ICES Journal of Marine Science, 58: 1133–1149. 2001
doi:10.1006/jmsc.2001.1103, available online at http://www.idealibrary.com on IDEAL®

Bayesian estimation of fish school cluster composition applied to a Bering Sea acoustic survey

T. R. Hammond, G. L. Swartzman, and
T. S. Richardson

Hammond, T. R., Swartzman, G. L., and Richardson, T. S. 2001. Bayesian estimation of fish school cluster composition applied to a Bering Sea acoustic survey. – ICES Journal of Marine Science, 58: 1133–1149.

This paper applies BASCET, a Bayesian Spatial Composition Estimation Tool for clusters of acoustically identified schools, to Bering Sea acoustic survey data collected during 1994. As the method employs prior information from an acoustic expert, procedures for eliciting such information are suggested and pitfalls of the process are indicated. Techniques for model checking using the posterior predictive distribution are employed, as is a multi-chain method for evaluating the convergence of the Markov-Chain Monte Carlo algorithm used in BASCET. Unlike methods based on neural networks, BASCET is able to provide confidence regions for its estimates of school cluster composition. In addition, it can indicate which school cluster attributes were most influential in determining a given estimate, a useful tool for model checking that is here demonstrated on a randomly selected cluster. Estimated abundance ratios of juvenile to adult pollock (*Theragra chalcogramma*) were compared, in two regions, to the values used by expert technicians. Ratios differed from expert values by less than 0.03 in both regions. The encouraging results reported here suggest that the BASCET method, originally tested on simulated data, may be usefully applied to real surveys.

Keywords: acoustic survey, Bayesian, prior elicitation, species discrimination.

Received 5 December 2000; accepted 18 April 2001; published electronically 29 August 2001.

T. R. Hammond, School of Fisheries, University of Washington, Seattle, WA, USA. G. L. Swartzman, Applied Physics Laboratory, University of Washington. T. S. Richardson, Department of Statistics, University of Washington. Correspondence to: T. R. Hammond, CEFAS, Pakefield Road, Lowestoft, Suffolk NR33 0HT, England, UK; tel: +44-1502-524213; fax: +44-1502-524511; e-mail: t.hammond@cefas.co.uk

Introduction

Species and size-class identification is an essential part of the analysis of acoustic survey data, with important impacts whenever acoustic surveys are used in stock assessment. The task is commonly done by eye, whereupon it is subject to inconsistency. Methods for automatically classifying fish schools are also available and are discussed in a companion paper (Hammond and Swartzman, 2001). The companion paper also describes BASCET, an algorithm unique in placing classification within a Bayesian framework, in focusing attention on training set construction, and in incorporating spatial structure into the estimation process.

The main objective of this paper is to evaluate the performance of BASCET by comparing the results it produces using data from the 1994 Bering Sea acoustic survey to corresponding values computed by expert technicians. This paper also serves as a blueprint for the implementation of BASCET on a real acoustic survey. While this Bayesian algorithm is fully specified by the equations in Hammond and Swartzman (2001), the user is still faced with the task of selecting prior distributions for a complex set of model parameters. Ideally, the posterior distributions from one year's survey should be used as priors for the next. Alternatively, one might use posterior distributions from a nearby survey to form the basis of the priors in the survey of interest. Of course, these suggestions are of little help where BASCET has never been applied. For such cases, priors should be elicited from acoustic experts. An example of how this can be done, and recommendations on how it might be done better, are provided here.

In any application of BASCET one should verify its assumptions and examine their impacts on estimation results. In this paper, test quantities from the posterior predictive distribution are simulated in order to check

1054–3139/01/061133+17 $35.00/0

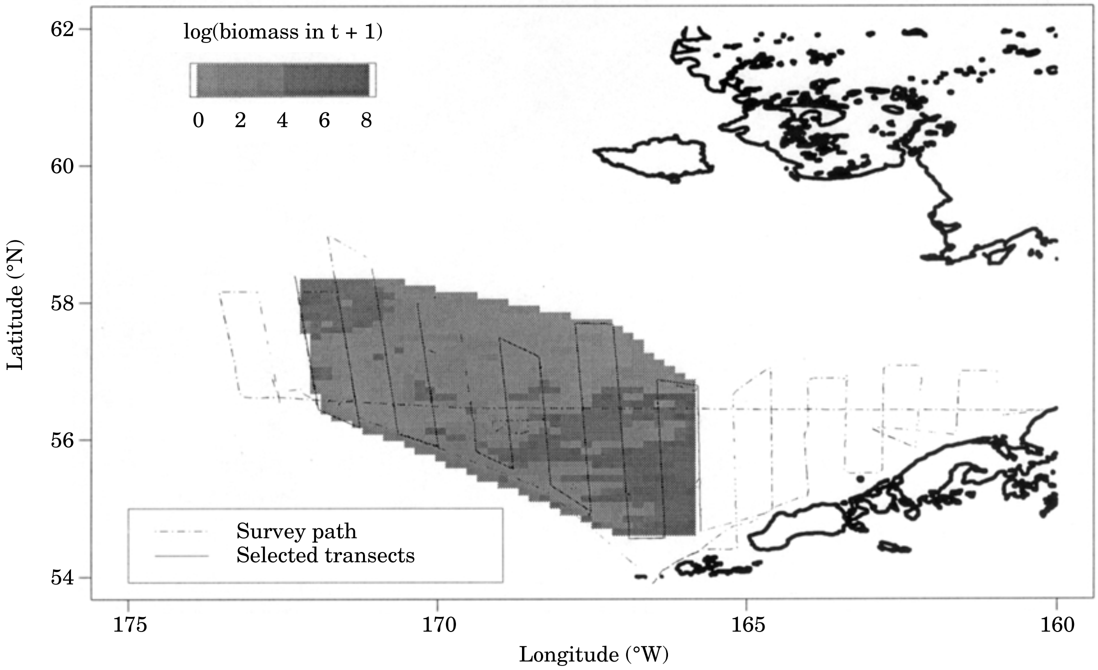

Figure 1. The 1994 Bering Sea acoustic survey track, shown as a broken line, selected transects as a solid line, and the distribution of biomass as estimated by the Bergen echo-integrator as a background. AFSC scientists believe these estimates represent adult and juvenile pollock.

model assumptions. BASCET can also give a measure of justification for its answers, a feature that can indicate problems (or provide reassurance) when these explanations are compared to those provided by human experts. Such explanations typically indicate which school cluster attributes most influenced a given composition estimate. They can also highlight attributes that suggest that the cluster may be composed of something else entirely. Finally, sensitivity to certain prior assumptions is investigated in order to examine the impact these assumptions have on estimation results.

Methods

This paper describes the analysis of acoustic survey data from the Bering Sea collected in summer 1994 on the ship "Miller Freeman" by the acoustic survey team of the Alaska Fisheries Science Center (AFSC). The acoustic images were created using a 38 kHz SIMRAD EK-500 split-beam transducer, and were recorded in Bergen Echo Integrator (BEI; Knudsen, 1990) format. Bottom depth was recorded continuously as part of the acoustic data, and temperature bathytherms (XBT/CTD) and trawl data were obtained as part of the survey.

The process of analysing a set of acoustic data with BASCET has 13 steps. These are:

(1) Decide which species or size classes are to be discriminated.
(2) Decide which acoustic school attributes to use in discrimination.
(3) Decide which environmental features to use.
(4) Extract schools and their attributes from backscatter images using image-processing techniques.
(5) Merge schools into school clusters.
(6) Obtain trawl catch by species class.
(7) Identify hotspots with "Hotspotgen" (or manually).
(8) Elicit priors for BASCET parameters.
(9) Determine proposal distributions for MCMC parameters (see Appendix 1).
(10) Run the BASCET estimation procedure.
(11) Investigate convergence of the MCMC algorithm (see Appendix 2).
(12) Test modelling assumptions using the posterior predictive distribution.
(13) Compare BASCET "reasoning" to that of an acoustic expert for selected school clusters.

The following paragraphs are intended to clarify and elaborate on these steps. Appendix 3 may be useful in understanding BASCET jargon.

