Methods in Oceanography 17 (2016) 187–205

Contents lists available at ScienceDirect

Methods in Oceanography

journal homepage: www.elsevier.com/locate/mio

Full length article

Acoustic identification of marine species using a feature library

Rolf J. Korneliussen [a,b,*], Yngve Heggelund [b], Gavin J. Macaulay [a], Daniel Patel [b], Espen Johnsen [a], Inge K. Eliassen [b]

[a] *Institute of Marine Research, P.O. Box 1870, 5817 Bergen, Norway*

[b] *Christian Michelsen Research AS, P.O. Box 6031, NO-5892 Bergen, Norway*

GRAPHICAL ABSTRACT

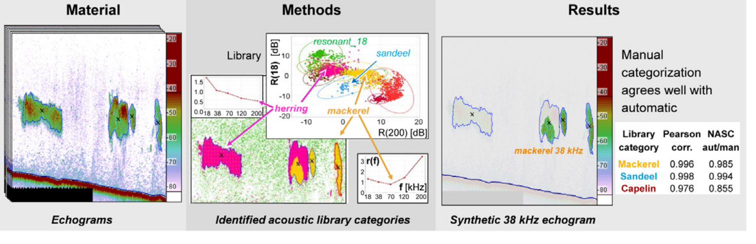

HIGHLIGHTS

- Echo data are categorized using an acoustic-feature library of ground-truthed target spectra.
- Automatic data-processing modules allow faster scrutiny, better quality and objectivity of results.
- Comparison of manual and automatic classification for three oceanic surveys in different ecosystems.
- High correlation between manual and automatic classification of backscatter for all case studies.

* Correspondence to: Institute of Marine Research, P.O. Box 1870 Nordnes, 5817 Bergen, Norway. Fax: +47 55 23 85 31.
E-mail address: rolf.korneliussen@imr.no (R.J. Korneliussen).

http://dx.doi.org/10.1016/j.mio.2016.09.002

ARTICLE INFO

Article history:
Received 11 December 2015
Received in revised form
12 August 2016
Accepted 6 September 2016

Keywords:
Relative frequency response
School-detection
Acoustic-based categorization
Acoustic feature library
Large Scale Survey System
LSSS

ABSTRACT

Sonars and echosounders are widely used for remote sensing of life in the marine environment. There is an ongoing need to make the acoustic identification of marine species more correct and objective and thereby reduce the uncertainty of acoustic abundance estimates. In our work, data from multi-frequency echosounders working simultaneously with nearly identical and overlapping acoustic beams are processed stepwise in a modular sequence to improve data, detect schools and categorize acoustic targets by means of the Large Scale Survey System software (LSSS). Categorization is based on the use of an acoustic feature library whose main components are the relative frequency responses. The results of the categorization are translated into acoustic abundance of species. The method is tested on acoustic data from the Barents Sea, the Norwegian Sea and the North Sea, where the target species were capelin (*Mallotus villosus* L.), Atlantic mackerel (*Scomber scombrus* L.) and sandeel (*Ammodytes marinus* L.), respectively. Manual categorization showed a high conformity with automatic categorization for all surveys, especially for schools.

1. Introduction

The sustainable management of commercially exploited marine species requires reliable knowledge of their abundance as well as of that of their prey and predator species. This is often informed by spatially extensive surveys using a variety of sampling tools that catch the targeted species, or remotely detect their presence (King, 2007). A commonly used remote sensing method is acoustics, where echosounders are used to generate directed pulses of sound and the resulting backscatter from marine organisms is measured and mapped (Simmonds and MacLennan, 2005). Given knowledge of the backscatter and other characteristics for species of interest, it is possible to assign species or species groups to the backscatter. Then, when combined with a relationship between backscatter amplitude and species size one can derive an estimate of the species biomass in the surveyed area (Simmonds and MacLennan, 2005). Estimating and understanding the sources of uncertainty in biomass estimates is an important means to reduce uncertainty and improve the inputs to the fisheries management process.

A key procedure in acoustic biomass estimation is the correct assigning of backscatter to species or species-group level (MacLennan and Holliday, 1996; Horne, 2000), a process typically called scrutinizing or categorization. Categorization, supported by biological sampling and human experience is time-consuming and subjective and there is a need to make the process more objective so as to reduce human-generated bias and uncertainty in acoustic abundance estimates. It has been estimated that the systematic error associated with the categorization of acoustic backscatter can be as high as ±80% (Simmonds and MacLennan, 2005). Automated species identification can be a valuable input to the categorization process and offers the potential for a more objective result.

The potential for objective classification of targets from backscatter was recognized in the 1970s (Holliday, 1977; Deuser et al., 1979). The acoustic scattering characteristics of some species may not be reliably distinguished from others by acoustic means only due to similarities in morphology and behaviour. Tilt distribution, i.e. swimming upwards or downwards, usually results in decreasing backscatter with increasing frequency. However, groups of similar species may be distinguished from other species or scattering groups (e.g. fish with swimbladder vs. zooplankton or fish without swimbladder).

The categorization methods can be divided into those that use multiple frequencies and those that primarily use a single frequency to identify school species. In some specific cases, school characteristics from a single frequency may provide sufficient information for species identification, but these techniques have not been universally applied because temporal changes in school shape and location (on scales from hours to years) can affect the results. For example, Scalabrin et al. (1996) concluded that their use of school structure and species-related spatial distribution did not give a high probability of species identification.

The attempts to utilize multiple echosounder frequencies to derive automated categorizing routines have been more successful. In the 1990s artificial neural feedback networks (ANN) were used to categorize acoustic data and are improved and expanded by the use of ground-truthed data for training the network (Haralabous and Georgakaros, 1996). A major problem with ANNs is that weights involved in the neural network are "hidden" (i.e. cannot be interpreted by humans), which makes systematic improvement to species identification difficult. A problem for many methods (Kloser et al., 2002; Korneliussen and Ona, 2002, 2003; Fernandes et al., 2002; Korneliussen, 2010) is the use of properties that are pre-defined and difficult to modify. The average relative frequency responses, $r(f)$, (Korneliussen and Ona, 2002), were measured for some species (Korneliussen and Ona, 2002, 2003; Fernandes et al., 2002; Korneliussen, 2010), and acoustic categories (representing the species) were then hardcoded into the BEI (Bergen Echo Integrator) postprocessing system (Korneliussen, 2004) to identify the acoustic categories. Pedersen and Korneliussen (2009) corrected for the beam-patterns of the transducer by using split-beam data of single-fish tracks so that species could be reliably distinguished from fewer samples than were needed from school measurements. De Robertis et al. (2010) used the Z-score as a summary to classify backscatter measurements based on $r(f)$, and used it in real-time with reasonable success. Although not explicitly stated, the Z-score implementation is likely to contain predefined features (given in their Table 2) in a similar manner to Fernandes et al. (2002) and Korneliussen (2010).

A more flexible solution is to use a range of features, such as $r(f)$, but let the statistical properties be derived from an expandable knowledge database. Many species occur in monospecific registrations that are the preferred type of acoustic registrations used to populate an acoustic library. A library could also contain non-acoustic features, such as geographical areas, depth and temperature ranges. Korneliussen et al. (2009a) organized $r(f)$ features in a library, but their focus was to show how aggregating and averaging acoustic measurements inside schools could improve the ability to distinguish species with similar acoustic properties. That work did not describe the details of the categorization method, so their results would be hard to reproduce and verify. Furthermore, they applied the method to a limited dataset where the school shapes were used as supporting information to distinguish between capelin and herring only, making the reliability of the method difficult to assess.

For some species it is difficult to find monospecific registrations and the categorization can be based on scattering models. The use of scattering models for zooplankton (Stanton et al., 1994; Demer and Conti, 2003) in a framework for estimating zooplankton specimen size (Holliday, 1977; Holliday et al., 1989; Fernandes et al., 2002) overcomes the problem of finding monospecific aggregations of zooplankton in-situ. Some large schools of euphausiids can be identified both from scattering models as well as from an acoustic library (for example, *Euphausia superba* L. in the southern ocean (Korneliussen et al., 2009b), or *Meganyctiphanes norwegica* L. in the North Atlantic Ocean), but such schools of zooplankton are the exception rather than the rule (Kaartvedt et al., 2005).

