, and one conducted by Canada since 1986 in the spring.

Are all indices equally reliable?

Often, we have several indices of the abundance of year-classes. This may be due to having several sources, for example more than one research survey, or by having age structured information so that we can track the abundance of a year-class at several ages during its life history. Though a situation where multiple indices are available is desirable, it requires that we compare the relative merits of the available indices. Even if we take the naive approach, that is to not weight the indices, we are implicitly assuming that all indices are equally informative. We should be cognizant therefore that even if indices are not weighted explicitly, an assumption is always implied regarding their relative importance.

We can usefully distinguish two types of error when considering how to weight indices, sampling error and model error. Sampling error arises because the observed variable values have been obtained from a sample rather than by measurement over the entire population. For example, a survey index is the result of a sample of tows. A different sample of tows would result in a different result. To obtain the true exact population value we would need to conduct tows at all possible locations. Model error arises from the specification of relationships which are not perfect even when sampling error is absent. The models we impute are imperfect representations of the real world. For example, we assume that research survey catchability is constant. Though it may be roughly constant it is not exactly constant. That is, even if we could tow at all possible locations, the ratio of survey abundance to population abundance would not be exactly constant from one year to the next. Residuals are an estimate of the composite sampling and model error.

If we only had to deal with sampling error and could assume that model error was negligible, we might be able to obtain a measure of the relative weight we should assign to our observations from consideration of sample design. Generally however, it will not be safe to assume that model error can be ignored. It is not feasible to estimate a relative weight to assign to each individual observation. One practical approach is to identify groups of observations of homogeneous error. An estimate of the relative weight of each group of observations can then be obtained as a function of the residuals (Judge et. al 1980). Natural groupings of observations can be exploited advantageously. One possibility is to assume that the error within any one index is homogeneous.

Examination of residuals from an unweighted analysis of the haddock data suggests that there are differences in their magnitude between age groups within and between the three available surveys (Table 1). Based on these results, Gavaris and Van Eeckhaute (1991) employed the technique referred to above to estimate relative weights for each of these groups and concluded that the parameter estimates from the weighted analysis were superior. This demonstrates how the residuals from an ADAPT analysis incorporating multiple indices can be used as a diagnostic in evaluating the relative merit of the indices.

Can we estimate all year classes or do we need to fix oldest age F?

It is a common practice to assume that the fishing mortality on the oldest age group is equal to the average fishing mortality over the fully recruited ages in the respective year. This model assumption is made in the interest of parsimony. To examine if this assumption is reasonable, a model formulation where the oldest age abundances were estimated in addition to the terminal year abundances and the catchabilities can be considered.

The results of such an analysis for the haddock data gave a better fit even after accounting for the lost degrees of freedom due to the greater number of parameters being estimated. However, most of the oldest age abundances for year-classes between 1961 and 1977 were poorly estimated (Table 2). The correlation matrix of parameters contained several high positive correlations between year-class abundance estimates which also suggests that individual parameters were ill determined. Further the fishing mortality rates at older ages displayed some disturbing fluctuations and at least one year-class was

estimated to be too small to account for all the catch during its history. These peculiarities in the fishing mortality may be attributed to an increased tendency for age misclassification at older ages in both the catch at age and survey abundance indices. Despite these problems, the resulting fishing mortality matrix indicated that the assumption about the age 8 fishing mortality being equal to the average of fully recruited ages might be a justifiable approximation. This demonstrates that the statistical properties of the parameter estimates, eg. variance and covariance, are useful diagnostics in determining what is estimable. Further, comparison of results, such as fishing mortality, from alternate model formulations can give insight into the validity of assumptions.

Should we consider alternatives to assuming that error in the catch at age is negligible?

We may want to examine the impact of assuming that catch at age is known without error. Models which recognize error in the catch at age require many more parameters. Generally, increasing the number of parameters greatly is associated with ill-determined parameter estimates because of intrinsic colinearities. A parsimonious model which does give relatively well determined estimates however, is one where we assume separability of the fishing mortality matrix. Thus, rather than having to estimate a fishing mortality for each age and year, we estimate an annual fully recruited fishing mortality and an average exploitation pattern at age.

Using this class of model with the haddock data, the year-class estimates were generally well determined as were the fully recruited fishing mortality rates, the average exploitation pattern and the catchabilities. The average exploitation pattern was unusual though, as age three displayed a higher exploitation rate than ages 4 and 5. Further, this model is not capable of emulating a higher exploitation rate on larger year classes at the partially recruited ages, a feature which is consistent with the perception of how the fishery operates. There were also some high correlations between parameters which require further exploration. Despite these problems, this model or modifications of it appear promising and worth pursuing. As this case demonstrates, a comparison of results may identify that competing models with alternative sets of assumptions can provide equally satisfactory fits to the information. If there are no compelling reasons to prefer one set of assumptions over others, then it would be wise to examine the implications of all acceptable models.

What can we do with unreliable landed weight?

The reliability of landed weight has been questioned on several occasions. If we suspect that the landed weight contains random error, then a variant of the model presented above which incorporates error in the catch at age should be considered. If however, it is suspected that the information is biased, such as might be the case if some landings are not reported, and we cannot obtain a measure of the bias, we may elect to exclude the data. Under these circumstances, the available information will generally comprise of survey abundance indices and size or age composition from the commercial fishery.