Figure 1 shows the path of the acoustic survey vessel. The parts of the path on which analysis was conducted are shown as a solid line and the rest of the path as a

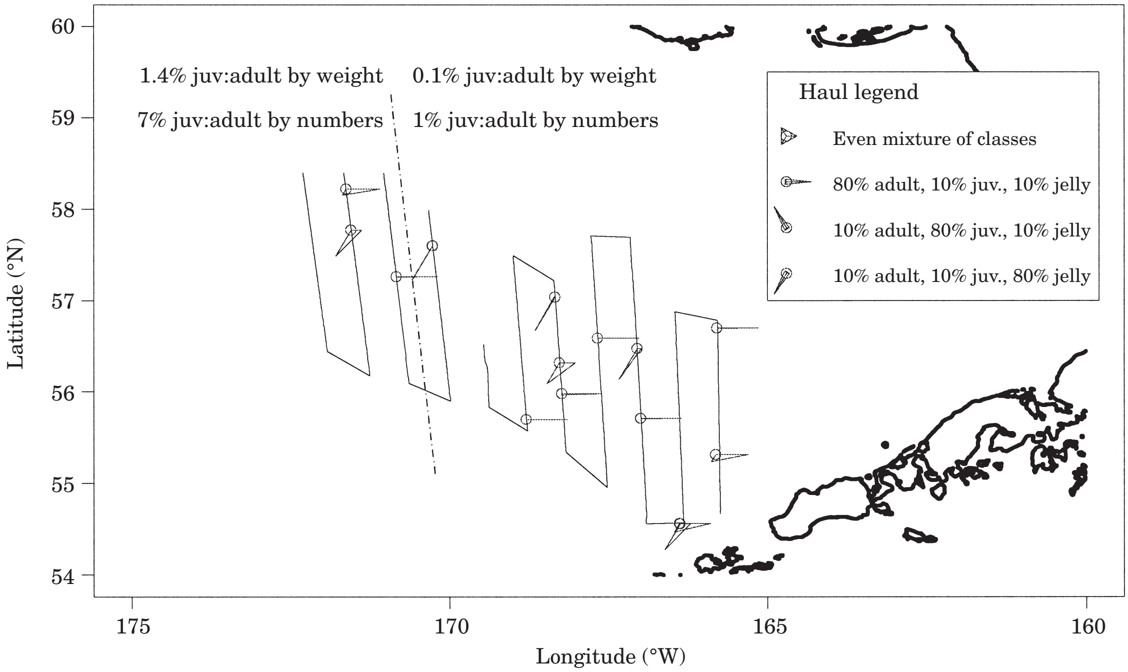

Figure 2. Location of trawl samples taken in the selected survey region and the composition of their catch. After examining these hauls and the acoustic images, AFSC scientists divided the survey region into two along the dashed line shown. On the right side of the line they decided that juvenile-to-adult pollock ratios were 0.01, and, on the left, the ratio was 0.07.

dashed line. Figure 1 also shows an image representing the distribution of BEI-estimated biomass over the selected area. AFSC scientists indicated that these estimates represent the biomass of adult and juvenile pollock, *Theragra chalcogramma* (T. Honkalehto, Alaska Fisheries Science Center, pers. comm.). In order to partition their biomass estimates between these two classes, scientists scrutinized backscatter images in the area, and examined the haul results.

Figure 2 shows the composition of the catch from hauls carried out in the survey region using a star plot. The shape of each plot symbol indicates the dominant class by weight in the catch, as indicated by the haul legend. On the basis of the haul and image information, the survey team decided to divide the region into two subareas along the dashed line shown in Figure 2. To the right of the line, they decided that there was one juvenile pollock for every 100 adults (>30 cm long), and, to the left, that there were seven juveniles for every 100 adults (T. Honkalehto, pers. comm.). Juvenile-to-adult biomass ratios are also shown on the Figure.

The BEI data were converted into a table of school parameters and environmental factors through morphological and other image-resolving techniques (Swartzman *et al.*, 1994). In order to simplify analysis, schools were merged into school clusters, as recommended in Hammond and Swartzman (2001). Following Swartzman (1997), the horizontal range threshold used for this operation was 500 m and the vertical range was 300 m.

The objective of BASCET is to identify the composition of school clusters. In this application, the algorithm was allowed to say that clusters were composed of adult pollock, juvenile pollock or jellyfish, categories suggested by AFSC scientists. In order to use the BASCET algorithm, prior distributions had to be obtained for the estimation model parameters, and the next section of this paper addresses general concerns about eliciting priors from an expert. Thereafter, techniques are suggested that may help elicit priors for BASCET. Proposal distributions were specified for use by the Markov Chain Monte Carlo (MCMC) algorithm that underlies BASCET. Recommendations for doing this are presented in Appendix 1. Finally, the number of iterations of MCMC was selected using the technique described in Appendix 2.

In order to compare BASCET results to the ratios used by AFSC, weighted averages of the school cluster composition estimates in each of the two regions defined by AFSC scientists (Figure 2) were computed. Each cluster composition estimate was weighted by the number of schools in the cluster. These average compositions were called $\mathbf{c}_1$ and $\mathbf{c}_2$ for the eastern and western regions respectively. The juvenile-to-adult pollock ratios in $\mathbf{c}_1$ and $\mathbf{c}_2$ were compared to the corresponding ratios used by AFSC (Figure 2).

Prior elicitation

The prior distribution for a parameter summarizes the information about it from all sources other than the data set itself. There are several pitfalls to watch for when eliciting priors from an expert. Above all, the expert must be interpreting each parameter correctly. For example, consider the problem of asking an angler what the probability is that he will catch a big fish during his holiday. In giving his answer, it is important to know whether or not he is assuming his planned fishing trip actually goes ahead. By raising such complications as these, one can assist an expert to express his or her uncertainty about parameters of interest properly.

Prior distributions elicited from an expert are not guaranteed to be consistent with each other (Punt and Hilborn, 1997). For example, suppose one is asking an angler for two pieces of information: the volume of the average fish he caught and how many of his fish would fit into a specified cooler. The two answers are negatively correlated; one may be subject to a little bias. Therefore, if the angler provides (or is led to provide) independent prior distributions for the two parameters, in some sense the priors will be inconsistent with each other. After all, the total volume of the fish in a cooler cannot exceed the volume of water the cooler can contain.

Eliciting priors for correlated parameters is a difficult task, but there is a technique that some experts find particularly helpful. The trick is to ask the expert to provide a sample from the joint prior distribution. This sample is then used to estimate the parameters of the prior (by maximum likelihood). The sampling approach is also helpful when parameters of the prior distribution do not have obvious physical interpretation.

Configuring BASCET

In this section, important parameters of BASCET are examined one at a time. The parameters are of two types: those for which prior information is required and those for which the user must specify a value (i.e. they are assumed known without error). In order to assist with prior specification, techniques for extracting beliefs about parameters from experts of acoustics are suggested. Then, to assist with setting user-specified parameters, interpretations are provided for them and trade-offs are examined.

The classes

The BASCET algorithm of Hammond and Swartzman (2001) was used to estimate the proportion of adult pollock, juvenile pollock and jellyfish in every school cluster. If any other types of fish appeared in acoustic backscatter images they were forced to fit into these three categories. By the same token, when other fish appeared in the trawl samples, as they occasionally did (<3% of total catch biomass for hauls shown in Figure 2), the catch composition was computed as if they were not present. Each of the candidate groups was called a class, and the classes were indexed by $k=1 \ldots K$. In this example, K was equal to 3. Note that, in this paper, the word "composition" means a vector of length K, whose entries add up to 1.

Hotspots and their composition

The BASCET algorithm can take advantage of the fact that acoustic survey data exhibit spatial and temporal structure. It handles such structure using regions of space-time called hotspots. These hotspots are taken to have a circular shape and finite duration; they are not allowed to overlap in both space and time (when hotspots do overlap in both space and time they are merged). The basic idea is that school clusters located within the same hotspot should have similar composition. The more hotspots there are in the survey region, the more diversity there can be in local species composition.

In BASCET, hotspots also provide the mechanism by which the catch from trawl samples is related to the composition of school clusters. In the algorithm, the catch from a particular haul provides information about the composition of any school clusters in the same hotspot. Thus, in the language of classification literature (James, 1985), the hotspots determine the training set with which the algorithm learns to estimate school cluster composition. Given the separation between trawls and corresponding acoustic images, spatial and temporal assumptions must be made when constructing an acoustic training set.

BASCET uses the hotspotgen algorithm (Hammond and Swartzman, 2001) to place hotspots in the survey region. The size of hotspots placed around both hauls and clusters can be tuned using hotspotgen arguments $R^T_{cluster}$, $R^S_{cluster}$, R^T_{trawl}, and R^S_{trawl}, and sensitivity to these values should be investigated. The more hotspots used, the less need there is to rely on spatial structure in the acoustic data, but the longer the algorithm will take to converge, i.e. the more iterations of MCMC will be required (see Appendix 2). This is because MCMC should take longer to converge when more parameters are being estimated. In order to ensure MCMC convergence (by using a small number of hotspots), the following values were chosen for the Bering Sea data: $R^T_{cluster}=8$ h, $R^S_{cluster}=15$ km, $R^S_{trawl}=25$ km and $R^T_{trawl}=10$ h.

In BASCET, every hotspot has a composition that represents the expected composition of school clusters contained within it (prior to observing cluster attributes). The composition $\mathbf{p}_i$ of hotspot i is assumed to be distributed as follows:

$$\mathbf{p}_i \sim \mathrm{ALN}(\mathrm{alr}(\mathbf{p}_H),\boldsymbol{\Sigma}_H$$

where the alr transformation is defined by:

$$\mathrm{alr}(\mathbf{z}) = \left[\log\left(\frac{z_1}{z_K}\right), \cdots, \log\left(\frac{z_k}{z_K}\right), \cdots, \log\left(\frac{z_{K-1}}{z_K}\right)\right]'$$

and the Additive Logistic Normal (ALN) distribution is as defined in Hammond and Swartzman (2001).

