

An Introduction to Statistical Algorithms Useful in Stock Composition Analysis

MICHAEL H. PRAGER AND KYLE W. SHERTZER

Population Dynamics Team, Center for Coastal Fisheries and Habitat Research, National Oceanic and Atmospheric Administration, Beaufort, North Carolina, USA

-
- I. The Problem and Its Terminology
 - II. Algorithms
 - A. Discriminant Analysis
 - B. Logistic Regression
 - C. Artificial Neural Networks
 - D. Finite Mixture Distribution (FMD) Methods
 - III. The Importance of Prior Knowledge
 - A. Priors and Discriminant Analysis
 - B. Priors and Logistic Regression
 - C. Priors and Neural Networks
 - D. Priors and FMD Methods
 - IV. Discussion
 - References

I. THE PROBLEM AND ITS TERMINOLOGY

In many fisheries, catches include fish that are conspecific but that originate in several spawning stocks. Because the population effects of fishing—and thus the choice of suitable management approaches—depend on which stock or stocks are harvested, estimates of stock composition of catches are needed. This need has given rise to the set of techniques often labeled *stock identification*. The focus of applications is usually on proportions in the harvest rather than on origin of individual fish; consequently, a more precise description of this work is *stock composition analysis*.

We define stock composition analysis as estimation of the stock composition of a mixed-stock sample (usually, some part of the harvest) taken from a known number J of component stocks. The proportion that originates in any given stock j is represented P_j . Thus $\sum_{j=1}^J P_j = 1$; within this constraint, any particular P_j may equal zero. The process by which the P_j are estimated constitutes the stock composition analysis.

Data used for such analyses are observations on characteristics of individual specimens; typical characteristics may include morphometrics, meristics, genetic characters, or chemical signatures (reviewed in Begg and Waldman, 1999). When we refer to characteristics here, we assume that they have been quantified in some reasonable way so that statistics (such as mean and variance) can be computed for the entire set (matrix) of observed characteristics. The probabilistic nature of the methods considered here is needed to account for overlap in the distribution of characteristics from different populations. When there is no overlap (as when using tags), the origin of each fish can be established with certainty, and much simpler methods can be used to define the stock composition of the catch.

This discussion also assumes the availability of a training sample of individuals whose stock membership is known. The training sample is used to fit a model by means of the investigator's choice of algorithm, a word used here to denote a statistical method or group of related methods. The fitted model is then used to estimate the stock composition of a mixed-stock sample (or samples) of the catch. The estimation can, but need not always, involve estimating the probability of stock membership of each individual in the mixed-stock sample. If in the course of estimation each individual is assigned membership in a particular stock, the method can be considered a classification method. Classification methods are a subset of all methods useful for stock composition analysis, because stock composition can be estimated without performing a classification.

The terminology of stock composition analysis is specific to fishery science, but the general problem is not. Analyzing characteristics of individual objects in a mixture to estimate the mixture's proportions is a general statistical problem known as finite mixture analysis. Constituent fish stocks in a mixed harvest are just one example of constituent classes or statistical populations of mixed objects; here, we tend to use the term *class* when describing algorithms generally, and *stock* when describing fisheries applications. Estimating stock composition is therefore a special case of the general problem of estimating mixing proportions.

We continue this chapter by introducing some algorithms useful for stock composition analysis. We then discuss issues involved in estimating performance of various algorithms, either in an absolute sense or relative to one another on a particular data set. We close with a few general comparative remarks.

An li

II.

Mar
anal
neu
rate
sinc
vari

A.

Am
and
con
ana
met
bee
ava
impmu
am
trib
vec
and
usefun
isti
j =
thewh
tity
fur
cla
δ,
lar,

II. ALGORITHMS

Many statistical algorithms have been proposed or used for stock composition analysis. We consider four—discriminant analysis, logistic regression, artificial neural networks, and finite mixture distributions—each of which may incorporate more than one observed characteristic. Variants of each algorithm exist, but since we focus on the conceptual basis of each algorithm, our treatment of such variants is generally brief.

