ELSEVIER

Computers and Electronics in Agriculture
31 (2001) 151–168

Computers and electronics in agriculture

www.elsevier.com/locate/compag

An automatic image-based system for estimating the mass of free-swimming fish

J.A. Lines [a,*], R.D. Tillett [a], L.G. Ross [b], D. Chan [a,b], S. Hockaday [b], N.J.B. McFarlane [a]

[a] *Silsoe Research Institute, Wrest Park, Silsoe, Bedford, MK45 4HS, UK*
[b] *Institute of Aquaculture, University of Stirling, Stirling, FK9 4LA, UK*

Received 3 August 2000; received in revised form 28 September 2000; accepted 30 October 2000

Abstract

The component parts of an image analysis based system for estimating the mass of swimming fish have been developed and tested under a limited range of conditions. They comprise automatic techniques for identifying video frames, which contain good images of fish, for identifying the outline of the fish in 3D space, and for determining mass from linear measurements taken from this outline. The major challenges in developing these techniques have been the variable and complex nature of images collected in sea cages, and in identifying morphological models of salmon which are invariant with strain, maturity and culture conditions. Tests show that salmon mass can be estimated from linear dimensions with a mean error of less than 0.5% and that linear dimensions can be extracted automatically from stereo images with a mean error of less than 10%. Preliminary tests of the complete system used 60 images of 17 fish in the weight range 0.7 to 5.7 kg. The mean mass measurement error was 18% with a standard deviation of only 9%. Development of the algorithms is continuing.

Keywords: Aquaculture; Biomass; Image analysis; Stereo vision

* Corresponding author. Tel.: +44 1525 860 000; fax: +44 1525 860 156.
E-mail address: jeff.lines@bbsrc.ac.uk (J.A. Lines).

0168-1699/01/$ - see front matter
PII: S0168-1699(00)00181-2

1. Introduction

Most farmed salmon in Norway, Scotland, Chile, Canada and Ireland are grown to harvest weight in sea cages. These are net pens floating in the sea, which typically enclose a sea volume of between 10^3 and 10^4 m^3 and contain between 10^4 and 10^5 fish. Salmon are grown in these cages for about two years.

One of the most significant pieces of information required by salmon farmers is the size of their fish. Information on both the average weight and on the weight distribution is required to guide decisions on grading, feeding and harvesting. The traditional measurement method is to catch a small sample of fish in a hand net. These fish are then anaesthetised and weighed. This process is time consuming, it stresses and damages the stock and it produces poor quality results since the sample taken is both small and biased. It is generally considered that weight estimates based on hand netting have an error of 1–12% for the sample and 15–25% for the population (Ross et al., 1998). In recent years, two types of device have become available which solve some of the problems of the traditional sampling methods. One of these comprises a frame somewhat like a thick picture frame through which the fish swim. This frame is crossed by a network of infrared beams. The profile of the fish as it swims through the unit is identified from the beams that are broken and an algorithm estimates the mass from linear dimensions of the fish. Questions remain however about the bias introduced into the sampling because of the fishes' aversion to swimming through a relatively small frame. The measurement head is also complex and has a rather small sampling volume. More recently, a system for estimating the mass of fish using stereo video cameras has become available (Petrell et al., 1997). An advantage of this approach is that the fish do not need to swim into an enclosed space for measuring. Its use is, however, labour intensive as an operator is required to collect suitable images of the fish and also, later, to identify four measurements points on each of two images of each fish. Since the images to be analysed are selected by the operators, care must be taken by the operator not to bias the sample. Further, since this is a manual system, the number of points to be identified on each fish has to be minimised to keep analysis time to an acceptable level. This places constraints on the choice of the morphological model, which is used to estimate mass from dimensions.

Two examples of stereo pairs of images of salmon in a sea cage are given in Fig. 1. The left-hand images were collected in relatively clear water with the sea surface illuminated by direct sunlight. The upper image is a slightly downward looking view from the upper camera of the stereo pair. The lower image was captured simultaneously with the upper by a horizontal looking camera about 0.5 m below the upper. The pair of images on the right was collected by the same equipment in slightly more turbid water when the water surface was in shadow.

Images captured in sea cages typically show low contrast between the fish and the background. The fish being examined may be either lighter or darker than the background and this may vary from point to point on a single fish. The surface features of the fish and the fins are also frequently obscure. Strong highlights may appear on the fish, due to specular reflection of the sunlight. Highlights are also

common in the background due to reflection of the sunlight off fish or debris. Where the water surface is visible, its brightness is likely to vary widely across the image due to refraction of the sunlight through the surface. The water surface curvature can focus the light creating diacaustics, which result in uneven and rapidly changing illumination. Due to the high densities of fish in a sea cage, it is also likely that some images of fish will overlap.