The expected hotspot composition ($\mathbf{p}_H$) was assigned the prior distribution:

$$\pi(\mathbf{p}_H) = \mathrm{ALN}(\boldsymbol{\mu}_H, \boldsymbol{\Phi}_H)$$

BASCET requires the ALN distribution solely for the sake of computational convenience. As $\mathbf{p}_H$ is a composition vector, obtaining a prior distribution for it from an expert is challenging. It can be helpful to ask the expert to create a sample of size 10–20 from the prior distribution. In this case, the expert was asked to identify 10 plausible values for the expected composition. These were centred around 54% adult pollock, 10% juvenile pollock, and 36% jellyfish, and $\boldsymbol{\mu}_H$ and $\boldsymbol{\Phi}_H$ were estimated from this sample. The results were:

$$\mu_H = \begin{bmatrix} 0.3918 \\ -1.2966 \end{bmatrix}$$

$$\Phi_H = \begin{bmatrix} 0.9896 & 0.8987 \\ 0.8987 & 1.0768 \end{bmatrix}$$

The same sampling technique was used to determine the prior for Σ_H, which controls the variability in hotspot composition. The expert was asked to imagine that the true $\mathbf{p}_H$ was equal to the expected value of the samples given in the previous step. Then he was asked to provide a second sample of ten hotspot compositions to represent how variable the individual hotspot compositions were about their global mean. A degrees-of-freedom value is also required. The prior distribution used in all computations was:

$$\pi(\Sigma_H^{-1}) = \text{Wishart } (15, \Omega_H)$$

$$\Omega_H = \begin{bmatrix} 0.1187 & -0.0647 \\ -0.0647 & 0.0724 \end{bmatrix}$$

The Wishart distribution is described in Gelman *et al.* (1995).

Acoustic attributes

Acoustic attributes of school clusters help the algorithm to discriminate between classes. Three acoustic attributes were used: mean within-school volume backscatter cross-section (S_v), range of within-school S_v, and school area. These features were placed into vectors denoted by **f** and indexed by a = 1 . . . 3. Thus, f_1 would refer to a measurement of mean within-school S_v, and f_3 would refer to a measurement of school area. The BASCET estimation model assumes that the distribution for acoustic feature f_a, for a cluster of composition **c**, is Normal with mean $\Sigma_{k=1}^{K} c_k \mu_{a,k}$ and variance σ_a^2 (K being the number of classes, 3 in this application). The assumption of normality should pose little problem when there are many schools in each cluster.

The interpretation of the nine $\mu_{a,k}$ values is best explained using examples: $\mu_{1,3}$ refers to the mean within-school S_v for jellyfish, whereas $\mu_{3,1}$ would be the mean school area of adult pollock. Each of these values was assigned a prior distribution as follows:

$$\pi(\mu_{a,k}) = N(\theta_{a,k}, \tau_{a,k}^2).$$

The three σ_a^2 values affect the importance of their respective acoustic attributes in discriminating between classes: larger values assign less weight to an attribute. The σ_a^2 values were assigned prior distributions from the inverse Gamma family (Gelman *et al.*, 1995):

$$\pi(\sigma_a^2) \sim IG(\alpha_a, \beta_a)$$

When setting prior parameters as described below, it is useful to know that $\pi(\sigma_a^2)$ is parameterized to have mean $((\alpha_a - 1)\beta_a)^{-1}$ and variance $((\alpha_a - 1)^2 (\alpha_a - 2)\beta_a)^{-1}$.

The prior parameters $\theta_{a,k}, \tau_{a,k}^2, \alpha_a$, and β_a must be specified, but there is no need to ask for them directly. The best way to obtain these may be to consult with a school-identification expert. First, select a number of backscatter images at random from a previous survey year or a nearby region. Have the expert identify the composition of the schools in these images and treat the identifications as being absolutely correct. Then one could use the acoustic attributes of the identified schools to estimate the required parameters.

The prior parameters used in this paper are listed below:

$$[\theta_{a,k}] = \begin{bmatrix} -43 & -45 & -55 \\ 2.5 & 3.0 & 1.5 \\ 1500 & 2500 & 1000 \end{bmatrix}$$

$$[\tau_{a,k}] = \begin{bmatrix} 5 & 5 & 2.5 \\ 0.5 & 1.0 & 0.4 \\ 300 & 350 & 100 \end{bmatrix}$$

In the matrices above, the rows indicate the acoustic parameters (average S_v in dB, S_v range in dB, and school area in m^2) while the columns indicate the classes (adult pollock, juvenile pollock, and jellyfish).

The parameters of an inverse gamma distribution should not be elicited directly. It is preferable to ask for the mean and SD and then derive the α_a and β_a from these. An expert might draw histograms of credible σ_a^2 values from which the mean and SD can be determined approximately. Once the mean and SD are obtained, the two equations provided above for the mean and variance of an inverse gamma distribution determine α_a and β_a uniquely. The numbers used in this paper are as follows:

$$[\alpha_a] = \begin{bmatrix} 10.000 \\ 2.977 \\ 8.250 \end{bmatrix}$$

$$[\beta_a] = \begin{bmatrix} 0.006 \\ 0.202 \\ 5.517 \times 10^{-07} \end{bmatrix}$$

As before, the rows indicate the acoustic parameter.

Trawl sample parameters

The treatment of trawl data in BASCET incorporates a weight-to-numbers conversion vector ($\mathbf{m}$) of length K. The entries of $\mathbf{m}$ represent the average weights of individual fish of each class. Given that the objective is identifying the species composition of clusters, the relative size of entries in $\mathbf{m}$ is more important that the absolute size. The values used in estimation were:

$$\mathbf{m} = \begin{bmatrix} 100 \\ 50 \\ 70 \end{bmatrix}$$

The BASCET algorithm does not admit any uncertainty in $\mathbf{m}$, so results are conditional on the values chosen. Whenever catch data are available as numbers of fish by class (as opposed to weight by class), the use of $\mathbf{m}$ can be dispensed with altogether.

BASCET uses selectivity vectors to account for the fact that the species that appear in backscatter images are likely to be different from those caught in trawls. There is one such vector for every type of gear in use (gears are indexed by $g=1 \ldots G$). These selectivity vectors represent the relative vulnerability of the fish observed in backscatter images to the trawl gears. They are vectors of length K whose entries sum up to 1. It must be stressed here that selectivity is relative to what appears in the acoustic images and not to what is in the water. The distinction is important because the acoustic equipment is itself highly selective.

Jellyfish are considerably more likely to appear in the trawl samples than in morphologically identified schools because they have low target strength at 38 kHz. With this tendency in mind, the gear selectivity vectors ($\mathbf{s}_g$, for $g=1$ to G) were set to the value:

$$\mathbf{s}_g = \begin{bmatrix} 0.03 \\ 0.02 \\ 0.95 \end{bmatrix}$$

Results should be interpreted as being conditional on the selectivity vectors chosen (by expert consultation).

Environmental features

Two environmental features were used for school cluster identification purposes: ocean depth and temperature at the cluster location. A prior for the composition vector ξ_{temp} that represents the change in composition one would expect with a one-degree increase in the local temperature was required. This vector works as follows: suppose a school cluster has composition $\mathbf{c}_1$ at temperature t_1, and suppose that the temperature changes to t_2, then the new expected composition $\mathbf{c}_2$ would be derived as follows:

$$\mathrm{alr}(\mathbf{c}_2) = \mathrm{alr}(\mathbf{c}_1) + (t_2 - t')\mathrm{alr}(\xi_{temp})$$

A prior distribution was assigned to ξ_{temp} from the ALN family for the sake of computational convenience. As ξ_{temp} is a composition vector, extracting a prior for it requires techniques similar to those used to determine mean hotspot composition priors.

To obtain a prior for ξ_{temp}, the expert might be told the average temperature for schools in the survey region. Then one might ask the expert to imagine a volume of water at that mean temperature that happens to contain some fish. Next one would ask the expert to assume that the composition of classes in the volume is known to be $\mathbf{c}$. The expert would then be asked to imagine how this composition would change in water 1° warmer. Finally, he could be asked to provide a sample of plausible values for the new composition, from which a prior for ξ_{temp} can readily be determined.

The prior for ξ_{temp} used in this paper is as follows:

$$\pi(\xi_{temp}) = \mathrm{ALN}(\mathrm{alr}(\bar{\boldsymbol{\omega}}_{temp}), \boldsymbol{\Psi}_{temp})$$

$$\bar{\boldsymbol{\omega}}_{temp} = \begin{bmatrix} 0.187 \\ 0.407 \\ 0.406 \end{bmatrix}$$

$$\boldsymbol{\Psi}_{temp} = \begin{bmatrix} 0.244 & 0.108 \\ 0.108 & 0.150 \end{bmatrix}$$

The other environmental feature was ocean depth. As this was thought *a priori* to have little or no effect on expected cluster composition, the following prior was chosen:

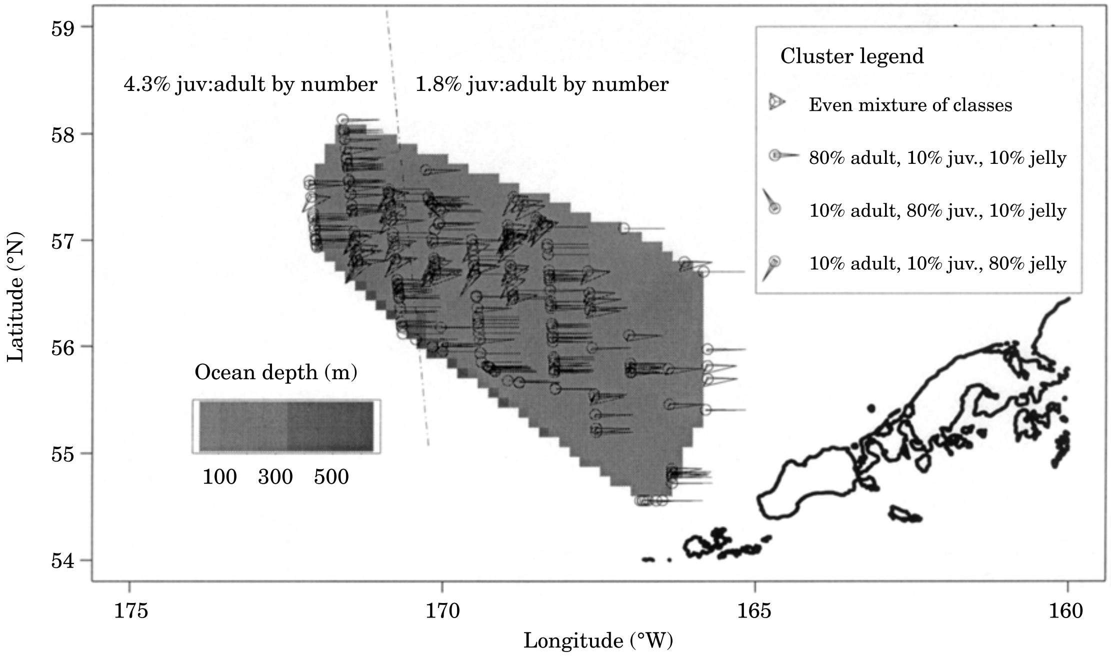

Figure 3. Location and composition of school clusters, as estimated by the BASCET estimation algorithm. When adult and juvenile pollock proportions were averaged over the regions separated by the broken line, the eastern area had 1.8 juveniles for every 100 adults and the western area had 4.3 juveniles per 100 adults.

$$\pi(\xi_{\text{depth}}) = \text{ALN}(\text{alr}(\bar{\boldsymbol{\omega}}_{\text{depth}}), \boldsymbol{\Psi}_{\text{depth}})$$

$$\bar{\boldsymbol{\omega}}_{\text{depth}} = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}$$

$$\boldsymbol{\Psi}_{\text{depth}} = \begin{bmatrix} 0.0002 & 0.0001 \\ 0.0001 & 0.0002 \end{bmatrix}$$

Depth preference parameters

In BASCET, depth distributions of the various classes are expressed as Beta distributions delimited by the local ocean depth. As the parameters of a Beta distribution are somewhat abstract, it may be best to obtain priors for them using the sampling approach. First, one would ask the expert at what percentage of local ocean depth each class of fish prefers to swim. Uncertainty in these preferences could be indicated using a sample of depth preferences for each species class. In this case, it was suggested that juvenile pollock prefer to swim at around 20% of the local ocean depth, jellyfish at about 25%, and adult pollock at about 90%. These preferences were considered accurate to within ±15%.

Information about the variability in depth distribution is also needed. To obtain this, the expert might be asked to assume that the expected preferences specified above were correct. Using the assumed preferences, the expert might then provide a sample showing the variability in school depth distribution about the preferred levels. In the Bering Sea, the expert felt that adult pollock were more widely distributed about their preferred depths than the other classes. For the adults, the variance in school depth divided by ocean depth was thought to be about 0.1 ± 0.05, whereas the other classes were assigned a variance of about 0.05 ± 0.03.

Parameters of the Beta distribution can then be reconstructed from the samples provided. To do this for a particular class, one can use each mean and variance sample value to compute a corresponding set of Beta distribution parameters. For example, if the mean sample is $\mu_{d,k}$ and the variance sample $\sigma^2_{d,k}$ for class k, then the Beta distribution parameters $\alpha_{d,k}$ and $\beta_{d,k}$ can be determined from the following equations:

$$\mu_{d,k} = \frac{\alpha_{d,k}}{\alpha_{d,k} + \beta_{d,k}}$$

and

$$\sigma^2_{d,k} = \frac{\alpha_{d,k}\beta_{d,k}}{(\alpha_{d,k} + \beta_{d,k})^2(\alpha_{d,k} + \beta_{d,k} + 1)}$$

The mean and variance samples thus determine a sample of $\alpha_{d,k}$ and $\beta_{d,k}$ values that may, in turn, be used to construct the priors.

The depth distribution priors used in this paper (1 is adult pollock, 2 juvenile pollock, 3 jellyfish) were:

Table 1. Effects of different user-specified parameter values on estimation results in the eastern and western regions of the survey area (defined in Figure 3). The values in $\mathbf{c}_1$ and $\mathbf{c}_2$ represent the average composition in the eastern and western regions respectively and r_1 and r_2 the juvenile-to-adult pollock ratios in $\mathbf{c}_1$ and $\mathbf{c}_2$ respectively. The compositions are indexed in the class order adult pollock, juvenile pollock, jellyfish.

Scenario	User-specified parameter values	$\mathbf{c}_1$	$\mathbf{c}_2$	r_1	r_2
Base case (hotspots as in Figure 4)	As defined in paper	$\begin{bmatrix}0.880\\0.016\\0.104\end{bmatrix}$	$\begin{bmatrix}0.786\\0.034\\0.180\end{bmatrix}$	0.018	0.043
Six hotspots	As defined in paper, but change random seeds for hotspotgen algorithm	$\begin{bmatrix}0.917\\0.025\\0.058\end{bmatrix}$	$\begin{bmatrix}0.818\\0.049\\0.132\end{bmatrix}$	0.027	0.060
13 hotspots	$R^S_{cluster}=10$ km $R^T_{cluster}=6$ h $R^S_{trawl}=20$ km $R^T_{trawl}=12$ h	$\begin{bmatrix}0.919\\0.017\\0.064\end{bmatrix}$	$\begin{bmatrix}0.822\\0.038\\0.140\end{bmatrix}$	0.018	0.046
Gear selectivity modification	$\mathbf{s}_g=\begin{bmatrix}0.06\\0.03\\0.91\end{bmatrix}$	$\begin{bmatrix}0.853\\0.014\\0.133\end{bmatrix}$	$\begin{bmatrix}0.751\\0.031\\0.218\end{bmatrix}$	0.016	0.041
Gear selectivity modification	$\mathbf{s}_g=\begin{bmatrix}0.1\\0.1\\0.8\end{bmatrix}$	$\begin{bmatrix}0.823\\0.010\\0.167\end{bmatrix}$	$\begin{bmatrix}0.721\\0.023\\0.256\end{bmatrix}$	0.012	0.032
Gear selectivity modification	$\mathbf{s}_g=\begin{bmatrix}0.02\\0.03\\0.95\end{bmatrix}$	$\begin{bmatrix}0.891\\0.013\\0.096\end{bmatrix}$	$\begin{bmatrix}0.805\\0.029\\0.166\end{bmatrix}$	0.015	0.036
Weight-to-numbers conversion vector modification	$\mathbf{m}=\begin{bmatrix}90\\30\\60\end{bmatrix}$	$\begin{bmatrix}0.811\\0.016\\0.173\end{bmatrix}$	$\begin{bmatrix}0.706\\0.031\\0.263\end{bmatrix}$	0.020	0.044

$$\pi(\alpha_{d,1})=\mathrm{LogNormal}(0.45,\ 0.28)$$

$$\pi(\alpha_{d,2})=\mathrm{LogNormal}(-1.03,\ 0.64)$$

$$\pi(\alpha_{d,3})=\mathrm{LogNormal}(-0.25,\ 0.62)$$

$$\pi(\beta_{d,1})=\mathrm{LogNormal}(-1.03,\ 0.59)$$

$$\pi(\beta_{d,2})=\mathrm{LogNormal}(0.70,\ 0.33)$$

$$\pi(\beta_{d,3})=\mathrm{LogNormal}(0.98,\ 0.36)$$

The first and second parameters in the specification of each LogNormal distribution are the mean and SD of the log transformed random variable. In contrast to other prior models in BASCET, the LogNormal form is not required in this instance (because the depth parameters were not Gibbs sampled; see Appendix 1 of Hammond and Swartzman, 2001).

Methods for model checking

Four different types of model checking are employed in this paper: sensitivity analysis for selected prior distributions and user-specified parameters, comparison of prior and posterior results, simulation of test quantities from the posterior predictive distribution, and justification of composition estimates for selected clusters. Of these methods, the last two require additional explanation.

Testing with the posterior predictive distribution

The posterior predictive distribution is the distribution one would use to predict new data $\mathbf{y}^{rep}$ after observing the real data $\mathbf{y}$. For parameter vector $\boldsymbol{\theta}$, the posterior predictive distribution is given by the equation:

$$p(\mathbf{y}^{rep}|\mathbf{y})=\int p(\mathbf{y}^{rep}|\boldsymbol{\theta})p(\boldsymbol{\theta}|\mathbf{y})d\boldsymbol{\theta}$$

Table 2. Effects of different prior distributions on estimation results in the eastern and western regions of the survey area (defined in Figure 3). The values in c_1 and c_2 represent the average composition in the eastern and western regions respectively and r_1 and r_2 the juvenile-to-adult pollock ratios in $\mathbf{c}_1$ and $\mathbf{c}_2$ respectively. The compositions are indexed in the class order adult pollock, juvenile pollock, jellyfish.