In this paper we describe categorization methods, methods to improve data quality, techniques for efficient processing of multi-frequency acoustic data, and automatic identification of biologically generated backscatter. These are all steps that make categorization of acoustic data practicable on operational surveys and provide a more objective scrutinizing process. The methods allow for processing large volumes of high resolution data from echosounders, which in turn makes it possible for a limited number of experts to analyse such data within restricted time periods. Special emphasis is given to methods for automatic categorization of acoustic features by means of statistical properties with reference to an acoustic feature library. We describe the realization of these methods for the efficient scrutinization of acoustic surveys. The reliability of the methods is illustrated with data from three surveys targeting Atlantic mackerel (*Scomber scombrus* L.), sandeel (*Ammodytes marinus* L.) and

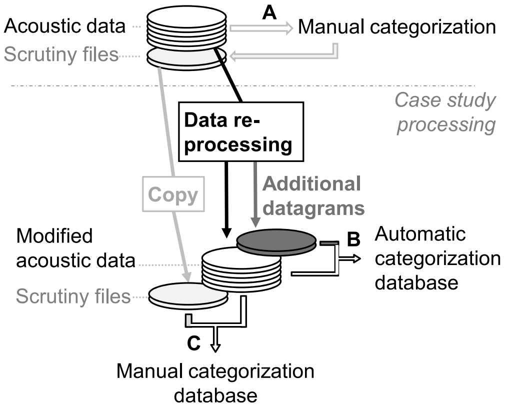

Fig. 1. Workflow used in the analysis. (A) Data were manually categorized and the results stored in scrutiny files. (B) Data were reprocessed, e.g. to remove noise. The scrutiny files were used with reprocessed acoustic data to generate the *manual categorization database*. (C) The acoustic data were further processed to add datagrams (grey in figure), e.g. automatically identified acoustic library categories and school-detections. The school detections were used only to improve acoustic library categorization in schools. The acoustic library categories were mapped to manual scrutiny categories so they could be stored in a database, in this case the *automatic categorization database*.

capelin (*Mallotus villosus* L.) in three different ecosystems (Norwegian Sea, North Sea, and Barents Sea, respectively).

2. Concepts and background

2.1. Overview of analysis system

Data were scrutinized on a daily basis by a team of experts within 24 h of data collection (Fig. 1(A)). The result of the scrutiny was nautical area scattering coefficient values (s_A, m^2 nmi^{-2}, MacLennan et al., 2002) allocated to categories and stored in a relational database. s_A is a measure of the backscattering intensity per unit surface area over some depth interval. The results of echogram categorization (Fig. 2) are used to estimate fish stock abundance for each of the surveys and are stored in scrutiny files together with the acoustic data for additional inspection.

2.2. The LSSS software

The Large Scale Survey System (LSSS) software was developed for the analysis of acoustic backscatter from marine organisms such as zooplankton, pelagic and demersal fishes. LSSS is a software package programmed in Java, developed specifically for the needs of the Institute of Marine Research (IMR), Norway. It is designed for efficient post-processing of large amounts of echosounder and sonar data, and contains tools to isolate echo-traces by creating regions and assigning categories to these regions. A design aim of LSSS was to achieve the categorization of acoustic data collected over 24 h within 2 h of operator time. LSSS has a pre-processing facility (called KORONA) that carries out time-intensive automated processing of acoustic data, such as removing noise, aligning data due to transducers offsets, and other operations that make manual categorization actions faster (such as school detection, single fish track detection and bottom detection; Anon, 1990). Furthermore, through automatic categorization KORONA provides information that allows for more objective scrutiny of the echograms. Some of this information is not accessible without pre-processing, e.g. the identification of specified acoustic categories from completely mixed acoustic aggregations. The information from KORONA is available during categorization, but it is the scrutinizing team that decides if any, parts or all the information from KORONA is used. The primary GUI of LSSS focusses on the acoustic data, with overlays providing additional information to the operator (Fig. 2).

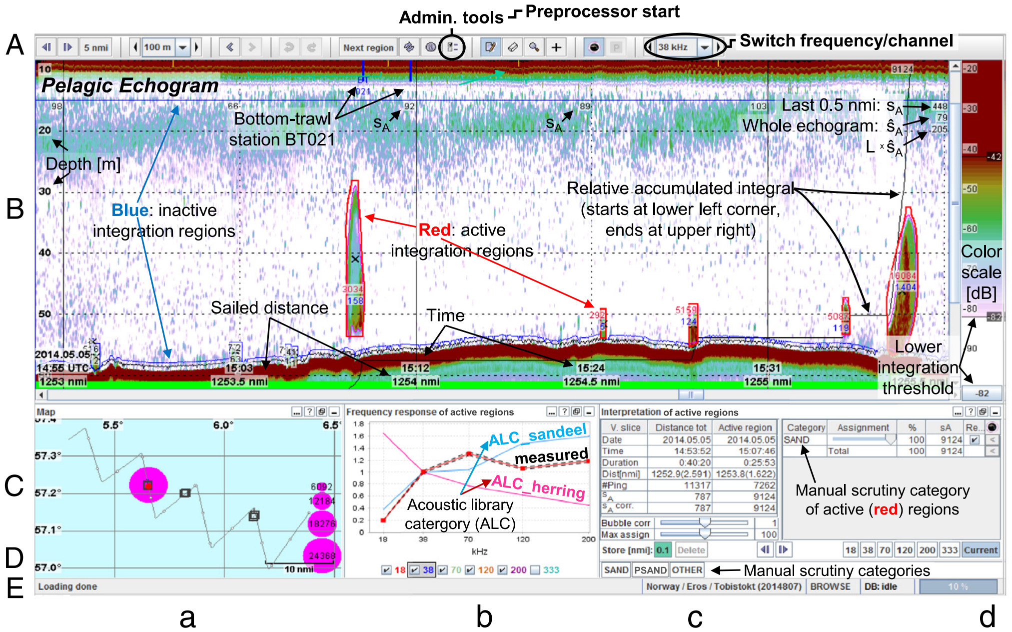

Fig. 2. Graphical user interface (GUI) of LSSS for interpreting backscatter data. The numbers (red for the active and black for inactive scatter-regions) on the echograms (B) show values of s_A. The s_A-values against each vertical black line refer to the whole water column within the block; the mean of these for the whole echogram is shown in the top right corner. Various scatter-regions outlined in red appear on the echogram, along with the corresponding s_A. The blue numbers, when present, show the mean s_A multiplied by the corresponding length of the scatter region. (A) Toolbar; (B) Pelagic echogram, i.e. showing backscatter with surface at the top. The areas indicated with red outlines are the regions being scrutinized. These are linked to the windows at position Cb and Cc. (Bd) Echogram colour scale showing volume backscattering strength, S_v. (Ca) Cruise track on a latitude/longitude grid; (Cb) The relative frequency response, $r(f)$. (Cc) Summary of interpretation results and general info; (Dc) Manual scrutiny categories (MSC). (E) Status bar that includes ship and survey name.

2.3. Echogram pre-processing

Repeated echoes from isolated single organisms, such as fish, are detected automatically (Handegard, 2007). Schools are detected as described below, where a school is defined by the boundary that encloses an aggregation of individually unresolved organisms with some minimum backscatter intensity.

In general, backscatter from two or more frequencies is combined to detect schools. Here we use the combined frequency channel $[s_v(38) + 3s_v(200)]/4$ for all case studies, where $s_v(38)$ and $s_v(200)$ are the volume backscattering coefficients at 38 kHz and 200 kHz. Although any combination of frequencies could be used, this combination is particularly well suited to the case studies used here: both 38 and 200 kHz are good for detecting swimbladdered fish such as herring and capelin, while 200 kHz is particularly useful for the detection of the two swimbladderless fish species used in the case studies (mackerel and sandeel). The use of combined frequency data, which is only used to detect the extent of the school, results in a smoother school boundary that is less dependent on frequency-dependent scattering. The combined frequency channel is smoothed with a 2D-Gaussian window (the normalized weights at ±12.5 m horizontally and ±0.5 m vertically are half the value of the centre weight). Median filters followed by erosion filters and dilation filters (Dougherty, 1992) are then applied to remove small scale variability while preserving the main features of the dataset.

Next, measured samples are grouped together in school regions and averaged. The school detection module of the pre-processor KORONA traces the borders of each region as defined by a mean volume backscattering strength (S_v, MacLennan et al., 2002) threshold of −62 dB re 1 m^{-1} for the examples presented here, and joins or splits borders at the correct location as soon as a new ping arrives, thereby providing near real-time detection of schools (i.e. within a few seconds). The traced borders are merged to form closed curves which constitute outlines of potential schools. The potential schools must match user-defined criteria to be accepted as a school. These are:

- Scattering strength: upper and lower limits on the average S_v inside the region;
- Overlap: the values inside the school border are compared with those just outside the region by fitting each to a Gaussian distribution. The percentage overlap of these distributions is calculated and used as a criterion. The thickness of this band is user-defined.
- Lower limits on the length, thickness and area of the region.

The density interval and overlap criteria make it possible to detect regions with subtle data variations, such as seen in plankton concentrations.

3. Categorization methods

The LSSS software provides three types of acoustic backscatter categories. These are:

1. Acoustic library categories (**ALC**), which are obtained from reference to a library of ground-truthed acoustic backscatter characteristics.
2. Model-based categories (**MBC**), which are obtained by comparison to backscatter frequency spectra obtained from acoustic backscattering models.
3. Manual scrutiny categories (MSC), where the operator manually attaches categories to acoustic backscatter.

Only the manual scrutiny categories can be used for echo integration and subsequent biomass estimation. The two other categories are either used to inform the manual categorization, or are mapped to manual scrutiny categories via an operator defined mapping matrix. Manual scrutiny categories correspond closely to species or species groups. Acoustic library categories and model-based categories do not necessarily correspond to species or species groups, but are more focussed on the distinguishable characteristics of the backscatter. The conversion of ALC and MBC to the manual scrutiny categories of relevance to particular surveys is an important function of the LSSS operator. There are currently 118 MSC, 25 ALC and five MBC available in LSSS. Acoustic categories should not be confused with species, which have a unique scientific name.