The interpretation of stereo images requires that corresponding points in each image of the pair are identified. Triangulation is then used to map a pair of matched images onto a real world 3D co-ordinate system. Because the geometrical relationship between the two cameras is fixed, a point in one image maps onto some point on a line in the second image. These lines, known as epipolar lines, are identified in the calibration of the cameras. Calibration of the stereo system taking account of the camera geometry and lens distortion is essential before 3D measurements can be made. This is achieved by identifying common points on a 3D grid of known geometry on each image of a stereo pair. Well-established methods for calibration, such as that presented by Tsai (1986), are available in the literature.

This article describes an approach to mass estimation, which avoids the limitations of existing methods by using stereo cameras, automatic image analysis, and improved morphometric analysis. The components of the system have been tested separately and, to a limited extent, together. Although all these components have been designed with the complex and variable nature of the sea cage environment in mind, most of the testing of the image analysis system has so far taken place under

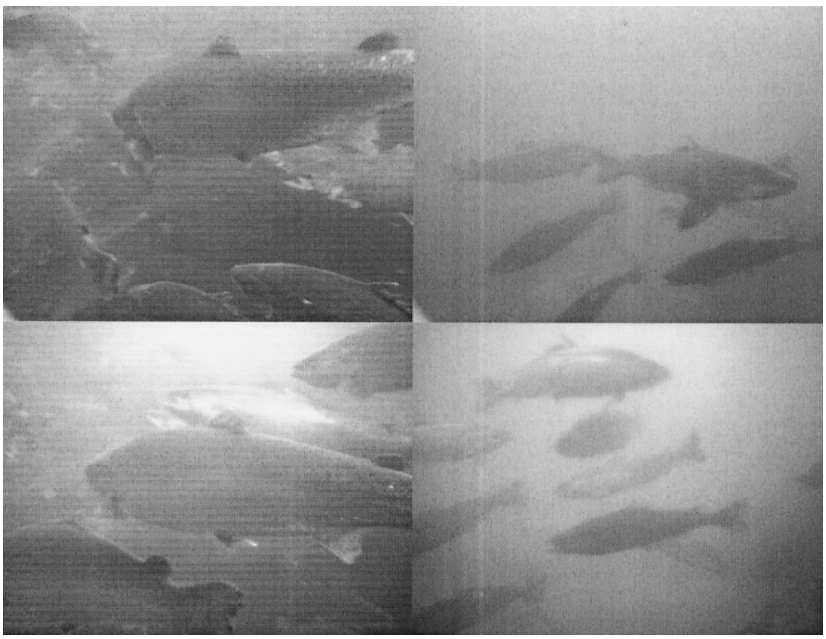

Fig. 1. Two pairs of stereo images of salmon in a sea cage.

more controlled tank conditions. This article provides an overview of the system in Section 2. Information on the separate components of the system, which has been published in a range of specialised publications, is here brought together and summarised to give the reader a general understanding of the functioning of the system. Section 3 presents some initial results from tests on the various components and on the performance of the system used as a unit.

2. System description

2.1. Overview

The hardware of the system comprises a stereo pair of waterproof video cameras, linked directly to a computer. Sequences of images are captured simultaneously from each camera and stored. Each frame is then automatically inspected to determine whether it is likely to contain any images of fish suitable for further analysis. Where suitable images are found, the approximate location and orientation of the candidate fish is identified. Model based analysis software then attempts to identify the boundary of the fish and locate it in 3D space. If a satisfactory fit is achieved, key measurement points on this boundary are identified and the distances between them calculated. A morphological model of the fish is then used to estimate its mass.

Since there is unlikely to be a shortage of images to analyse, the reliability of the analysis system is enhanced by using only images, which can be analysed with confidence. Marginal images are rejected since these might result in poor data, which could cause significant errors in the estimation of the mass distribution. There is a risk that such an error handling strategy might introduce bias into the sample, however the degree and nature of this bias are as yet unknown.

2.2. Image collection

Stereo pair video images are collected using a vertically oriented pair of video cameras with a base line of about 0.5 m and optical axes converging 1.5 m in front of the cameras. Vertical orientation allows the relatively well-defined top and bottom edges of the fish to be used to identify the distance of the fish from the cameras. Images of salmon are collected when they are between 1 and 2 m from the cameras. This represents a compromise between aiming to capture images from a zone close to the cameras which the fish might avoid, and capturing images from a greater distance where imaging might be prevented by poor water clarity and by fish between the camera and target. The chosen separation of the cameras represents a compromise between the requirement for precision in the depth estimates and the requirement that the two images are similar enough to enable corresponding points on the object to be identified in both images. Lens length was chosen to give a field of view width approximately twice the length of the fish. The system has been used with fish of length between 0.4 and 0.7 m. Fish significantly smaller or larger than this might require a stereo system of different dimensions.