A. DISCRIMINANT ANALYSIS (DA)

Among classification schemes, discriminant analysis (McLachlan, 1992; Johnson and Wichern, 1998; Hastie et al., 2001) boasts the longest history. Its two most common forms are linear discriminant analysis (LDA) and quadratic discriminant analysis (QDA). Linear discriminant analysis was the first formal statistical method used for stock composition analysis (Hill, 1959), and the method has been used many times and in many variations. One of its advantages is the wide availability of well-tested and flexible software, as discriminant analysis forms an important component of most major statistics packages.

The fundamental assumption of LDA is that observed characteristics follow a multivariate normal distribution with common variance–covariance structure among stocks. If the characteristics vector is represented \mathbf{X} , we can write this distribution for stock j as $f_j(\mathbf{X}) \sim \text{MVN}(\boldsymbol{\mu}_j, \boldsymbol{\Sigma})$ where $\boldsymbol{\mu}_j$ is the mean characteristics vector for stock j and $\boldsymbol{\Sigma}$ is the common variance–covariance matrix. Typically, $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}_j$, $j = \{1, \dots, J\}$ are estimated from the training set, and those estimates are used in forming discriminant functions.

Classification in LDA is determined from stock-specific linear discriminant functions, computed from three types of information: an individual's characteristics vector \mathbf{x} ; estimates of $\boldsymbol{\mu}_j$ and $\boldsymbol{\Sigma}$; and a set of prior probabilities, or *priors*, p_j , $j = \{1, \dots, J\}$. The priors are the analyst's estimates (which may be subjective) of the probabilities that a randomly chosen fish originates in each of the j stocks.

The discriminant function for stock j , evaluated for individual \mathbf{x} , is

$$\delta_j(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_j - \frac{1}{2} \boldsymbol{\mu}_j^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_j + \log(p_j) \quad (1)$$

where the notation \mathbf{M}^t indicates the transpose of vector or matrix \mathbf{M} . The quantity $-\delta_j(\mathbf{x})$ measures the distance from individual \mathbf{x} to the center of stock j (the function is conventionally in negative distance for computational reasons). In classification, each individual is assigned membership to the stock that maximizes $\delta_j(\mathbf{x})$, which is the stock with the closest center. This is also the stock with the largest posterior probability (eq. 3 of Pella and Masuda, this volume, Chapter 25).

The J discriminant functions thus define decision boundaries that classify an individual, from its observed characteristics, to the most likely stock of origin. The decision boundary between stocks j and k occurs where $\delta_j(x) = \delta_k(x)$, and it is this decision boundary that, when solved for x , is linear in the observed characteristics. If there are two measured characteristics, the decision boundary is a line; if three, a plane, and so on.

Quadratic discriminant analysis (QDA) is often considered preferable for problems in which the variance-covariance structure differs by class (Misra, 1985), as QDA does not assume equality of variance among classes (Kendall et al., 1983). However, estimates from QDA generally are of higher variance than those from LDA because of the additional parameters that must be estimated.

The quadratic discriminant function for stock j , evaluated for individual x , is

$$\delta_j(x) = -\frac{1}{2} \log |\Sigma_j| - \frac{1}{2} (x - \mu_j)^T \Sigma_j^{-1} (x - \mu_j) + \log(p_j) \quad (2)$$

where Σ_j is the stock-specific variance-covariance matrix and $|\Sigma_j|$ is its determinant. Classifications and posterior probabilities of class membership are computed as in LDA.

The decision boundary in QDA is defined as in LDA, but the resulting boundary is quadratic (curved) in the observed characteristics. For that reason, decision boundaries for LDA and QDA generally differ (Fig. 24-1).

In using either linear or quadratic discriminant analysis to estimate stock composition, one can proceed in two slightly different ways. The usual procedure, which we term *discrete classification*, is to classify each individual in the (mixed-stock) sample into the class for which its estimated membership probability is highest. The estimate of stock composition is then formed from the relative numbers of individuals classified into each class. In the second procedure, which we term *nondiscrete classification*, the probability of membership in a particular class is summed across all individuals. The estimate of stock composition is then obtained from the sums for each class divided by the total sample size.