3.1. Manual interpretation of echograms (in general)

The backscatter data are visualized as echograms, and manually categorized (e.g. Jech and Michaels, 2006) by means of a computer program. IMR used the BEI software (Korneliussen, 2004) for surveys prior to 2007, and then the Large Scale Survey System (LSSS) software. A range of features are considered during *manual categorization*. First, unwanted areas are excluded, then unwanted backscatter is removed (e.g. bottom echoes and noise spikes). The operators then identify and select regions with similar acoustic properties. Some fish species typically appear as single individuals, while other species are usually seen in schools, layers or loose aggregations. The spatial distribution of the targets and the distribution of target strength aid identification of single-fish. The size, density and shape of a school also help to indicate the species. The average $r(f)$ of an encircled region may be used in differentiating between species. Active use of the S_v threshold (volume backscattering strength, Fig. 2(Bd)) is useful to differentiate some species of zooplankton and fish, but also to partition s_A between those species. Some species of fish and zooplankton are associated with sharp changes in temperature profile. The geographical area and proximity to sea surface and bottom is also used to identify species. However, the results of directed sampling of the biology are the most important supportive information used when manually categorizing echograms.

3.2. Acoustic feature library

The acoustic feature library contains ground-truthed data from selected acoustic backscatter that has been processed in a standardized manner. The processing involves noise removal and smoothing.

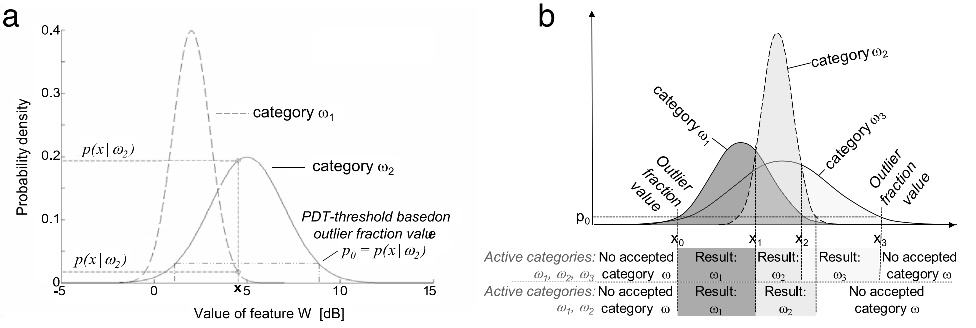

Fig. 3. (a) Probability density function of one acoustic feature W (i.e. one dimension) for two hypothetical ALC, ω_1 and ω_2, that is used to decide if measurement x (of acoustic feature W) is most likely to originate from category ω_1 or ω_2. Here $p(x|\omega_2) > p(x|\omega_1)$ indicates category ω_2 as the most likely. In practice many such comparisons between different features and categories contribute to the chosen category. (b) Example of results for value x of feature W when three acoustic-library categories ω_1, ω_2 and ω_3 are used. At values less than x_1, ω_1 is the most probable category (for probabilities above the threshold), but if $x < x_0$ the probability $p(x|\omega_1)$ is too small, being less than the threshold $p_0 = p(x_0|\omega_1)$, as determined by the specified outlier fraction (**cf. 3.2.2**). If ω_1 was not used, ω_3 would be accepted for $x < x_1$ (but with a different minimum value for x).

The statistical properties of each of the features of the acoustic library are calculated on several sample collection levels as described below. The library is organized by ALC, which for convenience are named after species since they are based on ground-truthed species. They are referred to as ALC due to the fact that the features of one species may also be used to identify species of similar acoustic properties (for example, a category derived from herring echoes may also identify the physically similar sardine). The average values of the acoustic measures $r(f)$ and s_v(38 kHz) are influenced by morphology and tilting behaviour.

3.2.1. *Populating the acoustic feature library*

The ALC are identified by their multi-frequency properties measured over some scattering volume. Currently, these are the volume backscattering coefficients, s_v, at 38 kHz and the ratio of s_v at the other frequencies to that at 38 kHz, that is $r(f) = s_v(f)/s_v(38)$. Thus, the variation of backscatter with frequency is covered by $r(f)$, and the scattering strength by $s_v(38)$. The logarithmic version $R(f) = 10\log_{10}(r(f))$ is used in most figures for convenience.

The acoustic feature library for each category contains a multi-dimensional normal probability density function (PDF) for the set of features, $W = W(R(18), S_v(38), R(70), \ldots, R(f), \ldots)$. The parameters of the PDFs are estimated from the training data by the maximum likelihood method (Theodoridis and Koutroumbas, 1999). Fig. 3(a) shows the fitted PDF for two hypothetical ALC, ω_1 and ω_2, in one dimension. The covariances among the individual acoustic features are considered as well as the variances. This means that the contours of the distributions will be hyper-ellipses, whose axes are not necessarily parallel to those of the property-space axes (Fig. 4, lower panel). The use of covariance in addition to variance provides a more comprehensive set of distribution functions and better separation between the different ALC.

The acoustic feature library used to identify categories is based on high quality acoustic measurements combined with trustworthy species identification, such as that obtained from targeted trawls that caught only one species, or were dominated by one species. Although trawl catches may not be representative of the local species mix, if one species dominates the catch then its abundance is clearly significant in that area. Thus, given some basic knowledge of biota scattering properties, the use of such ground-truthed measurements reduces the probability of assigning calculated acoustic features to the wrong ALC.

To develop an acoustic feature such as the $R(f)$, multi-frequency S_v training data are manually selected from rectangular sections of the echogram, and only values simultaneously above specified S_v thresholds at all available and chosen frequencies are used further in the training process. Next,

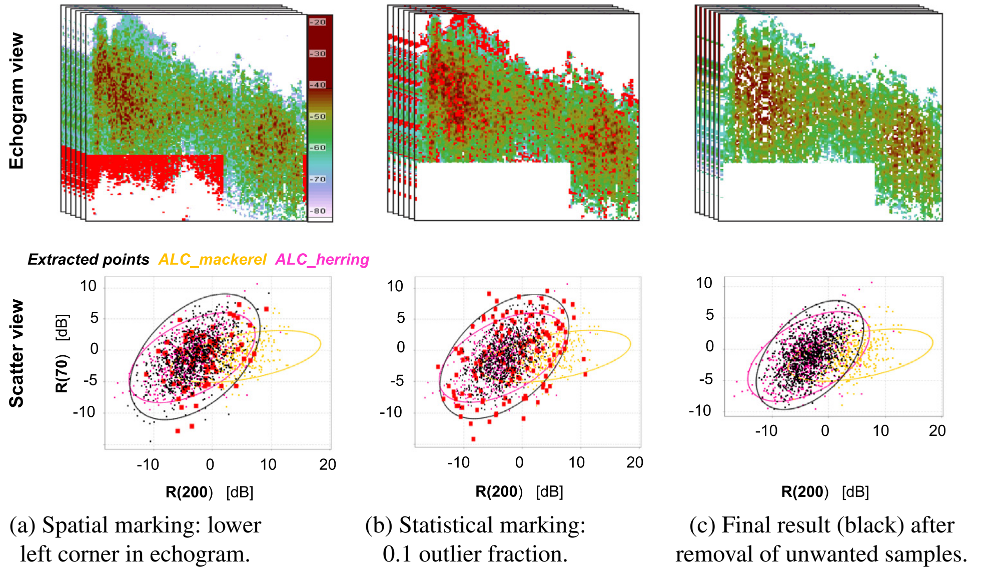

Fig. 4. Stepwise removal of unwanted samples from a school when populating the acoustic feature library. Echogram data from upper panel is shown as small black points in the lower panels. Fitted probability density functions (PDF) are shown as ellipses in the lower panels. *ALC_mackerel* (orange) and *ALC_herring* (purple) from the acoustic feature library are shown for comparison (Fig. 5). (a) Spatial selection of data: data to be removed is selected in the echogram. The selected data are visualized in red. (b) Statistical selection of data: statistical outliers are marked in red, being the data outside the hyper-ellipsoid encapsulating the 0.9 fraction of the data, i.e. 90% of the data should be inside. (c) Final result to be stored in the feature library.

unwanted data are identified and removed from the training dataset, such as observations near the edges of a school as these values may include sampling volume that are not representative of the acoustic target (Fig. 4(a)). Then, a statistical outlier fraction of the remaining data that deviates most from the multi-dimensional normal distribution (Fig. 4(b)) is removed from the training dataset. The remaining data (Fig. 4(c)) are stored in the feature library at different resolutions, referred to as pixel, cell, school, or track. The narrowest curves with the highest peak in Fig. 5 show the currently selected samples (in black) and the best-fit normal distribution for one feature, $R(200)$, for the same data (as in Fig. 4(c)). The acoustic library may be edited or expanded at any time, should additional high quality ground-truthed data become available.