The cameras which have been used are shuttered, raster scan, 1/2 in. monochrome CCD cameras with 4.5 mm lenses. In order to minimise the effects of fish motion only one scan from each image is used, giving an image resolution of 384×288. The two cameras use a single external sync to ensure that they collect simultaneous images. Since video recorders are unable to store pairs of simultaneous images, direct image collection by computer is used.

2.3. Identifying target fish

Identification of the fish outline requires a large amount of computer processing. It is, therefore, necessary to select suitable images and to estimate the location of the fish in the images prior to attempting to identify its outline. Simple image segmentation techniques such as thresholding and edge detection do not perform well due to the low levels of contrast typical of these images. The method developed for detecting objects that are likely to be fish uses a binary pattern classifier. This is outlined below and has been reported in more detail by Chan et al. (1998, 1999, 2000).

When an image of a swimming fish is subtracted from a second image slightly separated in time, a characteristic crescent shape consistently appears at the head of the fish. This is created by the change in image content from background to fish as the fish moves in front of the cameras. This pattern is robust enough to be enhanced by thresholding the processed image so reducing it to a simple binary representation which can then be recognised by a binary pattern classifier (Fig. 2). Recognition of the crescent appears to be relatively independent of the size of the fish image, since a small part of a large crescent and a large part of a smaller crescent can match equally well. The corresponding feature that is created at the tail of the fish is much more variable in form than the crescent and so is of less use for recognition.

Recognition of the crescent is achieved using an n-tuple binary pattern classifier. This established pattern recognition system is both simple and fast. The version used for this development is a simulation of the WISARD adaptive image classifier developed by Stonham (1986). The pixels in the image area to be examined are divided into groups of eight, randomly selected pixels, known as n-tuples. The pattern to be recognised is then learnt by examination of a small number of sample or training images. For each training image, the binary values of the pixels in each n-tuple define the value assigned to the n-tuple. For each image, each n-tuple can therefore take one of 256 possible values. The observed values for each n-tuple for the whole set of training images are calculated and stored. These values characterise the pattern.

During the recognition process the test image area is broken into the pre-determined n-tuples, and their values are calculated. These values are compared with those, which occurred in the set of training images. The proportion of n-tuples in the test image area which have a value which matches that of the corresponding n-tuple in any one of the set of training images, is used as a measure of the probability that the test image pattern belongs to the same class as the training image patterns.

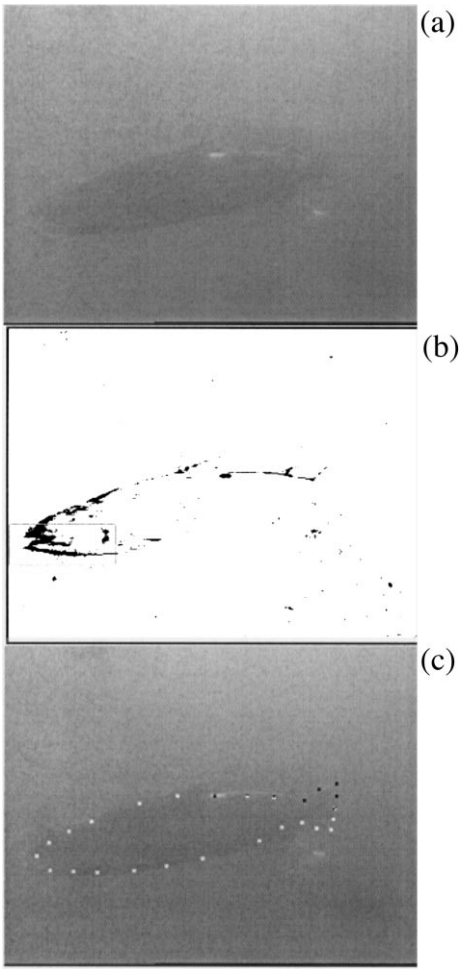

Fig. 2. Top camera image of salmon in turbid water in a sea cage. (a) Raw image from camera, (b) difference image with fish head automatically identified by pattern classifier algorithm, and (c) fish outline identified in 3D space using PDM.