Scenario	User-specified parameter values	$\mathbf{c}_1$	$\mathbf{c}_2$	r_1	r_2
Increase prior uncertainty about ξ_{temp}	$\Psi_{temp} = 1.5\begin{bmatrix} 0.244 & 0.108 \\ 0.108 & 0.150 \end{bmatrix}$	$\begin{bmatrix} 0.884 \\ 0.015 \\ 0.100 \end{bmatrix}$	$\begin{bmatrix} 0.793 \\ 0.033 \\ 0.173 \end{bmatrix}$	0.018	0.042
Increase prior uncertainty about depth distribution parameters $\alpha_{d,k}$ and $\beta_{d,k}$ for $k=1 \ldots K$	$\pi(\alpha_{d,1})=\text{LogNormal}(0.45, 0.42)$ $\pi(\alpha_{d,2})=\text{LogNormal}(-1.03, 0.96)$ $\pi(\alpha_{d,3})=\text{LogNormal}(-0.25, 0.92)$ $\pi(\beta_{d,1})=\text{LogNormal}(-1.03, 0.89)$ $\pi(\beta_{d,2})=\text{LogNormal}(0.70, 0.50)$ $\pi(\beta_{d,3})=\text{LogNormal}(0.98, 0.54)$	$\begin{bmatrix} 0.890 \\ 0.014 \\ 0.096 \end{bmatrix}$	$\begin{bmatrix} 0.805 \\ 0.028 \\ 0.166 \end{bmatrix}$	0.016	0.035
Increase prior uncertainty about σ_a^2 values	$\pi(\sigma_s^2)=\text{IG}(\alpha_a,\beta_a)$ $\alpha_a = \begin{bmatrix} 7.714 \\ 2.679 \\ 2.625 \end{bmatrix}$ $\beta_a = \begin{bmatrix} 0.007 \\ 0.238 \\ 2.462 \times 10^{-6} \end{bmatrix}$	$\begin{bmatrix} 0.881 \\ 0.017 \\ 0.103 \end{bmatrix}$	$\begin{bmatrix} 0.789 \\ 0.035 \\ 0.176 \end{bmatrix}$	0.019	0.044

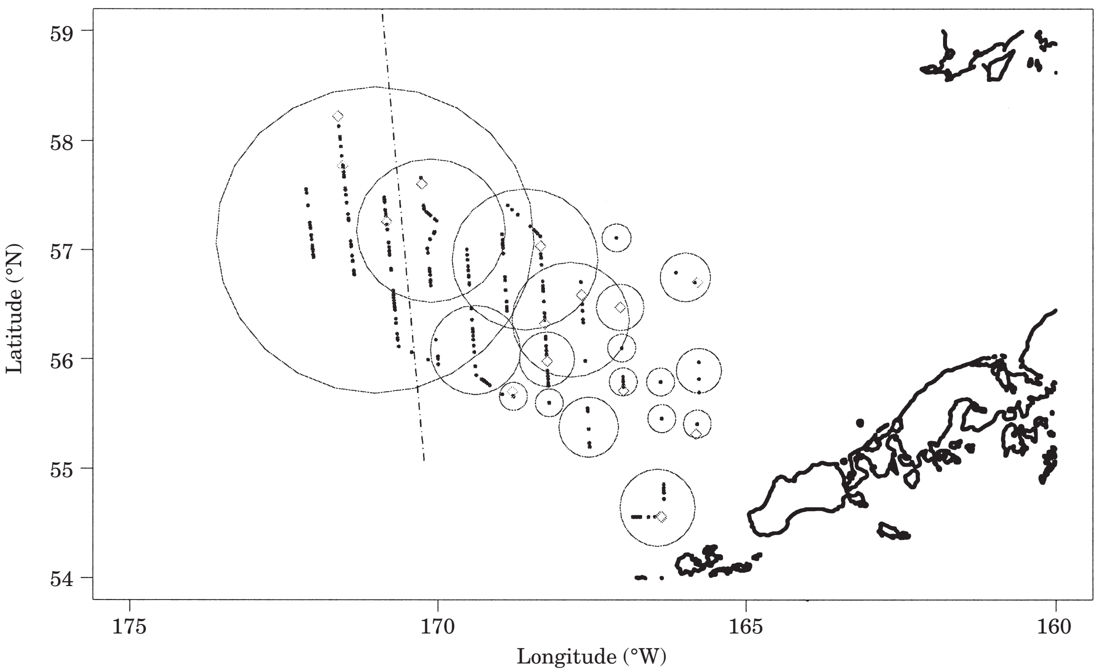

Figure 4. Hotspot locations used in the analysis, indicated by circles. The haul start locations are indicated by diamonds and school cluster locations by black dots.

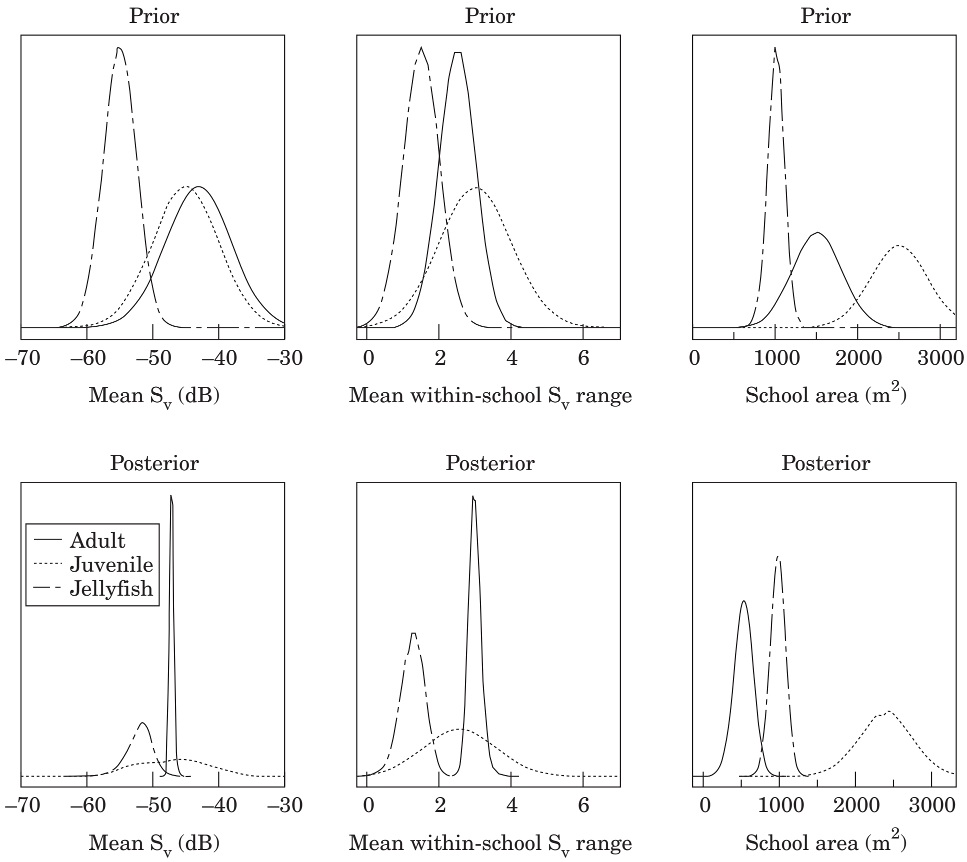

Figure 5. Prior and posterior distributions for acoustic feature means.

Following Gelman *et al.* (1995), the discrepancy between the model and data is measured using test quantities $T(y,\theta)$, which are scalar summaries of parameters and data. Test quantities play the same role as test statistics do in frequentist hypothesis testing. The basic idea is to compute the Bayes p-value for a particular test quantity, as follows:

$$\text{Bayes p-value} = \Pr(T(\mathbf{y}^{rep},\boldsymbol{\theta}) \geq T(\mathbf{y},\boldsymbol{\theta})|\mathbf{y})$$

where the probability is taken over the joint posterior distribution of θ and $\mathbf{y}^{rep}$. If the Bayes p-value is close to either 0 or 1, it suggests model failure. As in frequentist statistics, how close the Bayes p-value must be to these extremes before one concludes there is "significant" model failure is arbitrary. Here, values within 2.5% of the extremes were interpreted as cause for concern.

In practice, one draws samples from the posterior predictive distribution using simulation, as integration is often intractable. Given S samples ($\boldsymbol{\theta}^1 \ldots \boldsymbol{\theta}^S$) from the posterior distribution, one draws one value of $\mathbf{y}^{rep}$ from $p(\mathbf{y}|\boldsymbol{\theta}^s)$ for each sample $\boldsymbol{\theta}^s$, and then computes $T(\mathbf{y},\boldsymbol{\theta}^s)$ and $T(\mathbf{y}^{rep},\boldsymbol{\theta}^s)$. The Bayes p-value may then be estimated by computing the proportion of times in S samples that $T(\mathbf{y}^{rep},\boldsymbol{\theta}^s) \geq T(\mathbf{y},\boldsymbol{\theta}^s)$.

The models for the acoustic data were evaluated using test quantities. One test quantity was used for each acoustic feature. In order to define the test quantity for the acoustic data model, some notation will be introduced. For a particular set of model parameters, let the composition of school cluster j be $\mathbf{c}_j$. Under the BASCET model, the distribution of acoustic feature f_a in the cluster is supposed to be Normal with mean $\Sigma_{k=1}^{K} c_k \mu_{a,k}$ and variance σ_a^2.