Using only the haddock survey information and the age composition of the commercial catch, it was immediately evident that absolute population estimates could not be determined. Variances and covariances were extremely large and the converged solution was probably not globally optimal. These results confirm common wisdom that the information which allows us to calibrate survey trends to absolute population magnitude is the magnitude of the catch. A model formulation which attempted to estimate only relative year-class strengths rather than absolute abundance was much better behaved statistically. Nevertheless, individual parameters, year-class strengths and fishing mortality rates, were not as well determined as would be desired. The overall abundance trend (Fig. 1) was similar to the pattern from the formulation which incorporated landings data but there were substantial annual differences. The fishing mortality trend showed great annual fluctuation and several estimates of fishing mortality rate were negative. This is not unlike the kind of results obtained when calculating mortality rates from abundance indices alone. The adaptive framework can be molded to handle data poor situations, such as the

case presented here where landings were not incorporated. By virtue of the common framework, we have the same suite of diagnostics for comparative purposes with models which incorporate more or different types of information. This facility is useful for determining the impact of different types of information to the assessment results.

What can we do with total (no age structure) abundance indices?

Frequently, the abundance indices do not contain information on age structure as is the case for example with acoustic survey results. The survey indices for the haddock data were aggregated for ages 1 to 8 to evaluate how well the population abundance could be determined with this type of data. The results showed that though the estimate of total abundance was not unreasonable, the age structure of the population could not be determined. Indeed, there were several negative population estimates and the standard error of the estimates exceeded the mean for most age groups (Table 3). A model formulation which estimates the abundance of one year-class and determines the abundance of other year-classes from an assumed exploitation pattern gave acceptable results. This case is another example of how the diagnostics from ADAPT can help to determine what is estimable and what needs to be assumed.

How does the Laurec-Shepherd method relate to ADAPT?

The Laurec-Shepherd method uses a sequential/iterative approach for estimating population abundance. Can the model used in the Laurec-Shepherd method be extracted and employed within ADAPT? No. Models which come close to doing what the Laurec-Shepherd method does can be devised, including the incorporation of various popular weighting schemes like the 'wedge', but the exact model cannot be implemented. To implement a model within ADAPT it is necessary to distinguish between observations and unknown parameters. The Laurec-Shepherd method assumes a model for catchability, for example that catchability is constant, and then compares 'observations' of catchability to the model prediction. In fact, we do not have observations of catchability. The annual values of catchability which are being treated as observations are a function of population abundance. Population abundance is an unknown parameter and ADAPT, as well as other integrated methods, estimate population abundance and catchability simultaneously. With sequential methods catchability is modeled conditional on an assumed population abundance. The predicted catchability is then used to update the assumed population and this process is repeated. As alluded to in the introduction, simultaneous and sequential estimation, though fundamentally different, can probably be made to provide similar results. The way in which data are weighted is more likely to have significant consequences for the results and deserves the focus of our attention.

SUMMARY

The adaptive framework, ADAPT, provides a common statistical tool for simultaneously estimating all of the unknown parameters for the stock assessment problem. Its most significant positive feature, and one which makes it useful as a research tool in contrast to an assessment method, is the facility it provides for comparing alternative models and different treatments of the data. Standard statistical diagnostics as well as model output can play an important role in researching the information content of the various types of data available for stock assessment.

REFERENCES

- Gavaris, S. and L. Van Eeckhaute. 1991. Assessment of haddock on eastern Georges Bank. CAFSAC Res. Doc. 91/??: 30 p.
- Judge, G. G., W. E. Griffiths, R. C. Hill, and T. C. Lee. 1980. The theory and practice of econometrics. John Wiley and Sons, New York. 810 p.

Table 1. Average squared residuals from an unweighted ADAPT model formulation indicate that there are differences between age groups and surveys.

	Age							
	1	2	3	4	5	6	7	8
Survey 1	1.61	1.80	1.17	0.84	0.68	0.70	0.74	1.83
Survey 2	1.30	0.72	0.70	0.48	0.55	1.10	1.77	0.80
Survey 3	0.38	0.32	0.19	0.17	0.27	0.22	0.80	0.19

Table 2. Estimates of year-class abundance for age 8 were generally poorly determined.

Year-class	mean	se/mean	Year-class	mean	se/mean
1982	284	0.29	1971	252	0.49
1981	117	0.30	1970	34	0.75
1980	353	0.29	1969	89	0.52
1979	313	0.31	1968	22	0.68
1978	980	0.32	1967	29	0.59
1977	103	0.76	1966	161	0.49
1976	115	0.67	1965	106	0.77
1975	2113	0.50	1964	88	0.68
1974	106	1.05	1963	2633	0.57
1973	72	0.68	1962	2630	0.60
1972	729	0.45	1961	689	0.75

Table 3. Population abundance at the beginning of year 1991 was not well determined in the ADAPT model formulation using age aggregated indices.

Age	mean	se/mean
1	374	31.88
2	7653	1.29
3	11119	0.57
4	-2762	-1.43
5	-256	-19.98
6	1456	3.49
7	3842	0.88
8	-2778	-0.53




FIGURE 1: Comparison of population numbers from two alternative ADAPT model formulations.