This pattern recognition process is performed on all possible areas of one of the images of the stereo pair and the area with the highest probability score is identified. A search is then performed in the second image for a corresponding object. In this image, it is, however, only necessary to search for the highest scoring area close to the epipolar line associated with the highest scoring area of the first image. If the scores of these two areas are sufficiently high and if they are similar in value, then the 3D position of the area identified is calculated and considered to be a candidate position for a fish head. In practice, since fish are likely to be swimming either way, two searches will be needed using mirror images of the characteristic crescent. The position and direction of the candidate fish identified in this way is used to initiate the process of searching for the outline of the fish.

Fig. 2(a) shows an image of a salmon in a sea cage collected by the upper camera of a stereo pair. Fig. 2(b) shows a thresholded image created by subtracting the image in Fig. 2(a) from a similar image collected 0.16 s earlier. The fish head has been recognised by the pattern recognition algorithm and its position is identified by the window automatically marked on the image.

2.4. Locating the boundaries of fish

The edges of the fish in sea cage images can be difficult to identify due to the poor contrast of the fish with the background, the uneven illumination, and partial occlusion by other fish. A model-based approach to image segmentation is therefore used since this enables weak fish-shaped edges to be selected in preference to other, stronger, edges and allows areas where no edge is visible to be interpreted correctly. The method selected is the point distribution model (PDM). In this technique, the computer holds a shape template, which comprises the relative locations of a number of points on the fish shape together with the principal modes of variation of this shape. This shape is then fitted by working through a series of iterations where the strength and proximity of local edges are used to identify candidate fish edges which are then tested against the template. This work is described in more detail by McFarlane and Tillett (1997) and Tillett et al. (2000).

The shape template comprises 26 landmark points on the boundary of the fish. Some of the landmarks correspond to physical features of the fish, such as the tip of the tail or the junction of a fin with the body, while others are added to fill in the gaps along featureless parts of the shape. This template is built from data derived from a small set of training images. For each fish in this training set, the landmarks are placed by hand in both images of the stereo pair and the 3D positions of the landmarks are calculated. These fish shapes are then normalised to the same scale, rotation and translation, and the main modes of variation of the normalised shape are calculated using principal component analysis. Examination of these modes of variation suggests that the most significant mode is due to the swimming motion of the fish. The first four modes of variation are illustrated in Fig. 3.

The PDM is fitted to the candidate fish by minimising an energy-like function. This function comprises three components representing the energy required to deform and move the model to a particular position, shape and orientation, the image energy (edge strength) pulling the model to specific edge points and a residual energy due to the distance between the model and the edge points. Each component is configured as of a probability distribution, to avoid the need for empirical weightings of the individual energy terms.

The energy required to deform and move the model is calculated assuming that there is a Gaussian distribution of values of scale, translation, rotation and shape variation in the training set.

The energy due to the residual distance between the landmark points and the point on the candidate edge is calculated in a similar way but in this case, a Gaussian distribution of likelihood vs. distance is not suitable because of the

presence of outlier points. When incorrect candidate edges have been selected, they can occur at significant distances from the corresponding landmark points. To avoid a small number of such errors exerting too much influence on the overall fit of the PDM, a function corresponding to the Cauchy distribution is used since this does not increase as fast a Gaussian distribution when distance errors become large.

The image energy is calculated from the sum of the absolute values of the grey level gradient at the edges and the deviation of the angle of orientation of the edges from the angle of the edge of the model. The absolute grey levels and knowledge as to whether the fish are lighter or darker than the background are not used, because of the danger of learning features which are only applicable to specific lighting conditions and camera angles.

The PDM is fitted to the fish using a two-step iteration. In the first step, the PDM is held fixed, while a set of candidate edges is identified for each landmark point. In each image, a search is performed along the normal to the landmarks, identifying edges. At this point, the edge candidates for each landmark consist of a list of 2D image co-ordinates for each of the images in the stereo pair. These are combined into a list of 3D positions. The most likely edge to be associated with each point is then identified as that with the minimum sum of image energy and residual distance energy. In the second step, these candidate edges are held fixed while the PDM is fitted to them by minimising the sum of the model energy and the residual energy. These two steps are iterated several times to allow the model to converge to a stable position. Fig. 2(c) shows the landmark points identifying the edge of the fish after a series of such iterations. These have been projected onto the upper camera image.