The pixels represent the original and highest spatial resolution, but they show the poorest separation between ALC due to having the largest variance and covariance between the acoustic features $R(f)$ and $S_v(38)$ (compare Fig. 6(a)–(c)). Here, a cell is defined as 16 pings horizontally by 50 pixels vertically, giving 800 samples covering approximately 82 m $\times$ 9.5 m of the echogram, given a ping rate of 1 s^{-1}, a cruising speed of 10 knots and a vertical resolution of 0.19 m. The cells are required to be at least 50% full, that is with 400–800 samples in each cell with $s_v(38) > -82$ dB re 1 m^{-1} being available simultaneously at all frequencies. Data at school resolution are averaged over a school, and data at track resolution are averaged over a single fish track.

The acoustic features of several ALC may not be well separated (e.g. *ALC_herring* and *ALC_capelin*, Fig. 6(a)). Backscatter from a single fish track clearly belongs to a particular species, and those from a region identified as a school are likely to belong to just one species (Pitcher, 1993). Averaging the acoustic data over school cells or entire schools reduces the variance of the $s_v(38)$ and $r(f)$ data, thus the acoustic features of the ALC become better separated as the data aggregation level goes from pixel to cell to school (Fig. 6(a)–(c)). The splitting of a school into cells where each cell contains a fixed number of samples is partly used to check if a school seems to contain more than one species (Korneliussen et al., 2009a), and partly to standardize the number of samples used to calculate the acoustic features.

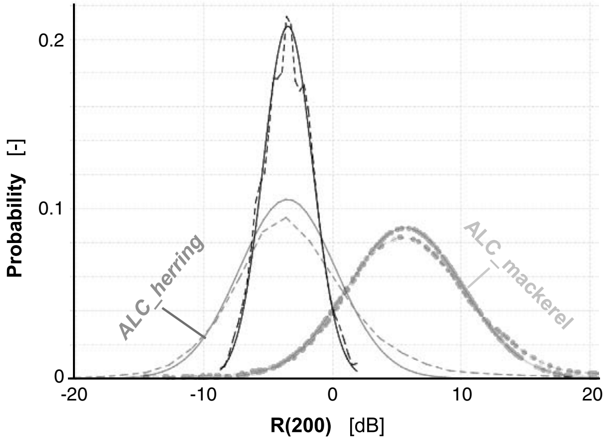

Fig. 5. Final result of data in Fig. 4(c) (black) after removal of unwanted samples. Normal distributions (continuous lines) fitted to real data (stippled). Selected data in narrow curves left, *ALC_herring* in grey wide curves to left, *ALC_mackerel* in curves to right.

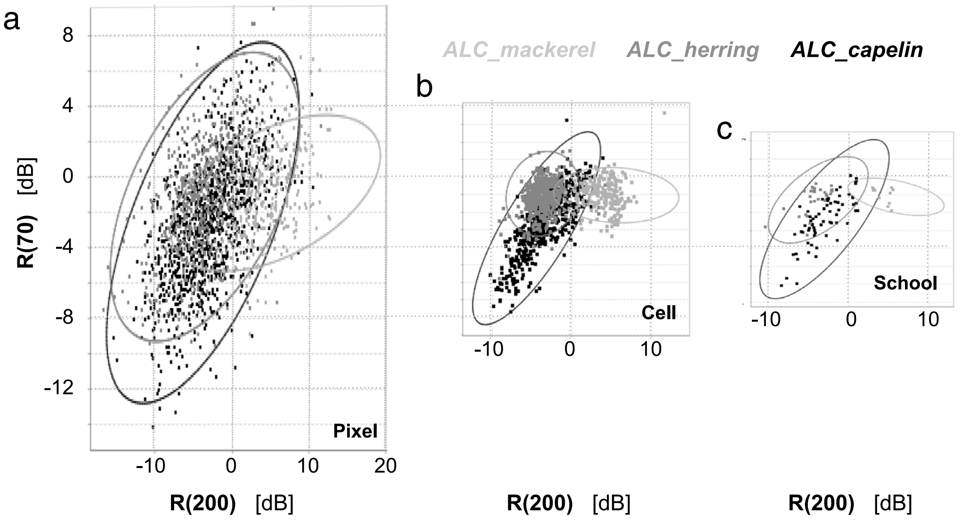

Fig. 6. Acoustic features of three ALC and the $R(70)$ and $R(200)$ different spatial resolutions: (a) Pixel resolution; (b) Cell resolution (800 pixels averaged spatially); (c) School resolution (one value per identified school). The separation of the categories improves as the resolution decreases.

Fig. 6 shows two of the hyper-ellipse dimensions, $R(70)$ and $R(200)$, in the training dataset for *ALC_mackerel*, *ALC_herring* and *ALC_capelin*, and compares the results at three resolutions—pixels, cells, and entire schools. The variances (the 10% percentile) that appear as sizes of the ellipses are clearly larger at pixel resolutions than at cell and school. Several of the ALC can be completely separated when visualized at cell resolution, in which case the automatic separation of these categories is considered reliable (Fig. 7). The reliable identification of some ALC depends on particular features. For example the category *ALC_mackerel* needs $R(200)$ and *ALC_resonant_18* needs $R(18)$, so there is no attempt to identify *ALC_mackerel* beyond the maximum range accessible at 200 kHz. Similarly, there is no attempt to identify *ALC_resonant_18* if 18 kHz data is not available. ALC such as *ALC_herring* may be identified beyond 300 m by using the features based on frequencies still valid at larger ranges. To compensate for this range limitation, KORONA can merge data from hull-mounted

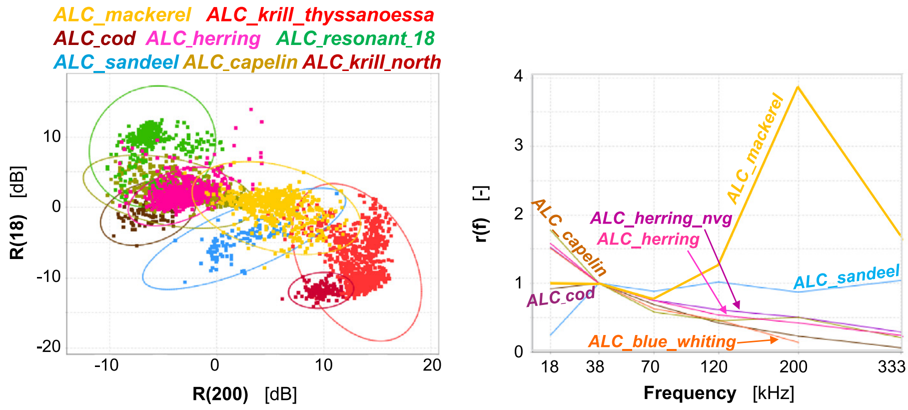

Fig. 7. (a) Relative frequency responses $R(18)$ vs. $R(200)$ at cell resolution (800 samples averaged spatially) for the ALC of importance for the case studies. The ellipses are statistical boundaries on each category that contain 90% of the observations. The plot shows that for the features $R(18)$ vs. $R(200)$ *ALC_mackerel* is almost completely separated from *ALC_resonant_18*, while *ALC_herring* has partial overlap with both *ALC_mackerel* and *ALC_resonant_18*. The use of more features in addition to $R(18)$ and $R(200)$ improves the separation of categories. (b) The categories are generated from ground-truthed backscatter of *Scomber scombrus* L. (*ALC_mackerel*), *Ammodytes marinus* L. (*ALC_sandeel*), *Mallotus villosus* L. (*ALC_capelin*), *Gadus morhua* L. (*ALC_cod*), *Micromesistius poutassou* L. (*ALC_blue_whiting*) and *Clupea harengus* L. (*ALC_herring* and *ALC_herring_nvg*). The *ALC_herring_nvg* category is trained on Norwegian spring spawning herring, while the *ALC_herring* category also includes North Sea herring.

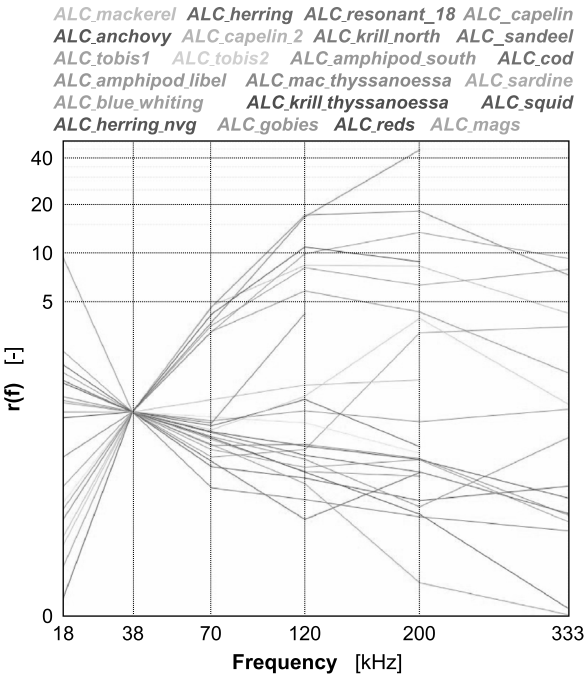

Fig. 8. $r(f)$ of acoustic library categories currently available in LSSS.

systems and deep towed systems, in which case data at higher frequencies may be used at any depth within the range limitations of the frequencies from the towed vehicle.