Upon simulating a value f_a^{rep} from this distribution, one can compute the following statistics for the cluster:

$$d_{a,j}^{rep} = \frac{\left(f_a^{rep} - \sum_{k=1}^{K} c_k \mu_{a,k}\right)^2}{\sigma_a^2} \text{ and } d_{a,j} = \frac{\left(f_a - \sum_{k=1}^{K} c_k \mu_{a,k}\right)^2}{\sigma_a^2}$$

The test quantities for the acoustic features are found by summing over the school clusters: $T_a^{rep} = \Sigma_{j=1}^{J} d_{a,j}^{rep}$ and $T_a = \Sigma_{j=1}^{J} d_{a,j}$. The Bayes p-value for each acoustic feature is computed by examining the proportion of posterior parameter simulations under

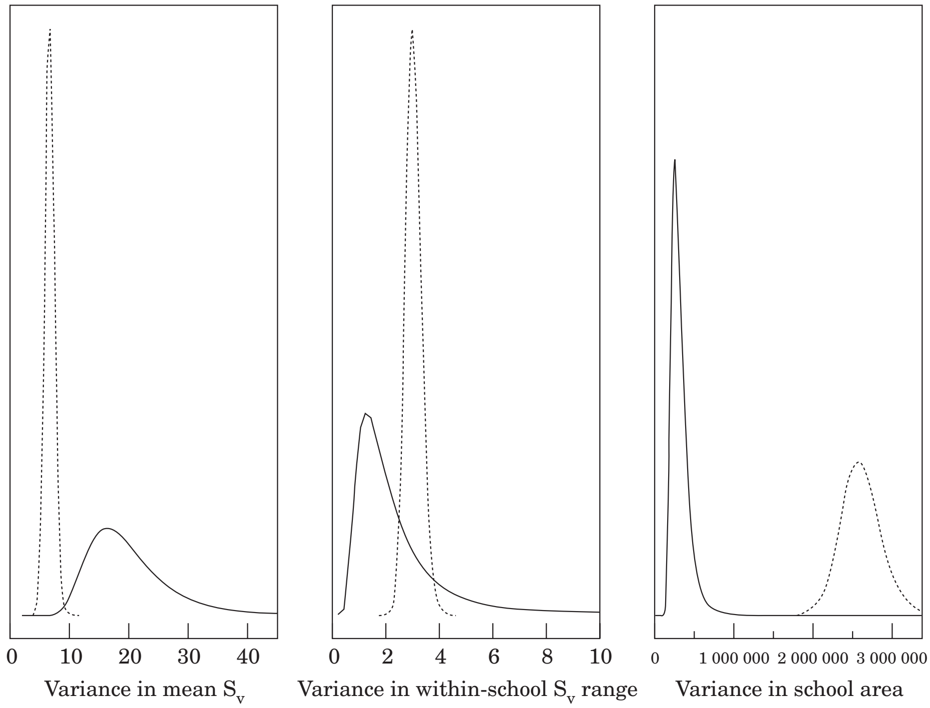

Figure 6. Prior (solid line) and posterior (broken line) distributions for acoustic feature variances.

which $T_a^{rep} \geq T_a$. If there appears to be model failure, one should investigate whether the predicted variance in any particular cluster (assumed to be σ_a^2 for all clusters) actually depends on either the number of schools in the cluster or the cluster composition. If so, appropriate changes to the model may be required.

A different test quantity was employed to evaluate the trawl sample model. Let us assume that, under a particular set of model parameters, the expected (numeric) composition of haul l is $\mathbf{c}_{T,l}$ (Hammond and Swartzman, 2001), the weight-to-numbers conversion vector is $\mathbf{m}$ (see above), and the observed vector of catch weights by class is $\mathbf{w}_l$. Under BASCET assumptions, when $\mathbf{w}_l$ is divided elementwise by $\mathbf{m}$, the resulting vector $\mathbf{z}_l$ is supposed to be distributed multinomially:

$$\mathbf{z}_l \sim \text{Multinomial}(Z_l, \mathbf{c}_{T,l})$$

for

$$Z_l = \sum_{k=1}^{K} z_{l,k} \text{ (Hammond and Swartzman, 2001)}$$

Thus, one can compute Bayes p-values using $\mathbf{z}_l^{rep}$ vectors drawn randomly from the Multinomial distribution above. In the Bering Sea application, interest centres on the adult-to-juvenile pollock ratio, so the test quantity was based on the following values:

$$d_l^{rep} = (\log(z_{l,adult}^{rep}) - \log(z_{l,juvenile}^{rep}))^2$$

and

$$d_l = (\log(z_{l,adult}) - \log(z_{l,juvenile}))^2$$

These values were summed over all the hauls giving $T_{hauls}^{rep} = \Sigma_{l=1}^{L} d_l^{rep}$ and $T_{hauls} = \Sigma_{l=1}^{L} d_l$. From these two test quantities, the Bayes p-value can be computed as above. If there appears to be model failure, one should consider altering the user-specified gear selectivity vectors $\mathbf{s}_g$ or the weight conversion vector $\mathbf{m}$.

Justification of composition estimates

Perhaps the greatest strength of the BASCET approach is its ability to "justify" the composition estimates it makes. Such justification involves an exploration of the cluster attributes most important in determining the final result. Methods based on neural networks are not designed to facilitate such exploration. Central to this advantage is an ability to provide posterior confidence regions for all composition estimates, a task that neural networks cannot achieve. These regions indicate how certain the algorithm is about each cluster composition point estimate.

Justification also relies on the fact that BASCET can indicate posterior credible regions for composition estimates that exclude the information in any particular set of cluster attributes. For example, the algorithm could provide a posterior credible region for the composition

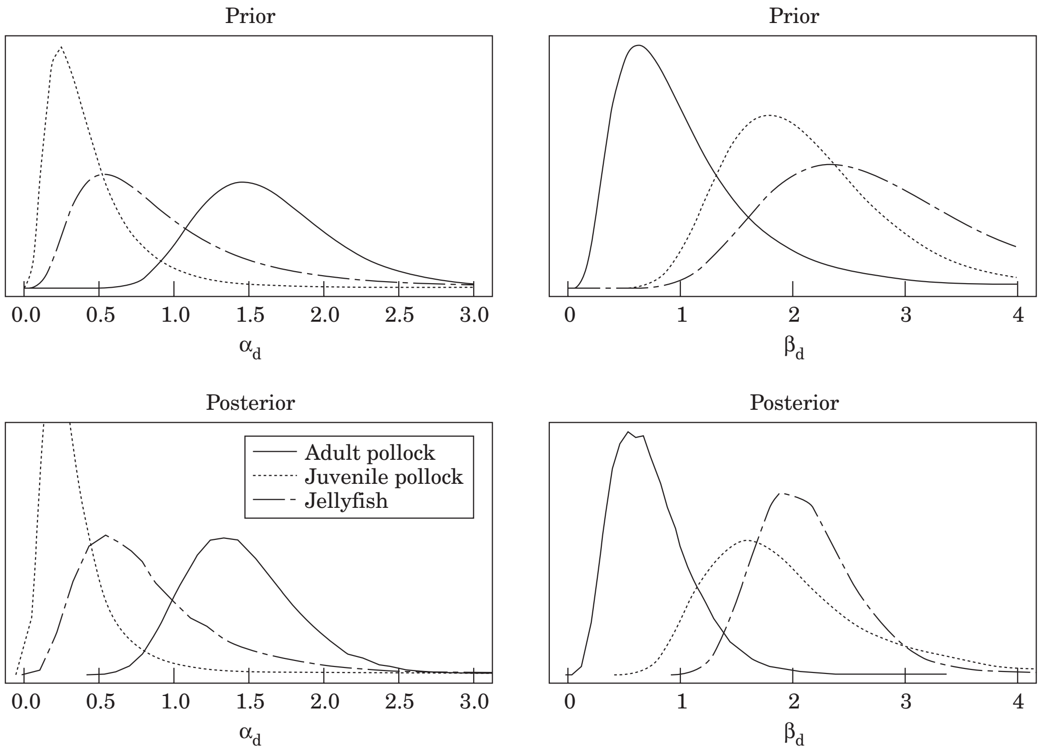

Figure 7. Prior and posterior results for depth distribution parameters. The parameters determine the Beta distribution that describes the depth preferences of each class. The mean of a Beta distribution is $\alpha_d/(\alpha_d+\beta_d)$.

of a particular school cluster that uses all acoustic features of the cluster, but none of its environmental attributes. Therefore, the user can investigate which particular attribute was most decisive in determining the final answer. Furthermore, such analysis can determine which cluster attributes are pointing towards the same dominant class and which are suggesting that the cluster may be composed primarily of something else entirely. The ability to examine the algorithm's "reasoning" in this manner allows comparison with the way a human expert reasons. Such comparison can provide a powerful tool for model checking because it allows one to investigate whether, when expert and algorithm agree on an answer, they also agree on the reason for that answer.

Results

Figure 3 shows the locations of all the school clusters whose composition was estimated using BASCET. The estimated composition of these clusters is indicated using a star plot. An image showing the ocean depth in the area is also given. BASCET predicted that, in the eastern region, the juvenile-to-adult pollock ratio (by numbers) was 1.8%, while the ratio in the western region was 4.3%. AFSC figures were 1% and 7%, respectively. Table 1 indicates the changes in these predictions under different choices for user-specified parameters, and Table 2 shows the same for different priors. Figure 4 shows the locations of all the school clusters and trawl samples observed over the course of the Bering Sea acoustic survey. It also shows the location of all hotspots resulting from the hotspotgen algorithm (Hammond and Swartzman, 2001).

Bayes p-values for the three acoustic feature test quantities were 0.90, 0.46, and 0.31 for average S_v, S_v range, and school area respectively. For the trawl sample test quantity, the Bayes p-value was 0.59. None of these results suggests significant model failure.