2.5. Estimating fish mass

A morphological model was built to enable measurements that could be extracted from the stereo images to be used to estimate the mass of as wide a population of

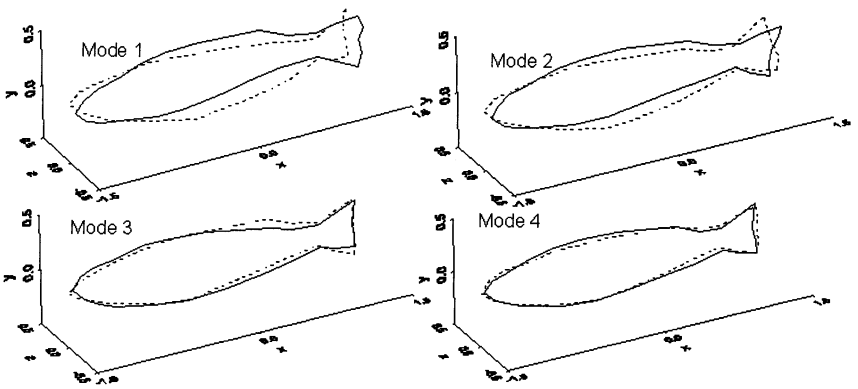

Fig. 3. The first four modes of variation of the point distribution model showing ± 3 S.D. about the mean shape.

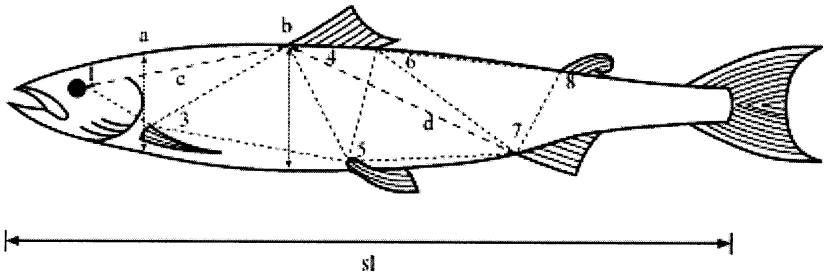

Fig. 4. Network of truss measurements made on the salmon.

Atlantic Salmon as possible. The model has followed Humphries et al. (1981) in identifying the lengths of a series of trusses linking discernible anatomic features on the fish and covering the whole body. Beddow and Ross (1996) developed a series of linear regression equations that used morphometric truss dimensions to predict the mean mass of individual Atlantic Salmon to within $\pm 2\%$. When these equations were applied to data taken from stereo video images, overall group biomass was predicted to within 0.4% of the real value and individual weight estimates were determined with a mean error of $-0.1 \pm 9.0\%$ (Beddow et al., 1996). However, in comparing Scottish and Norwegian strains of salmon, morphological differences were found that affected this predictive accuracy (Beddow and Ross, 1996).

A database of measurements was built to enable new mass models to be identified which could handle variations in genetic strain with more reliability (Hockaday et al., 2000). Just over 1300 Atlantic Salmon were measured to collect this data. The fish had a mass of between 30 and 6150 g and were grown on one of several fish farms on the West Coast of Scotland (Table 1). Fish were anaesthetised, weighed, and the distances between landmark positions (truss lengths) were measured. The truss lengths measured were chosen because they had been shown by Beddow et al. (1996) to be useful in predicting the biomass of Atlantic Salmon from stereo-image pairs. They are shown in Fig. 4 as the dimensions *a*, *b*, *c*, *d*, sl, and the marked distances between points 1, 3, 4, 5, 6, 7 and 8.

Model identification was approached by seeking combinations of these lengths that showed a near-linear relationship with mass. The relationships were then optimised using regression analysis. In choosing which variables to use in a model, preference was given to lengths, which could be measured with low measurement errors, and those that contained different information about the shape of the fish.

Two simple mathematical relationships were used for building models:

$$\text{Mass} = p + \sum_i p_i X_i^{q_i}, \tag{1}$$

$$\text{Mass} = p + p_0 \prod_i X_i^{p_i}, \tag{2}$$

Table 1
Source location and genetic strain of Atlantic Salmon used for building morphological models

Strain (and origin)	Location							
	Loch Eil	Loch Duich	Loch Sunart	Firth of Lorne	Machrihanish	Otterferry	Kilberry	Loch Fyne
Homebred (East coast of Scotland)	108	165	21					
Lochy (West coast of Scotland)	188	107						
Otterferry (Otterferry fish farm)				8	99	31		
Mowi (Norwegian)		133	9				16	
Namsen (Namsen river in Norway)		120	12					
Misc. Norwegian strains	212							
Various (mixture of several strains)								30
80:20 Norwegian : Scottish crosses							11	
Tasmanian/Canadian strains		49						

where p, p_0, p_i and q_i are constants, and X_i are lengths measurements on the fish.

After fish data with missing values and outliers were removed from the data set, there remained 1002 individuals. These were split randomly into a training data set (70%) and a test data set (30%). Many models were proposed and optimised using regression. Among the best of these models are the first seven shown in Table 2.