As a simple example, consider a two-stock problem in which three fish are in the mixed-stock sample. Let the estimated probabilities of membership in class I for the three fish be {0.55, 0.45, 0.8}. The discrete estimate of mixing proportions would be 2/3 from class I and 1/3 from class II. The nondiscrete estimate would be 0.6 from class I and 0.4 from class II. The discrete estimate is obtained because two of the three fish are thought more likely to belong to class I; the nondiscrete is the mean of the three probabilities given.

Although discrete and nondiscrete classification usually produce similar estimates, it seems logical to prefer nondiscrete classification. There is no necessity to round estimated membership probabilities to whole numbers, as done in discrete classification, when the objective is to estimate mixing proportions.

4 -
2 -
0 -
x²

-2

-4

FIGURE 24
circles) from
observed ch

and $\mu_i =$

$$\Sigma_i = \begin{pmatrix} 1.8 \\ -0.2 \end{pmatrix}$$

the mean 0

Severe
tification
age-inval
analysis
Andrade
results o
discrimi




FIGURE 24-1. Example of decision boundaries between stock j (filled circles) and stock i (open circles) from linear (dotted line) and quadratic (solid line) discriminant analysis based on two observed characteristics (X_1, X_2). The prior probabilities are $p_j = p_i = 0.5$; the means are $\mu_j = \begin{pmatrix} 0.6 \\ 0.75 \end{pmatrix}$ and $\mu_i = \begin{pmatrix} -0.8 \\ -0.5 \end{pmatrix}$; the variance-covariance matrices are $\Sigma_j = \begin{pmatrix} 1.5 & -0.25 \\ -0.25 & 1.25 \end{pmatrix}$ and $\Sigma_i = \begin{pmatrix} 1.8 & -0.2 \\ -0.2 & 2.0 \end{pmatrix}$; and, with equal sample sizes, the pooled matrix Σ for the linear boundary is the mean of Σ_j and Σ_i .

Several other variants of discriminant analysis have been applied to stock identification; these include polynomial discriminant analysis (Cook and Lord, 1978), age-invariant discriminant analysis (Fabrizio, 1987), jackknife discriminant analysis (Small et al., 1998), and stepwise discriminant analysis (Palma and Andrade, 2002). Correction matrices, which can be computed from classification results on a test data set, are frequently used to improve mixture estimates from discriminant analyses (Cook and Lord, 1978; Pella and Robertson, 1978) and

ager and Kyle W. Shertzer

boundaries that classify an likely stock of origin. re $\delta_i(x) = \delta_k(x)$, and it in the observed char- decision boundary is a

ed preferable for prob- y class (Misra, 1985), (Kendall et al., 1983). ance than those from imated.

ed for individual x , is $+ \log(p_j)$ (2)

and $|\Sigma_i|$ is its deter- ass membership are

the resulting bound- or that reason, deci- -1).

sis to estimate stock it ways. The usual fy each individual in imated membership hen formed from the n the second proce- / of membership in a ate of stock compo- lided by the total

hich three fish are in membership in class te of mixing propor- nondiscrete estimate estimate is obtained along to class I; the

roduce similar esti- here is no necessity ers, as done in dis- proportions.

might be used to correct estimates from other classification-based methods as well. Millar (1987) demonstrated that the use of classification with correction is a special case of maximum-likelihood finite mixture distribution methods (described below).

B. LOGISTIC REGRESSION

Logistic regression (Aldrich and Nelson, 1984; Hosmer and Lemeshow, 1989; Agresti, 2002) is a type of generalized linear model (McCullagh and Nelder, 1989). It was suggested for stock identification by Prager and Fabrizio (1990), who found the method promising. Its chief theoretical advantage is that it assumes neither multivariate normality of input data nor equality of variances and is appropriate for a wide variety of distributions (Kendall et al., 1983). It can also handle input data that are continuous, categorical, or a mix rather than continuous only, as in discriminant analysis. Logistic regression is applied most often to problems with a binary response, as when analyzing mixtures of two source stocks. But its use is not limited to binary problems, and indeed logistic regression has been applied to stock identification problems with more than two stocks (Prager and Fabrizio, 1990; Waldman et al., 1997).