The content of some of the ALC can be visualized as scatter plots (Fig. 7(a)) or as $r(f)$ (Fig. 7(b), Fig. 8).

3.2.2. Identification of ALC in small volumes

The estimation of the most probable ALC to be associated with a pixel is to some extent subjective, but we still use the term probability even if it is not a rigorous statistical probability. The subjectivity comes from, for example, operator decisions about how strong the requirement of belonging to the same category as the adjacent pixels should be (**cf. step 7 below**), and how much (if any) adjustment of the á priori probability (**cf. step 5**) there should be (e.g. that capelin is less likely to be found in the Norwegian Sea than in the Barents Sea, and thereby subjectively setting the maximum possible probability $P_{\mathrm{apriori}} < 1$), although $P_{\mathrm{apriori}} = 1$ for all case studies here.

The categorization method is based on Bayes decision theory (Theodoridis and Koutroumbas, 1999). New backscatter measurements are compared to the content of the feature library to decide which, if any, of the existing ALC the new measurements are most similar to. The categorization of $R(f)$ and $S_v(38)$ measurements from a pixel proceeds as described below. The ALC with the highest probability is allocated to a pixel (but the other candidates are also kept), provided the probability is above a specified minimum value (**cf. step 2**).

The outlier fraction sets the cut-off for the measured pixel. When there are several ALC active, the most likely category (as shown in Fig. 3(b)) is associated with the pixel. This is done for all pixels. Further, a pixel is more likely to be occupied by the same species as the majority of nearby pixels, which is why the probability is adjusted according to the categories of the nearest spatial neighbours (**cf. step 7**).

Algorithm

Multi-frequency data from each pixel are tested against the properties of the ALC. The most probable ALC is associated with that pixel. This is determined by the probability that the tested data belongs to a particular ALC. The ALC categorization proceeds as follows for all pixels by means of features selected and valid at the current range:

(0) Find the subset of useable ALC, $\Omega = \{\omega_1, \omega_2, \ldots, \omega_n\}$, based on user specified criteria in terms of available features. For example, categorization of *ALC_mackerel* needs $R(200)$, which is not valid for hull-mounted transducers beyond a depth of about 300 m.
(1) Assign special ALC to pixels that should not be categorized using the acoustic features. These special categories include transducer near-field, return from the seafloor, noise, etc.
(2) Find the probability density function (PDF) threshold for each ALC based on an outlier fraction, normally 0.1. Figs. 4, 6 and 7 show the borders corresponding to the probability minima visualized as ellipses.
(3) Compute the feature vector x, for example $[R(18), S_v(38), R(70), R(200)]$ in case of a 4-dimensional feature space.
(4) For each ω_i (i.e. ALC), calculate the ALC-dependent probability density $p(x|\omega_i)$. Fig. 3(a) shows probability density functions (PDF)s for the feature W for categories ω_1 and ω_2 for a 1-dimension feature space.
(5) Calculate a discriminant function $g(\omega_i|x) = P(\omega_i) \cdot p(x|\omega_i)$, where $P(\omega_i)$ is a user-defined á priori probability for the category ω_i. For the case studies $P(\omega_i) = 1$.
(6) Assign to the pixel the ALC ω_i with the largest discriminant g. Leave the volume unassigned if the measurement is not above the PDF-threshold from step (2) of any category.
(7) Update the á priori probability of each category by considering neighbouring pixels (Besag, 1986). The updated probability, $P'(\omega_i)$, is calculated using the formula $P'(\omega_i) = P_0(\omega_i)e^{\beta(n_i-3)}$, where $P_0(\omega_i)$ is the original á priori probability, n_i is the number of neighbouring pixels assigned category ω_i (maximum of 8), and β, with default value 0.1, is a parameter that determines how much the probability should change. Thus, the probability is modified by the factor $e^{\beta(n_i-3)}$, which is greater than unity if more than three neighbouring pixels have the same category.
(8) Steps (5)–(7) may be repeated several times (usually three). New categorizations hence take into account previous changes to the categorization of neighbouring pixels.

3.2.3. Identification of ALC in sample collections

The algorithm for identification of ALC in sample collections is the same as for small volumes except that step (7) is omitted. Categorization of data inside detected schools and single fish tracks are

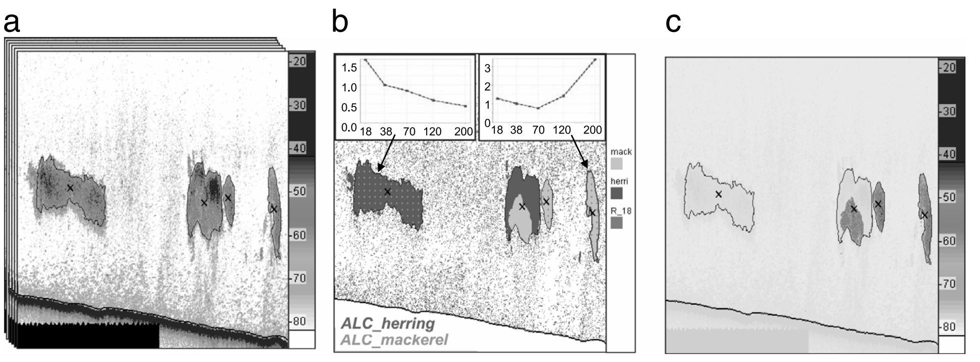

Fig. 9. Results of acoustic library categorization. (a) Echograms at six frequencies 18–364 kHz on overlaid sheets, 38 kHz on top. (b) Identified ALC shown with solid colours (grey for *ALC_herring*, light grey for *ALC_mackerel*) with relative frequency responses, $r(f)$, for two schools shown. (c) Synthetic 38 kHz echogram with *ALC_mackerel* kept and all others removed.

more highly averaged than at pixel resolution. The categories *ALC_capelin* and *ALC_herring* are hardly distinguishable at pixel level (Fig. 6(a)), but are better distinguished at the cell level (Fig. 6(b)) are even better distinguished at the school level (Fig. 6(c)). To ensure that a school is monospecific, it is split into cells and each cell is categorized as one unit (as described by Korneliussen et al., 2009a). The school categorization uses all pixels inside an aggregation, (e.g. School 1 in Fig. 9(b)), including those that would be uncategorized at pixel level. If the cell categorization indicates a multiple species aggregation, then pixel categorization is used as for School 2 in Fig. 9(b). Trawl samples targeting the schools in Fig. 9 contained both mackerel and herring, and therefore supported the results of a two species aggregation (or two schools in close proximity).

A disadvantage of testing all schools for multiple species as described above is that the result may be inconclusive, in which case the pixel-categorization will be used for the school, which in turn results in many uncategorized pixels inside the schools. Examination of echograms indicates that schools in Norwegian waters almost always contain only one species. Therefore, the schools are considered to be monospecific in the analysis.

The analysis of data from a single fish track is similar to that of schools, but the analysis is based on the backscattering cross section, σ_{bs}, instead of s_v. The knowledge that a track clearly contains a single specimen, combined with the use of within-beam location data to correct for the transducer beam pattern, reduces the variance of the relative frequency responses. Schools require more samples than single fish tracks for good separation between acoustic categories, but a school contains far more samples than a single fish track.

3.2.4. Model-based categorization of aggregated samples

The MBC algorithm uses several scattering models to estimate backscatter characteristics from zooplankton, such as euphausiids and amphipods. These currently include the randomly oriented fluid bent cylinder model (Stanton et al., 1993) and the stochastic distorted-wave Born approximation (Demer and Conti, 2003). Those models have been implemented in an optimized framework (Fernandes et al., 2002, sections 6.2.3 to 6.2.4). Zooplankton classification may be applied to pixels, but is preferably done on cells or schools to avoid large variations in estimated sizes. We use the term pixel to refer to this minimal volume being the highest resolution of the echosounder/transducer combination. To speed up the processing, data allocated to ALC that are clearly not zooplankton are ignored. Note that the results of model-based categorization are not used in the case studies in this paper.

4. Case studies

4.1. Material and methods

Multi-frequency acoustic data were collected according to existing protocols (Korneliussen et al., 2008), with the intention of combining and analysing co-incident observations from several echosounders operating at different frequencies. This requires the use of calibrated systems (Foote et al., 1987; Demer et al., 2015) operating at the same pulse-duration (in these data 1 ms) and ensonifying nearly the same volume of water at all frequencies. Data were collected during three surveys using Simrad EK60 echosounders mounted in a tightly-packed configuration on a protruding keel (Korneliussen, 2010). The 18, 38, 70, 120, and 200 kHz echosounder frequencies were used on all surveys. In addition, a 364 kHz echosounder was used in 2004 and a 333 kHz on the 2014 sandeel survey. The 18 kHz system provided a transducer beam width of 11° between the half-power points, while the other frequencies provided 7° beamwidths. All surveys were used for stock abundance estimation based on the manual categorization. The surveys were:

1. RV G.O. Sars in October–November 2004 in the northern part of the North Sea and southern part of the Norwegian Sea targeting Atlantic mackerel. The survey was selected to compare the acoustic library categorization method with an earlier method (Korneliussen, 2010).
2. FV Eros in 2014 in the North Sea targeting sandeel. The survey was selected because it observed small schools on and just above the seabed, which provides a challenge to the categorization process.
3. RV G.O. Sars in the Barents Sea in 2014 for investigation of the major components of the ecosystem. The survey was selected for identification of capelin, which is acoustically similar to herring and cod in the same waters.