Table 2
Errors in mean mass of 304 salmon estimated using various morphological models[a]

Model	Model equation	Correlation coefficient	Error in mass estimate
R0	$\text{mass} = n_1 + n_2 \cdot b^{n_3} \cdot \text{sl}^{n_4}$	0.996	−0.3%
R2	$\text{mass} = n_5 \cdot b^{n_6} \cdot c^{n_7} \cdot d^{n_8}$	0.997	0.6%
R4	$\text{mass} = n_9 \cdot b^{n_{10}} \cdot \text{sl}^{n_{11}}$	0.997	2.0%
R5	$\text{mass} = n_{12} \cdot a^{n_{13}} \cdot \text{sl}^{n_{14}}$	0.994	−0.4%
R6	$\text{mass} = n_{15} + n_{16} \cdot b^3 + n_{17} \cdot c^3 + n_{18} \cdot d^3$	0.997	−0.2%
R10	$\text{mass} = n_{19} + n_{20} \cdot b^2 \cdot \text{sl}$	0.995	−0.2%
R11	$\text{mass} = n_{21} + n_{22} \cdot a^2 \cdot \text{sl}$	0.991	0.4%
B5	$\text{mass} = n_{23} + n_{24} \cdot b^3 + n_{25} \cdot c^3 + n_{26} \cdot d^3$	0.995	−0.3%
B6	$\text{mass} = n_{27} + n_{28} \cdot b^3 + n_{29} \cdot \text{sl}^3$	0.995	−0.2%

[a] The directly measured mean mass was 1472 g. In the model equations, $n_1,\ldots,n_{29}$ are fitted constants and a, b, c, d and sl are lengths on the salmon as shown in Fig. 4.

Each predicted the mass of the test set to within $\pm 2\%$ of the weighed mass, with five of the models predicting to within $\pm 0.5\%$ and two models to within $\pm 0.2\%$. The individual mass predictions using models R6, R10 and R11 were poorest for fish smaller than 400 g.

The data set was also split into groups, comprising fish weighing more than, and less than, 500 g. Two more sets of equations based on Eqs. (1) and (2) were then produced. These enabled the masses of both test sets to be predicted to better than $\pm 1.4\%$. Examples of models developed for fish weighing more than 500 g include B5 and B6 shown in Table 2. A more detailed description of the morphological models, including the results of testing models using single strains of salmon is given by Hockaday et al. (1997, 2000).

3. Results of preliminary trials

Some testing of the system components has taken place, using images, which were captured, in a tank stocked with salmon of known mass and dimensions. These salmon were marked to enable them to be individually identified from the video images. The tank was 3.5 m in diameter with a water depth of 0.9 m. A black back-board was placed 1.8 m from the cameras and a similar base-board was placed on the floor of the tank. These boards provided a more uniform background to the images than the fibreglass tank wall, and also encouraged the fish to position themselves in front of the cameras. Shading cloth covering the tank reduced and diffused the natural illumination. The entire sequence of tasks from collection of the images to prediction of fish mass has been followed through using this facility. These images are simpler to analyse than many images collected in sea cages because of the uniform lighting, the background and the high water clarity.

Some images were also collected on the west coast of Scotland in commercial sea cages at depths of between 0.5 and 5 m. These were of salmon of unknown mass. These images have enabled the functioning of the image analysis routines to be observed under more challenging conditions. Insufficient experience has however been gained to allow generalisations to be drawn regarding the likely characteristics of sea cage images or to allow the image analysis algorithms to be optimised for this environment.

3.1. Images

A typical stereo pair of images collected in the tank is given in Fig. 5. The high water clarity and uniform background results in relatively clear images. However, contrast between fish and background is still relatively low so shading across the body of the fish can sometimes result in stronger edges than occur at the edge of the fish.

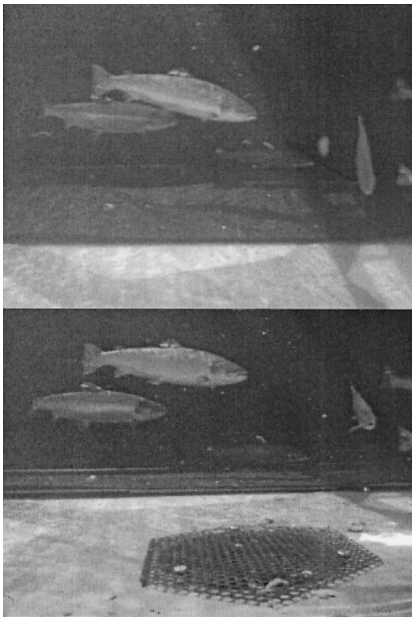

Fig. 5. Example of stereo image pair collected in the tank.