Figure 5 shows prior and posterior results computed for the acoustic feature means ($\mu_{a,k}$). The BASCET algorithm learned a considerable amount about the adult pollock parameters from the data, but the juvenile pollock parameters resemble prior values. This tendency is explained by the relatively low representation of juvenile pollock in the composition estimates. The problem occurs also, but to a lesser extent, for jellyfish.

The prior and posterior results for acoustic feature variances (σ_a^2) are shown in Figure 6. This figure indicates that the posterior distributions for average S_v and school area are in the tails of their respective prior distributions. In other words, the results suggest that

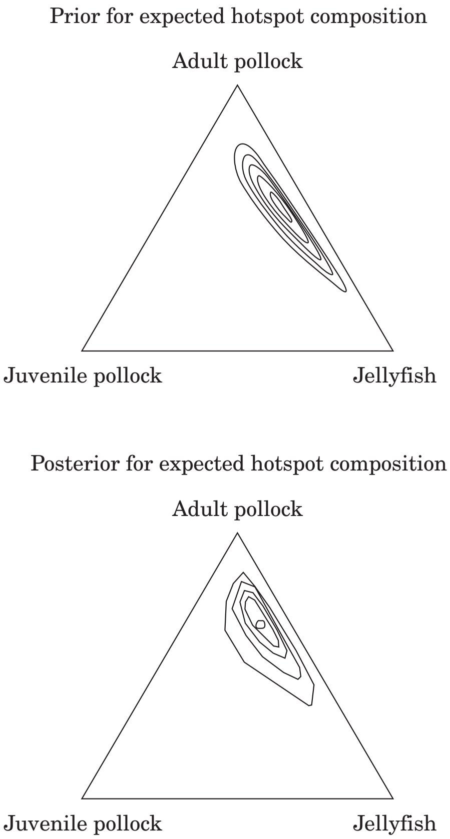

Figure 8. Prior and posterior results for expected hotspot composition.

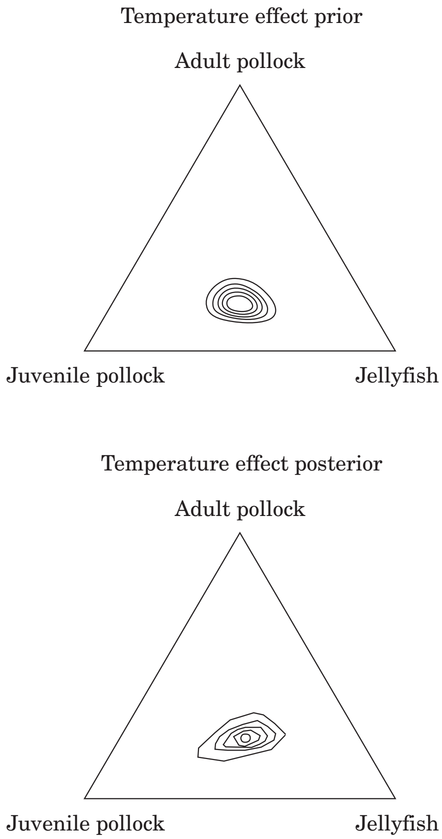

Figure 9. Temperature effect prior and posterior.

unwarranted precision was assigned to the prior distributions for these quantities. The impact of widening these particular prior distributions was therefore examined in Table 2.

Figure 7 shows the prior and posterior distributions for cluster depth related parameters ($\alpha_{d,k}$ and $\beta_{d,k}$). For the most part these appear little changed from the prior values. The effect of increasing prior uncertainty about these parameters was examined in Table 2.

The prior and posterior results for expected hotspot composition ($\mathbf{p}_H$) are revealed in Figure 8 using triangle plots. Triangle plots are used in this paper to display composition estimates. Each vertex of the triangle represents a pure composition, whereas the centre indicates an even mixture of classes. The prior is indicated using contour lines, defined using 95%, 75%, 50%, 25%, and 5% credible regions. The posterior is shifted slightly towards the adult pollock, suggesting that these were more abundant than was thought *a priori*.

Figure 9 shows the prior and posterior results for the composition vector that defines the effect of temperature on cluster composition (ξ_{temp}). While there is little shift in location between the prior and posterior, the shape of the posterior distribution is slightly different. The posterior implies that increases in temperature favour jellyfish a bit less than was thought *a priori*.

The process of justifying composition estimates on a school cluster chosen at random (number 78) is illustrated in Figure 10. The first triangle plot suggests the cluster is composed of a mixture of jellyfish and adult pollock (as shown by the 50%, 25%, and 5% credible regions), but it also indicates a good deal of uncertainty. According to these results, it is unlikely that the cluster could be composed of more than 40% juvenile pollock.

The next triangle plot in Figure 10 shows the composition of the hotspot that contains cluster 78. This hotspot suggests that jellyfish and adult pollock represent the dominant classes in the cluster, which the algorithm suspects upon examining other clusters in the hotspot. This triangle plot also indicates that the hotspot prediction is a rather uncertain one, although the possibility that the cluster might be dominated by juveniles is already ruled out. The next six triangle plots indicate how the hotspot prediction is refined upon examining different cluster attributes (one at a time). The average school area, the cluster depth (44 m) and the average S_v

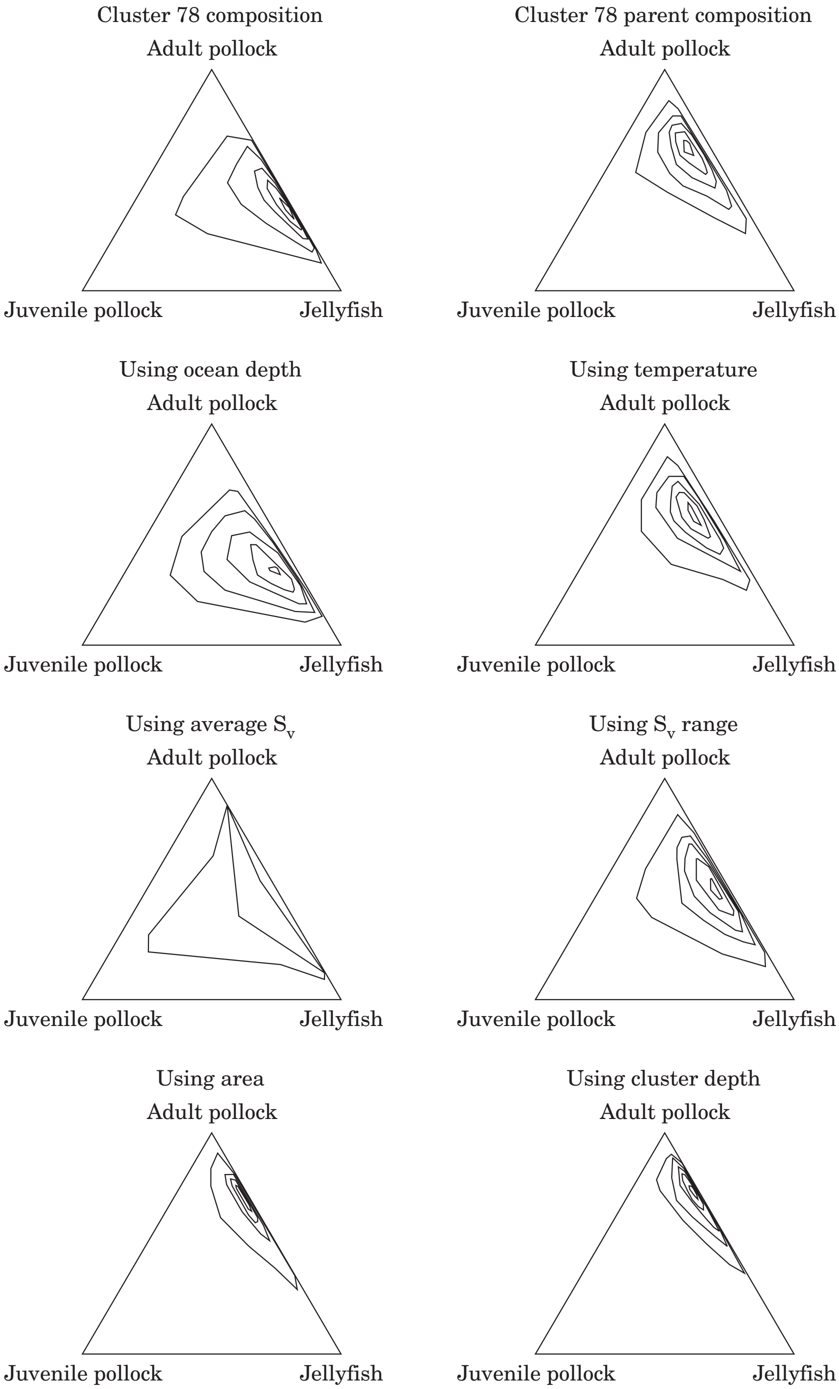

Figure 10. Illustration of the ''justification of cluster composition estimates''. The cluster depth is 43.6 m, the average S_v − 51.4 dB, the S_v range 1.03 dB, the area 94.38 m^2, the ocean depth 79.7 m, and the temperature is 4.18°C. The ''parent composition'' refers to the composition of the hotspot containing this cluster.

all suggest that there are few juvenile pollock. The relatively shallow ocean depth (80 m), the S_v range and the low average S_v (− 51.4 dB) suggest there may be a high proportion of jellyfish.