The acoustic data were manually categorized during the surveys (Fig. 1(A)) and were reprocessed with a standardized setup (Fig. 1(B) and (C)). LSSS used the same processing setup in all cases, to remove noise spikes and ambient noise (Korneliussen, 2000). Spatial offsets between the transducers and pulse-transmission delay between frequencies were corrected (Korneliussen et al., 2008) for the purpose of comparing data at different frequencies in small volumes. These reprocessed data were used during the manual categorization (Fig. 1(C)).

The *manual categorization* was carried out daily within 24 h after it was collected and then stored in scrutiny files. Those scrutiny files were loaded into LSSS and applied to the reprocessed acoustic data (Fig. 1(C)) and stored into a *manual categorization database*. The results of the automatic categorization were ignored.

The reprocessed data were also used during the automatic acoustic categorization (Fig. 1(B)), but the processing continued by detecting ALC at pixel-level, and by detecting the extent of echo-traces. These are investigated as a unit when detecting ALC. Finally, the MBC are detected. An echo-trace based ALC is used before a pixel-based ALC. Each pixel becomes associated with an ALC thus:

- All pixels: ALC with probability based on pixel-data (even within echo-traces);
- Pixels within echo-traces: ALC with probability based on echo-trace data;
- All pixels: A single MBC with probability (not used in case studies).

The settings for pixel categorization and school categorization were the same for all surveys, except that the selection of ALC was different due to geographical considerations (see Table 1).

The process of *automatic acoustic categorization* is automated (i.e. no operator intervention), and the resulting categories (MSC) stored in an *automatic categorization database* (Fig. 2(B)). The process starts with automatic acoustic categorization, which results in ALC and MBC for each pixel. Furthermore, those categories are mapped into manual scrutiny categories (MSC) chosen by the operator prior to the start of the process. For example, in the 2014 Barents Sea survey, *ALC_nvg_herring* was mapped to the herring MSC.

The s_A was stored to the databases at a horizontal resolution of 0.1 nautical miles both for the original manual scrutiny and for the automatic categorization. Table 1 shows the ALC used by the categorization modules in the different case studies.

4.2. Results

The s_A from the automatic categorization and the manual categorization for the three case studies are similar (Fig. 10), especially for large s_A. A notable result especially for capelin, is that the values from the automatic categorization are equal to or slightly smaller than the manual categorization

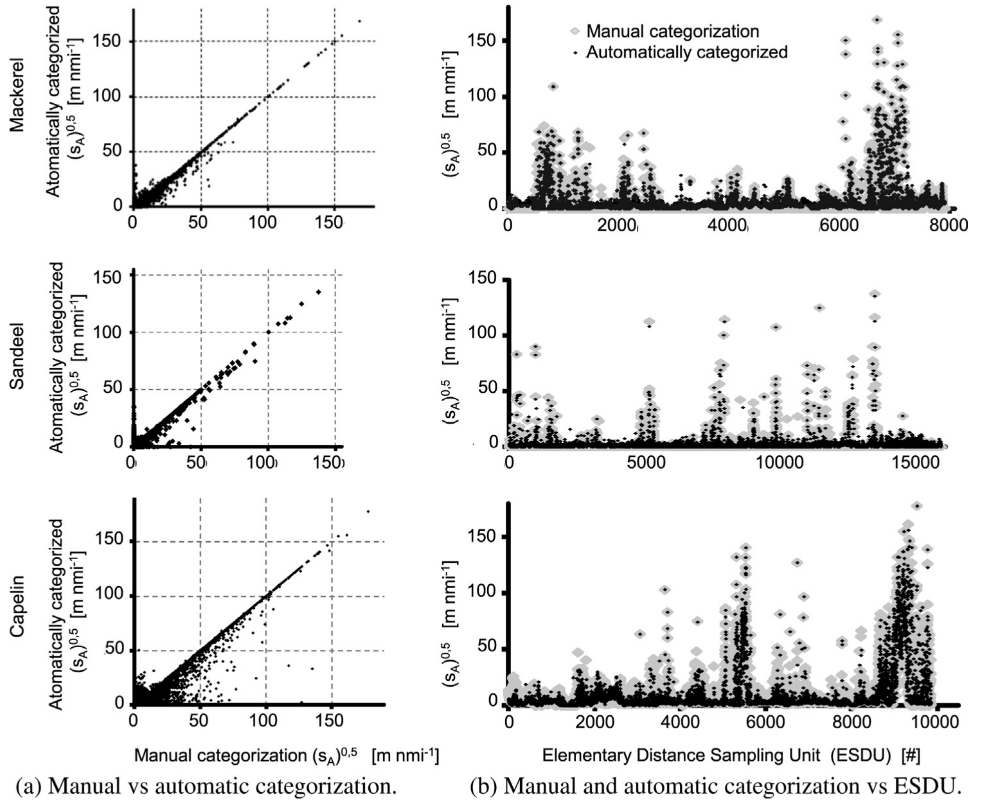

Fig. 10. Comparison of automatic categorization (ALC) and manual categorization (MSC) for Atlantic mackerel, sandeel, and capelin. s_A are stored in the database at a horizontal resolution of 0.1 nautical miles, and shown here as $s_A^{0.5}$ to better show the dynamics. (a) Automatically categorized ALC compared to manually categorized. (b) ALC (black) and MSC (light grey) as a function of distance travelled.

except for small values of s_A. Furthermore, there are many small to moderate values from the automatic categorization that were given zero s_A for the manual categorization, especially for sandeel. Subsequent inspection of the manually categorized sandeel survey echograms showed there were some sandeel schools that were clearly categorized incorrectly. This objective check of the manual categorization is one advantage of the automatic process. There was high spatial conformity between the manual categorization and the automatic categorization (Fig. 10(b)) except for low values of s_v where the conformity varied considerably.

The overall Pearson correlation coefficients between manual scrutiny and automatic categorization are very high for all surveys (Table 2). For mackerel and sandeel the total acoustic abundance of the automatic categorization and the manual categorization were within 2% of each other. This slight difference is explained by the observation that manually drawn school integration regions are always slightly larger than the automatically drawn school regions. The capelin data correlation is close to unity (Table 2), but the automatic categorization resulted in almost 15% lower s_A than the manual categorization. This may be due to school categorization failing so that some schools are not being categorized as a unit but as individual pixels. This in turn leaves many pixels uncategorized and some of the pixels being wrongly identified as cod or herring instead of capelin, which is quite acoustically similar. In that case, masking the measurements with pixels identified as capelin will give s_A values that are too low. This problem may be reduced if additional features are used in the categorization process although some uncertainty will always remain.

Table 1
Categories and species compared in the case studies. The second column is the target species, and the third and fourth columns are the similar scrutinize category and ALC believed to originate from the target species. The final column contains ALC that were identified simultaneously with the target library category. If the backscatter of a pixel is not most likely to originate from the target ALC, or if the probability is below the probability threshold, $s_v = 0$ for that pixel.

Case study survey number	Target species	MSC[g]	ALC[g]	Number of training schools	Training schools collection year	ALC used to remove backscatter
2004113	*Scomber scombrus* L.	Mackerel	*ALC_mackerel*	35	1999 2005	*ALC_resonant_18*[a] *ALC_krill_north*[b] *ALC_krill_thyssanoessa*[c] *ALC_herring*[d]
2014807	*Ammodytes marinus* L.	Sandeel	*ALC_sandeel*	44	2008 2009	*ALC_resonant_18*[a] *ALC_krill_north*[b] *ALC_krill_thyssanoessa*[c] *ALC_herring*[d]
2014116	*Mallotus villosus* L.	Capelin	*ALC_capelin*	67	2008 2014	*ALC_resonant_18*[a] *ALC_krill_north*[b] *ALC_krill_thyssanoessa*[c] *ALC_herring_nvg*[e] *ALC_cod*[f]

[a] *ALC_resonant_18* were trained from 4 shoals of zooplankton that were resonant at 18 kHz.
[b] *ALC_krill_north* were trained from 24 schools of *Meganictiphanes norwegica* L.
[c] *ALC_krill_thyssanoessa* were trained from 8 shoals of *Thyssanoessa inermis* L.
[d] *ALC_herring* were trained from 40 schools of *Clupea harengus* L.
[e] *ALC_herring_nvg* were trained from 9 schools of Norwegian Spring Spawning herring.
[f] *ALC_cod* were trained from 10 schools of *Gadus morhua* L. These were actually schools of cod.
[g] ALC are always in italic and preceded by *ALC_*, MSC are not italic.