3.2. Identifying the locations of candidate fish

Fourteen sequences of stereo image captured in the tank were used to assess how the pattern classifier might identify the location of fish. Each sequence comprised 25 image pairs representing two seconds of time. The n-tuple pattern classifier system was trained using nine training images, which had been collected, earlier in sea cages (Chan et al., 1999).

Initial observations suggested that the binary image representing the difference of two images should be created using images separated by 160 ms and that the threshold for detecting changes in the images with a total range of 256 grey levels should be set at 15 grey levels. Examination of the scores generated by the pattern classifier suggested a fish head was likely to result in a score of greater than 0.6 in both images of the stereo pair. In the 322 processed image pairs, 86 pairs of images were found where a fish head was identified by the algorithm. Of these images, 55 were correctly classified and 31 were wrongly classified (i.e. the feature identified in the images was not a fish head). In order to reduce the number of wrong classifications, a second requirement was introduced that the score should be similar for each member of the pair. When the maximum permitted difference was 0.15, the number of correct classifications was reduced to 51 and the number of incorrect classifications to 18. By reducing the permitted difference in score to 0.05, the number of correct classifications was further reduced to 47 and the number of incorrect classifications was reduced to only two.

It would also be of interest to identify the number of strong images of fish heads, which were not identified by the pattern classifier. This, however, requires independent criteria for judging the strength of the images. However an indication of this failure rate can be obtained by observing that all the image sequences were selected because at some point, they contained images of well-presented fish. In 4 out of 14 image sequences, no fish heads were identified by the pattern classifier with an acceptance criteria of 0.6. When the pair similarity test was performed (criteria, 0.15 or 0.05), a further two image sequences were rejected as containing no good images. In all, therefore, 8 out of 14 sequences were identified as containing suitable images and six were rejected.

3.3. Fitting the model

The point distribution model was applied to a set of tank images of tagged fish for which manual size and weight measurements had also been taken. The data set comprised 60 images of 17 fish ranging in mass from 0.7 to 5.7 kg. Twenty-six of the images were used to train the model. The model was then applied to all 60 images. A model fit, which appeared to be correct could be achieved from 42 out of 60 images, when the starting position of the model was selected by hand. The model failed to converge onto the fish in 5 out of 26 training images and 13 out of 34 test images.

The PDM was also applied to images using the position of the fish head identified by the pattern classifier. Eleven images were used where the pattern classifier had correctly identified a fish head. In 7 out of 11 image pairs, the model converged successfully onto the outline of the fish.

Some examples of the landmark points of the PDM fitted onto images are shown in Fig. 6. Images (a) and (b) show well-fitting models, (c) shows a model, which has fitted to an edge slightly away from the lower self occluding edge of the fish due to strong local shading patterns, and (d) shows a fit which has failed due to the extreme position of the fish.

It is generally more difficult to fit the point distribution model to sea cage images than to the images collected in the tank. In order to prevent the model attaching itself to spurious edges, it can be necessary to reduce the distance over which it searches for edges. More accuracy in the initial placement of the model is therefore required under this circumstance. Indistinct edges due to low water clarity and strongly directional or uneven illumination can result in poor fitting of the model.

3.4. Estimating the dimensions of free-swimming fish

Three sets of linear measurements were made on each fish in the tank. The first set is the truss measurements made using callipers on the anaesthetised fish. These we regarded as reference measurements. The second set of measurements was collected from the images by manually positioning a cursor on the landmark points and at end of each truss and using the stereo calibration to calculate distances. The third set was calculated from points, which were automatically located on the