In binary logistic regression, the response (Y_i) for fish i takes one of two values: $Y_i = 0$ implies membership in the first stock, and $Y_i = 1$ implies membership in the second. The probability that fish i with N measured characteristics $\mathbf{x}_i = (x_{i1}, x_{i2}, \dots, x_{iN})$ belongs to the second stock is estimated by the continuous logistic function π ,

$$P(Y_i = 1 | \mathbf{x}_i) = \pi(\mathbf{x}_i) = \frac{\exp(z_i)}{1 + \exp(z_i)} \quad (3)$$

where

$$z_i = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_N x_{iN} \quad (4)$$

and the β 's are parameters to be estimated (Fig. 24-2). The analysis defined by eq. 4 is called multiple logistic regression, which refers not to the number of stocks, but to analyzing more than one characteristic per fish (i.e., $N > 1$). A suitable transformation, accomplished by use of a link function, converts the model of eq. 3 into one that is linear. Several standard links exist for binary data, such as logit, probit, or log-log (Agresti, 2002); we present here the logit link because it is most often applied in polytomous logistic regression (more than two classes). The logit link is

$$\log\left(\frac{\pi(\mathbf{x}_i)}{1 - \pi(\mathbf{x}_i)}\right) = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_N x_{iN} \quad (5)$$

FIGURE 24
tion, z is an
formation π

After tran
meters ca
cal softwa

If there
to polyto
The majo
mially rat
functions
 $j = 1, \dots$

where z_j
is constr

based methods as with correction is ribution methods

Lemeshow, 1989; Blag and Nelder, 1989; Blag and Fabrizio (1990), advantage is that it handles the equality of variances (McCullough et al., 1983). It can be used rather than compared most often for regression of two source classes. The estimated logistic regression is better than two stocks

one of two values: membership in the first stock. Characteristics $x_i = (x_{i1}, \dots, x_{in})$ are continuous logistic




FIGURE 24-2. Logistic transformation shown in eq. 3. In logistic regression for stock identification, z is an unbounded linear combination of the measured characteristics. The value of the transformation π is the estimated probability of membership in the second of two stocks.

(3)

After transformation, errors are assumed to be distributed binomially, and parameters can readily be estimated by maximum likelihood using standard statistical software.

(4)

If there are more than two classes, binary logistic regression can be extended to polytomous logistic regression (Hosmer and Lemeshow, 1989; Agresti, 2002). The major difference is that now errors are assumed to be distributed multinomially rather than binomially. Classification among J stocks requires $J - 1$ link functions, which is no different from the binary case. For stock membership $j = 1, \dots, J$,

analysis defined by the number of stocks ($N > 1$). A suitable link function converts the model to a multinomial logit link because the data are binary (two classes).

$$P(Y_i = j | x_i) = \pi_j(x_i) = \frac{\exp(z_{ij})}{\sum_{h=1}^J \exp(z_{ih})} \quad (6)$$

(5)

where z_{ij} are analogous to eq. 4 but with parameters $\beta_{0j}, \beta_{1j}, \dots, \beta_{Nj}$. The problem is constrained by the requirement that the probabilities of stock membership sum to 1.

to one. Given $J - 1$ estimates of π_i , the J th estimate must be $\pi_J(x_i) = 1 - \sum_{j=1}^{J-1} \pi_j(x_i)$. That constraint is implemented by defining parameters for the J th stock to be zero (i.e., $\beta_{0J}, \beta_{1J}, \dots, \beta_{NJ} = 0$). As a result, eq. 6 simplifies because $\exp(z_{ij}) = 1$, and the π_i 's can be estimated under transformation by a link function. As in the binary case, there are $J - 1$ unique logit equations,

$$\log\left(\frac{\pi_j(x_i)}{\pi_j(x_i)}\right) = \beta_{0j} + \beta_{1j}x_{i1} + \beta_{2j}x_{i2} + \dots + \beta_{Nj}x_{iN} \quad (7)$$

Computer programs for logistic regression are readily available in the major statistics software packages. In using such software, the analyst usually wishes to specify that the stock designations are nominal rather than ordinal values. Careful reading of the software's documentation may be needed to effect that choice, which is not always the default.