Table 2
Comparison of manual scrutiny and automatic categorization. The $s_A(200)$ samples are averaged over 0.1 nautical miles. Column 6 gives number of samples, column 7 the total s_A of automatic detection relative to s_A of manual scrutiny.

Survey	Target species	MSC	ALC	Pearson correlation coefficient	$s_A(200)$ samples	$\Sigma s_A(200)_{auto} / \Sigma s_A(200)_{man}$
2004113	*Scomber scombrus* L.	Mackerel	*ALC_mackerel*	0.996	7 913	0.985
2014807	*Ammodytes marinus* L.	Sandeel	*ALC_sandeel*	0.998	17 925	0.994
2014116	*Mallotus villosus* L.	Capelin	*ALC_capelin*	0.976	9 869	0.855

The high correlation coefficients (Table 2) show that, at least for these surveys, the s_A (and thereby also the biological abundance) from the target species can be calculated with high conformity using automatic categorization method. For sandeel and Atlantic mackerel, there is insignificant difference in s_A, while automatic categorization of capelin was 15% less than manual categorization.

4.3. Discussion

The target ALC (*ALC_mackerel*, *ALC_sandeel* and *ALC_capelin*) were used in the three case studies, along with *ALC_herring*, *ALC_cod*, *ALC_blue_whiting* that have features with properties not very different from *ALC_capelin*. The mean $r(f)$ values of *ALC_mackerel* (Korneliussen, 2010) and *ALC_sandeel* (Johnsen et al., 2009) are quite different from the other ALC (Fig. 7), which made reliable automatic categorization for schools a not unexpected outcome. However, the sample variation of

$R(f)$ illustrated by the ellipses of the scatterplots (Figs. 4, 6 and 7) is much larger than the variation of mean $R(f)$ at school, cell and pixel resolution. This makes occasional misclassifications likely even for *ALC_mackerel* and *ALC_sandeel*, especially at pixel level. In other words, misclassification is more likely if the scattering regions cannot be grouped into single fish tracks or schools. The sample correlation coefficient is high for large s_A, thus the total correlation also becomes large. The less common alternative, the rank correlation coefficient (Walpole et al., 2002) gives equal weight to all values, and results in a somewhat lower correlation. However, it is the large values that have most influence on the abundance estimates, and it is therefore felt that the use of sample correlation is more appropriate.

The high correlation between manual categorization and automatic categorization of capelin was unexpected, especially since the $R(f)$ of *ALC_capelin* in the feature library is similar to that of other fish with swimbladder, such as *ALC_cod* and *ALC_herring_nvg* (Fig. 7(b) and Fernandes et al., 2002). The manual categorization of capelin is considered to be reliable in this case due to the low water temperature in which capelin is often found, the location in the water column, the schooling behaviour and the results of the biological sampling. It was expected that additional, non-acoustic characteristics would be needed for a high correlation between manual categorization and automatic categorization. Differences in school shapes have previously been used to improve automatic categorization (Scalabrin et al., 1996; Korneliussen et al., 2009a), but that was not used in this case as the scattering regions mostly formed large layers for capelin in the 2014 Barents Sea cruise, and not schools of limited size, so morphology was not considered to be reliable. A slight reduction in the á priori probability of herring and cod in the 2014 Barents Sea case study, for example from 1 to 0.9, while keeping the á priori probability of capelin as 1, may be defended due to the low water temperature. Such a reduction in the á priori probability did in fact make the s_A resulting from automatic categorization of capelin closer to manual categorization. However, the quantification of the á priori probability is not easy, and the á priori probability is therefore set to $P = 1$ for all ALC in all case studies.

The statistical properties of the features should be calculated from a sufficiently large quantity of data for each ALC, and preferably from more than one survey, to make the calculated values representative. Poor training of one ALC in the acoustic feature library can lead to biased results, as would the use of data from uncalibrated echosounders. Uncalibrated instruments result in biased s_v, and therefore an incorrect $r(f)$ (unless all $s_v(f)$ are equally wrong at all frequencies f). The uncritical use of fully automated categorization during operational surveys is inadvisable. Although the case studies showed good agreement between automatic categorization and manual categorization, the $r(f)$ may potentially vary with water temperature, fat content, specimen size and species behaviour, so that similarly good results may not be found in other cases. Ona (1990) and Ona et al. (2001) showed that TS (and thereby s_v) vary with physiology and with season, and unless $s_v(f)$ varies similarly at all frequencies, $r(f)$ will also vary somewhat. A solution to this problem is to build an acoustic library that allows for seasonal variations of scattering properties. Alternatively, one species may be used to build up several ALC, such as *ALC_mackerel_spring*, *ALC_mackerel_summer*, etc.

A comprehensive study of acoustic features, such as $r(f)$, was made by Fernandes et al. (2002) for several species, and those $r(f)$ may be used as a reference during scrutiny of multi-frequency echograms. De Robertis et al. (2010) also have a list of scattering classes that could be used during categorization (their Fig. 4), but they were able to integrate all $r(f)$ into a single Z-score, and used that to make an operational system for multi-frequency backscatter classification. Like De Robertis et al. (2010), we also integrate many features into a discriminant function used to identify ALC in real-time, and like them we may expand the features we use with additional $r(f)$ for additional frequencies, f. Our method has additional flexibility in that we can also use $s_v(38)$, and may expand the features used (for example depth). Unlike De Robertis et al. (2010), we extract the statistical properties of the features directly from measured data, and can easily update the statistical properties during the survey in an operational sense.

A future enhancement of our method will incorporate more of the information that the operator has at hand, such as the local geography, and use it to automatically adjust the á priori probability (e.g. reducing the probability of finding Atlantic mackerel in the Barents Sea during the winter). Based on geography (and an ALC mapped to a species), the á priori probability could be multiplied by 1 within the main residential area of a species, and gradually reduced outside that area. The methods for acoustic categorization presented here can be applied to wideband acoustic data, although this has not been shown here. Wideband data can improve range resolution and hence improve single fish detection and tracking. Wideband backscatter can also increase the number of points in the $r(f)$ curve, potentially leading to improved species identification. Here wideband is defined to consist of one or several continuous sub-bands, where each sub-band (called broadband) obeys approximately to the ratio bandwidth/center_frequency ≥ 0.5.

The two methods both give high conformity between manual categorization and automatic categorization, although the presented results used different Elementary Sampling Distance Unit (ESDU). A natural application of the high conformity between manual categorization and automatic categorization shown here and by De Robertis et al. (2010) is that automatic categorization may be used to reanalyse data from old surveys to obtain a new and more objective view of those data. Furthermore, our method may be used to analyse data that has no temporally relevant ground-truthed information available, such as backscatter data collected by autonomous underwater vehicles or ocean observatories.

Although no comprehensive comparison for a survey has been done, the 2014 sandeel survey typically required 10 h per 24 h data for manual scrutiny of echograms for the few days with the most complex registrations, while it was carried out in less than 1 h when manual categorization was assisted with automatic categorization, and 5 min (excluding 20 min processing) when fully automated. The 5 min is time needed for loading data and storing the results of the categorization to the database.

5. Summary and future expectations

Automatic categorization is reliable for the large s_A of the ALC (or species) investigated in the case studies. Although the manual categorizations, are kept separate from automatic categorization for these case studies, that would not be the case for most future survey analyses. The automatic acoustic categorization provides a more objective view that assists the manual categorization and creates a common baseline for categorization. The speed and objectivity of the categorization naturally also improves with the use of reliable echo trace detection, as the acoustic library categorization becomes more reliable, and because less time is needed to draw school regions. Thus, more time is available for the most difficult part of the categorization. During routine categorization it is a design aim of the software that the operators (not the software) decide if the automatic acoustic categorization should be used or whether the categorization should be wholly manual, or a combination of the two. One reason for this is that automatic acoustic categorization may identify species in small volumes that are difficult to see or too labour intensive for manual categorization.

Automatic categorization can provide a more objective view to assist manual categorization. Efficient and effective categorization of acoustic data in a resource limited environment requires that attention be paid to the time required to produce the estimates of acoustic abundance, s_A, used to estimate fish stock size. The use of automatic categorization contributes to this aim for surveys which require estimates of fish stock size at completion.

The improved scrutiny objectivity obtained from automatic or semi-automatic data analysis and the ability to process, in a timely manner, the increased volume of acoustic data during surveys is advantageous. Automatic analysis produces consistent and reproducible results, and enables improved assessment and understanding of the spatial distribution of ecosystem components. Such automatic processing makes practical the timely analysis of the very large quantities of data that will be collected by autonomous vehicles (Fernandes et al., 2003; Ghani et al., 2014; Suberg et al., 2014); and stationary systems (Pawlowicz and McClure, 2010). Much of these data can be processed automatically as described here, with the manual work being simply verification of the processing.

6. Funding

The work is funded by the Institute of Marine Research, which is a Norwegian governmental institution, and by Christian Michelsen Research.

References

Anon 1990. SIMRAD EK500 Scientific echo sounder. Instruction manual. P2172E, 1990–1997. Contents: Operators manual, P2170E Installation Manual P2180E, Service Manual 2161E.