Discussion

BASCET provided results fairly close (2% vs. 1% and 4% vs. 7%) to the values chosen by AFSC scientists, when the results were compared at AFSC resolution. However, BASCET is able to provide far more detail in its answers, making them more useful than the averages currently used by AFSC. The application documented here provides no reason to suspect that the performance results attributed to the algorithm on simulated data (Hammond and Swartzman, 2001) might be lost if the algorithm were applied to other real surveys.

The ability to reproduce results is, of course, an important component of the scientific method. BASCET represents a step forward for acoustic survey science because it enables survey technicians to compute reproducible results (from given parameters), something that cannot be expected from current *ad hoc* methods. It is doubtful, for example, that acoustic technicians from another laboratory could reproduce the AFSC juvenile to adult pollock ratios from the same data. Moreover, constructing the training set required by most classification algorithms from trawl data is not a straightforward process (Hammond and Swartzman, 2001). BASCET is the first algorithm to propose a general procedure for such construction. Therefore, the higher standard of clarity about training set construction in BASCET must be maintained in acoustic survey papers for them to be widely accepted as appropriate scientific procedures.

Unlike all other school classification algorithms to date, BASCET is Bayesian in inspiration and flavour. This means that it can readily incorporate expertise developed by acoustic scientists and fishers. It also means that it must do so. The ability to incorporate human expertise is valuable because it implies that the algorithm can function even when trawl data are missing altogether, a situation that would completely defeat most competing algorithms. This implies that BASCET has a higher degree of generality in its possible applications. As long as the information elicited from human experts is accurate, BASCET should also perform better than competing algorithms (as it has more information available, though, admittedly, no challengers have tested this claim).

BASCET is designed to indicate not only how confident it is about the estimates it makes, but also to justify those answers if they are challenged. In other words, the algorithm is designed to explain its answers, not as a human expert would, but nonetheless in a manner that provides insight into the way the algorithm weighs different sources of information. These capabilities set BASCET apart from "black box" algorithms based on neural networks.

Acknowledgements

A Washington Sea Grant obtained by Gordon Swartzman and Don Gunderson supported the initial research presented in this document. The writing of this manuscript and the majority of the research were funded by the Center for Environment, Fisheries and Aquaculture Sciences in Lowestoft. We thank Neal Williamson and Taina Honkalehto of the AFSC for providing data.

References

Billheimer, D. D. 1995. Statistical analysis of biological monitoring data: state-space models for species compositions. PhD thesis, University of Washington, Seattle.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. 1995. Bayesian data analysis. Chapman & Hall, London.

Hammond, T. R., and Swartzman, G. L. 2001. A general procedure for estimating the composition of fish school clusters using standard acoustic survey data. ICES Journal of Marine Science doi: 10.1006/jmsc.2001.1103.

James, M. 1985. Classification algorithms. John Wiley, New York.

Knudsen, H. P. 1990. The Bergen echo integrator: an introduction. Journal du Conseil International pour l'Exploration de la Mer, 47: 167–174.

Punt, A. E., and Hilborn, R. 1997. Fisheries stock assessment and decision analysis: the Bayesian approach. Reviews in Fish Biology and Fisheries, 7: 35–64.

Swartzman, G. L. 1997. Analysis of the summer distribution of fish schools in the Pacific Eastern Boundary Current. ICES Journal of Marine Science, 54: 105–116.

Swartzman, G. L., Stuetzle, W., Kulman, K., and Wen, N. 1994. Modeling the distribution of fish schools in the Bering Sea: morphological school identification. Natural Resource Modeling, 8: 177–194.

Appendix 1. Determining proposal distributions for MCMC

The performance of the BASCET algorithm is affected by a number of user-specified parameters (Hammond and Swartzman, 2001), some of which determine the distributions used to propose moves to particular estimation model parameters. In effect, they tune the MCMC estimation method that underlies the BASCET approach. A simple rule of thumb is suggested for determining these proposal distributions.

In MCMC methods, proposals are made to model parameters from a transition probability function (q), and these proposals are either accepted or rejected according to a particular rule. Experience with MCMC has suggested that it is desirable to tune proposal distributions so as to achieve acceptance rates between 30% and 80% (Billheimer, 1995). Low acceptance rates suggest that the proposal distributions are too broad, whereas high ones suggest that they are too narrow. Neither extreme case will allow for the efficient traversal of parameter space that is required to achieve timely convergence of the algorithm.

This simple rule was applied by running test chains and examining the acceptance rates for various proposal distributions. Proposal distributions were adjusted until the desired acceptance rates were achieved.

Appendix 2. How many iterations of MCMC?

It is of interest to determine not only how many iterations are required by MCMC methods to achieve satisfactory results, but also how many iterations are required to "forget" the arbitrary initial value for the chain. The latter number is often referred to as the burn-in period. These questions were addressed using a technique suggested by Gelman *et al.* (1995, p. 330). The idea of the technique is to run several Markov Chains of a given length, each starting from a different initial value. The start values are supposed to be overdispersed relative to the posterior distribution, a requirement that was met by drawing them randomly from the joint prior distribution. If the chains have both converged, then box plots constructed from parameter values in each chain should be similar. Gelman *et al.* (1995) suggest a test statistic called $\sqrt{\hat{R}}$ (computed using the between- and within-chain variance), which should be close to 1 if convergence has occurred. According to the authors, values below 1.2 are acceptable for most applications.

Two test chains of length 300 000 (after a burn-in of 100 000), each sampled every 10 iterations, were used in the computation of $\sqrt{\hat{R}}$ statistics. These statistics were computed for all of the classification parameters depicted in Figures 5–9, as well as for a random selection of cluster composition estimates (clusters 24, 55, 78, and 159). All values were below 1.1 and thus our test showed no sign of failed convergence. Both chains were combined for final inference.

Appendix 3. Glossary of BASCET jargon

α_d – A parameter vector that affects the depth distribution of schools. The kth element is denoted by $\alpha_{d,k}$; it determines the depth distribution of the kth species class

β_d – A parameter vector that affects the depth distribution of schools. The kth element is denoted by $\beta_{d,k}$; it determines the depth distribution of the kth species class

σ_a^2 – The variance of acoustic feature a

σ_s^2 – Variance in hotspot log spatial range

σ_t^2 – Variance in hotspot log temporal range

μ_a – A vector of length K indicating the mean value of acoustic feature a for a pure school of each species class. The kth element is denoted by $\mu_{a,k}$; it determines the mean of acoustic feature a for the kth species class

ξ_b – A composition vector that determines the effect of environmental attribute b on school cluster composition. Symbols ξ_{temp} and ξ_{depth} stand for the effects of temperature and ocean depth respectively

Σ_H – Determines the variability in hotspot composition

Ω_H – A parameter of the prior distribution for Σ_H

$\boldsymbol{\omega}_{temp}$ – A parameter of the prior distribution for ξ_{temp}

$\boldsymbol{\Psi}_{temp}$ – A parameter of the prior distribution for ξ_{temp}

$\boldsymbol{\omega}_{depth}$ – A parameter of the prior distribution for ξ_{depth}.

$\boldsymbol{\Psi}_{depth}$ – A parameter of the prior distribution for ξ_{depth}.

λ_H – Hotspot intensity. Determines the expected number of hotspots in the survey region

μ_s – Expected hotspot log spatial range

λ_S – School intensity. Determines the expected number of schools in a hotspot

μ_t – Expected hotspot log temporal range

A – The number of acoustic attributes of a school cluster

Backscatter image – An image in which each pixel represents the S_v value in a rectangular curtain of water beneath the path of the acoustic survey vessel.

$\mathbf{c}_j$ – The composition of school cluster j. A vector of length K whose entries sum to 1

Class – A category of fish defined either by size or by species. A class may also consist of several species. The number of classes is K

Cluster – A school cluster

Composition – A vector whose individual entries sum up to 1

$\mathbf{c}_{T,l}$ – The composition of trawl sample l

d_h – A horizontal threshold used in the definition of school clusters

d_v – A vertical threshold used in the definition of school clusters

Haul – A trawl sample

Hotspot – A region of space and time within which the composition of size or species classes can be different from the overall class composition in the acoustic survey region

Hotspotgen – An algorithm used to generate hotspots using the observed locations of school clusters and trawl samples. Takes four arguments: $R^T_{cluster}$, $R^S_{cluster}$, R^S_{trawl}, and R^S_{trawl}

K – The number of species classes

m – A numbers to weight conversion vector of length K

MCMC – Markov Chain Monte Carlo, a method for generating a sample from a Bayesian posterior distribution

$\mathbf{p}_H$ – Expected hotspot composition

$\mathbf{p}_i$ – The composition of hotspot i

Posterior – A distribution that represents uncertainty about the value of a parameter after observing the data

Prior – A distribution that represents uncertainty about the value of a parameter prior to observing the data. A Bayesian concept

$R^{S}_{cluster}$ – A parameter used in creating hotspot configurations using the hotspotgen algorithm

R^{S}_{trawl} – A parameter used in creating hotspot configurations using the hotspotgen algorithm

$R^{T}_{cluster}$ – A parameter used in creating hotspot configurations using the hotspotgen algorithm

R^{T}_{trawl} – A parameter used in creating hotspot configurations using the hotspotgen algorithm

School cluster – A collection of schools that are near each other

School – A data record consisting of a location, a depth, a time, and a collection of acoustic and environmental attributes that have been extracted from an acoustic backscatter image

$\mathbf{s}_g$ – A composition vector representing the relative vulnerability of classes to gear g

S_v – Volume reverberation

B – The number of environmental attributes of a school cluster