C. ARTIFICIAL NEURAL NETWORKS

The term *artificial neural network* (ANN) is not a precise one, but refers to a group of computational algorithms that sift through and combine many models to arrive at a model of optimum (in some sense) complexity (Ripley, 1996; Hastie et al., 2001). Unlike the other classification schemes presented here, ANNs are non-parametric. They require no assumptions about the distribution of data nor the particular functional relationship between model input and output. This can be a major advantage when such assumptions would be violated. Nonetheless, the success of such methods still depends on similarity of the data in the mixed-stock sample to the data in the training set.

Artificial neural networks have been developed in analogy to the structure of the human brain. Like the brain, an ANN consists of interconnected layers that process information provided by neurons. The input layer performs computations on the input data, and the results, along with a constant (bias), are then passed to a hidden layer. That procedure is iterated among a series of hidden layers until finally results are passed to the output layer. The number of hidden layers and the number of neurons in each can be adjusted to reflect the complexity of the problem. Once the architecture is established, values of network parameters are chosen as those that minimize some fitting criterion, a task usually accomplished with a learning algorithm.

Neural networks have proved useful in numerous fields, such as artificial intelligence, image compression, medical diagnosis, nondestructive testing, signal processing, and terrain classification, to name only a few. To our knowledge, the first published application in stock composition analysis was by Prager (1984, 1988) to estimate stock composition of striped bass and American shad. That study used

X₁X₂X₃

FIGURE 24-3.
ables (X_j) are th
processed by the
The final result is

the group me
of neural netw
performances
cessful, at lea
Thorrold et a

In the cont
the measured
illustrates a n
two compone
layer sums th
The result is
signal for the
an output-lay

$\pi_j(x_i) = 1 - \sum_{i=1}^j \pi_i(x_i)$.
 or the j th stock to be
 because $\exp(z_{ij}) = 1$,
 link function. As in the

$$\beta_{N_j} x_{iN} \quad (7)$$

available in the major
 analyst usually wishes to
 ordinal values. Careful
 to effect that choice,

, but refers to a group
 many models to arrive
 y, 1996; Hastie et al.,
 here, ANNs are non-
 of data nor the
 l output. This can be
 ted. Nonetheless, the
 ita in the mixed-stock

gy to the structure of
 connected layers that
 r performs computa-
 tant (bias), are then
 g a series of hidden
 ne number of hidden
 d to reflect the com-
 d, values of network
 iteration, a task usually

such as artificial intel-
 ve testing, signal pro-
 r knowledge, the first
 Prager (1984, 1988)
 had. That study used




FIGURE 24-3. Example of an artificial neural network with a single hidden layer. The input variables (X_i) are the measure characteristics. Their weighted (w_{ab}) values plus a bias term (β_b) are processed by the hidden layer, and in turn, those results are processed similarly by the output layer. The final result is probabilities of stock membership, among two stocks in this example.

the group method of data handling (Ivakhnenko and Ivakhnenko, 1974), a type of neural network based on polynomials. More recent studies have compared the performances of ANNs and LDA and have found ANNs to be slightly more successful, at least on the particular data sets analyzed (Taylor and Beacham, 1994; Thorrold et al., 1998; Wells et al., 2000).

In the context of stock composition analysis, the input variables for an ANN are the measured characteristics; the output is the stock classification. Figure 24-3 illustrates a neural network with a single hidden layer that classifies a sample into two component stocks based on three characteristics. Each neuron in the hidden layer sums the weighted (w_{ab}) input signals (X_1, X_2, X_3) and adds a bias term (β_b). The result is then processed by a hidden-layer function (f_H) to produce an input signal for the output layer. The procedure is the same at the output layer, but with an output-layer function (f_O). Figure 24-3 can be expressed mathematically as

$$Y_k = f_O \left(\beta_k + \sum_b w_{bk} f_H \left(\beta_b + \sum_a w_{ab} X_a \right) \right) \quad (8)$$