Besag, J., 1986. On the statistical analysis of dirty pictures. J. R. Stat. Soc. Ser. B Stat. Methodol. 48 (3), 259–302.

De Robertis, A., McKelvey, D.R., Ressler, P.H., 2010. Development and application of an empirical multifrequency method for backscatter classification. Can. J. Fish. Aquat. Sci. 67, 1459–1474.

Demer, D.A., Berger, L., Bernasconi, M., Bethke, E., Boswell, K., Chu, D., Domokos, R., et al., 2015. Calibration of acoustic instruments. ICES Cooperative Research Report No. 326. p. 130.

Demer, D.A., Conti, S.G., 2003. Reconciling theoretical versus empirical target strengths of krill: effects of phase variability on the distorted-wave Born approximation. ICES J. Mar. Sci. 60, 429–434.

Deuser, L., Middleton, D., Plemons, T., Vaughan, J., 1979. On the classification of underwater acoustic signals. II. Experimental applications involving fish. J. Acoust. Soc. Am. 65, 444–445.

Dougherty, E.R., 1992. An Introduction to Morphological Image Processing. ISBN: 0-8194-0845-X.

Fernandes, P.G., Korneliussen, R.J., Lebourges-Dhaussy, A., Masse, J., Iglesias, M., Diner, N., Ona, E., et al., 2002. The SIMFAMI project: species identification methods from acoustic multi-frequency information. Final Report to the EC Number Q5RS-2001-02054 (Report accessible from any of the authors).

Fernandes, P.G., Stevenson, P., Brierley, A.S., Armstrong, F., Simmonds, E.J., 2003. Autonomous underwater vehicles: future platforms for fisheries acoustics. ICES J. Mar. Sci. 60, 684–691.

Foote, K.G., Knudsen, H.P., Vestnes, G., MacLennan, D.N., Simmonds, E.J., 1987. Calibration of acoustic instruments for fish density estimation: a practical guide. ICES Cooperative Research Report, 144. p. 69.

Ghani, M.H., Hole, L.R., Fer, I., Kourafalou, V.H., Wienders, N., Kang, H., Drushka, K., Peddie, D., 2014. The SailBuoy remotely-controlled unmanned vessel: Measurements of near surface temperature, salinity and oxygen concentration in the Northern Gulf of Mexico. Methods Oceanogr. 10, 104–121.

Handegard, N.O., 2007. Observing individual fish behavior in fish aggregations: Tracking in dense fish aggregations using a split-beam echosounder. J. Acoust. Soc. Am. 122 (1), 177–187.

Haralabous, J., Georgakaros, S., 1996. Artificial neural network as a tool for species identification fish schools. ICES J. Mar. Sci. 53, 173–180.

Holliday, D.V., 1977. Extracting bio-physical information from the acoustic signature of marine organisms. In: Andersen, N.R., Zahuranec, B.J. (Eds.), Oceanic Sound Scattering Prediction. Plenum, NY, pp. 619–624.

Holliday, D.V., Pieper, R.E., Kleppel, , 1989. Determination of zooplankton size and distribution with multifrequency acoustic technology. J. Int. Counc. Explor. Sea 46, 52–61.

Horne, J.K., 2000. Acoustic approaches to remote species identification: a review. Fisheries Oceanography 9 (4), 356–371.

Jech, J.M., Michaels, W.L., 2006. A multifrequency method to classify and evaluate fisheries acoustics data. Can. J. Fish. Aquat. Sci. 63 (10), 2225–2235.

Johnsen, E., Pedersen, R., Ona, E., 2009. Size-dependent frequency response of sandeel schools. ICES J. Mar. Sci. 66, 1100–1105.

Kaartvedt, S., Røstad, A., Fiksen, Ø., Melle, W., Torgersen, T., Breien, M.T., Klevjer, T.A., 2005. Piscivorous fish patrol krill swarms. Mar. Ecol. Prog. Ser. 299, 1–5.

King, M., 2007. Fisheries Biology, Assessment and Management. Wiley-Blackwell, Oxford, p. 400.

Kloser, R.J., Ryan, T., Sakov, P., Williams, A., Koslow, J.A., 2002. Species identification in deep water using multiple acoustic frequencies. Can. J. Fish. Aquat. Sci. 59, 1065–1077.

Korneliussen, R.J., 2000. Measurement and removal of echo integration noise. ICES J. Mar. Sci. 57, 1204–1217.

Korneliussen, R.J., 2004. The Bergen echo integrator post-processing system, with focus on recent improvements. Fish. Res. 68, 159–169.

Korneliussen, R.J., 2010. The acoustic identification of Atlantic Mackerel. ICES J. Mar. Sci. 67, 1749–1758.

Korneliussen, R.J., Diner, N., Ona, E., Berger, L., Fernandes, P.G., 2008. Proposals for the collection of multifrequency acoustic data. ICES J. Mar. Sci. 65, 982–994.

Korneliussen, R.J., Heggelund, Y., Eliassen, I.K., Johansen, G.O., 2009a. Acoustic species identification of schooling fish. ICES J. Mar. Sci. 66, 1111–1118.

Korneliussen, R.J., Heggelund, Y., Eliassen, I.K., Øye, O.K., Knutsen, T., Dalen, J., 2009b. Combining multibeam-sonar and multifrequency echosounder data: examples of the analysis and imaging of large euphausiid schools. ICES J. Mar. Sci. 66, 991–997.

Korneliussen, R.J., Ona, E., 2002. An operational system for processing and visualising multi-frequency acoustic data. ICES J. Mar. Sci. 59, 293–313.

Korneliussen, R.J., Ona, E., 2003. Synthetic echograms generated from the relative frequency response. ICES J. Mar. Sci. 60, 636–640.

MacLennan, D.N., Fernandes, P.G., Dalen, J., 2002. A consistent approach to definitions and symbols in fisheries acoustics. ICES J. Mar. Sci. 59, 365–369.

MacLennan, D.N., Holliday, D.V., 1996. Fisheries and plankton acoustics: past, present, and future. ICES J. Mar. Sci. 53, 513–516.

Ona, E., 1990. Physiological factors causing natural variations in acoustic target strength of fish. J. Mar. Biol. Assoc. UK 70, 107–127.

Ona, E., Zhao, X., Svellingen, I., Fosseidengen, J.E., 2001. Seasonal Variation in Herring Target Strength. Herring: Expectations for a New Millennium 461 Alaska Sea Grant College Program AK-SG-01-04, 461-487.

Pawlowicz, R., McClure, B., 2010. Inverted echosounder for continuous high-resolution water column profiling from the NEPTUNE (Canada) ocean observatory, OCEANS 2010 MTS/IEEE SEATTLE. Presented at the OCEANS 2010 MTS/IEEE SEATTLE, pp. 1–8.

Pedersen, G., Korneliussen, R.J., 2009. The relative frequency response derived from individually separated targets of northeast Arctic cod (Gadus morhua), saithe (Pollachius virens), and Norway pout (Trisopterus esmarkii). ICES J. Mar. Sci. 66, 1149–1154.

Pitcher, T.J., 1993. Behaviour of Teleost Fishes. In: Fish and Fisheries Series, vol. 7. Chapman and Hall, London, p. 715.

Scalabrin, C., Diner, N., Weill, A., Hillion, A., Mouchot, M.-C., 1996. Narrowband acoustic identification of monospecific fish shoals. ICES J. Mar. Sci. 53, 181–188.
Simmonds, E.J., MacLennan, D.N., 2005. Fisheries Acoustics: Theory and Practice, second ed. Blackwell, Oxford, p. 437.
Stanton, T.K., Chu, D., Wiebe, P.H., Benfield, M.C., Clay, C.S., 1993. Average echoes from randomly oriented random-length finite cylinders: Zooplankton models. J. Acoust. Soc. Am. 94 (6), 3463–3472.
Stanton, T.K., Wiebe, P.H., Chu, D., Benfield, M.C., Scanlon, L., Martin, l., Eastwood, R.L., 1994. On acoustic estimates of zooplankton biomass. ICES J. Mar. Sci. 51, 505–512.
Suberg, L., Wynn, R.B., Kooij, J. van der, Fernand, L., Fielding, S., Guihen, D., Gillespie, D., Johnson, M., Gkikopoulou, K.C., Allan, I.J., Vrana, B., Miller, P.I., Smeed, D., Jones, A.R., 2014. Assessing the potential of autonomous submarine gliders for ecosystem monitoring across multiple trophic levels (plankton to cetaceans) and pollutants in shallow shelf seas. Methods in Oceanogr. 10, 70–89.
Theodoridis, S., Koutroumbas, K., 1999. Classifiers based on Bayes theory. In: Pattern Recognition, third ed. Academic press, Oval Road, London NW1 7DX, UK, pp. 24–28 (Chapter 2). Classifiers based on Bayes decision theory. http://www.amazon.com/Pattern-Recognition-Sergios-Theodoridis/dp/0126861404.
Walpole, R.E., Myers, R.H., Myers, S., Ye, K., 2002. Probability & Statistics for Engineers & Scientists. Prentice-Hall, New Jersey, p. 730.