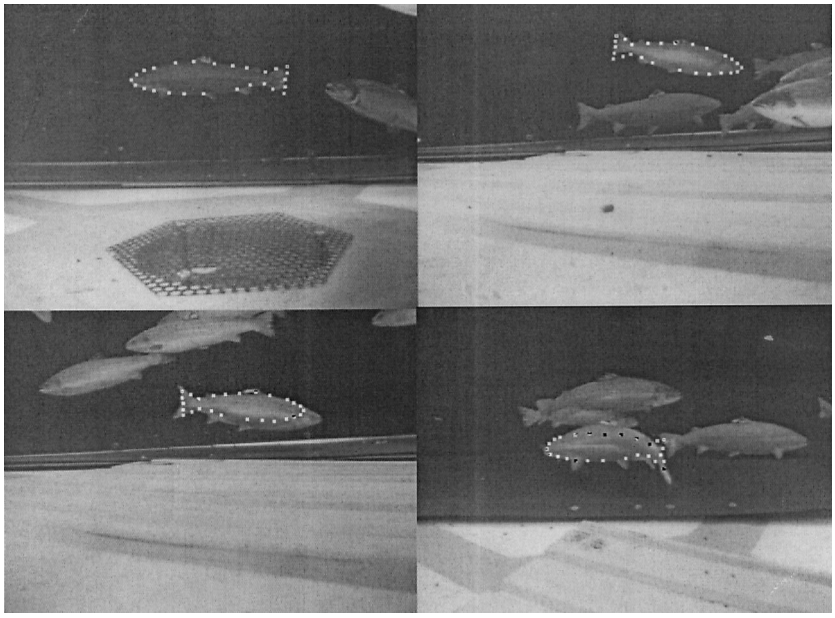

Fig. 6. Examples of PDM fitted to stereo tank images illustrating good fits (top), a poor fit due to shading on the fish (lower left) and poor fit due to extreme fish posture (lower right).

boundary of the fish images by the PDM. These data, therefore, provide an opportunity to separately test the morphological model, the accuracy with which distances can be recovered from stereo images, and the accuracy of data derived from the automatically identified fish outline. In some cases, positions on the fish may be required which are not supplied by the PDM. An example of this is the position of the eye. This was predicted from the expected relationship of the eye location with the locations of landmark points identified by the PDM.

Lengths generated from the hand picked points gave reasonable estimates of the truss lengths with mean errors of between -2% and 8% with standard deviations of up to 5%. The points fitted automatically were slightly less accurate with errors ranging from -8% to 1% and a standard deviation of up to 9%. (Table 3).

Table 3
Means and standard deviations of the errors in measuring four key truss lengths from stereo images using hand fitted points and using the point distribution model

Length	*b*	*c*	*d*	sl
Hand positioned points	$-2\% \pm 3\%$	$8\% \pm 4\%$	$0.3\% \pm 5\%$	$-2\% \pm 1\%$
Point distribution model	$-8\% \pm 9\%$	$1\% \pm 8\%$	$-6\% \pm 9\%$	$-8\% \pm 7\%$

Table 4
Mean and standard deviation of errors in estimated mass of fish using three mass models and using truss lengths derived in three ways[a]

	Calliper measurements	Hand fitted points	Point distribution model
R2 model	$-2\% \pm 7\%$	$1\% \pm 8\%$	$-8\% \pm 10\%$
B5 model	$-1\% \pm 7\%$	$5\% \pm 8\%$	$-6\% \pm 11\%$
B6 model	$-2\% \pm 5\%$	$-9\% \pm 7\%$	$-18\% \pm 9\%$

[a] Model R2 is outlined in Table 2. Models B5 and B6 were among the models created based on data derived from fish with mass greater than 500 g.

The data available is drawn from 60 images of 17 fish so these results give only an indication of the potential of this technique. However, it appears that the model-based approach will position the landmarks reliably so that, the key dimensions on the fish can be estimated. A specific calibration may be required to allow for bias in how the model fits to the fish images.

3.5. Mass estimations

The truss length data generated in the various ways were introduced into three of the morphological models to obtain estimates of mass (Table 4). The three models chosen were all suitable for the weight range of fish examined. The estimated mass of the individual fish based on calliper measurements had a mean error of 2% or less with a standard deviation of 7%. Errors in the estimated masses based on the automatically interpreted image data were substantially greater than this, however, the standard deviation of the errors remained relatively small and constant. This suggests that there are some systematic errors in the processing of the image data and the fitting of the model, which have yet to be identified. The relationship between the estimated masses and directly measured fish masses are shown in Fig. 7.

4. Discussion

Although the basic components of an automatic mass estimation system have been developed, enhancements are likely to be required to achieve adequate accuracy and reliability for use in commercial sea cages. Suitable techniques for achieving this objective can only be identified following a fuller assessment of the range of image characteristics, which are likely to be encountered under a normal range of operational conditions. These techniques are likely to include methods to detect when the point distribution model fails to fit. More robust fitting of the model might possibly be achieved through the use of some internal landmark points, modelling of the shading patterns across the fish and improved edge detection techniques using oriented filters. A substantial collection of new sea cage images under a range of realistic conditions is required before a program of

optimisation and testing can proceed. This process will also help to identify conditions under which the best images are likely to be collected.

The quality of the images is also dependent on the proximity of the fish. Both too many and too few fish can lead to a shortage of good images. It is, therefore, of some importance to learn how the cameras should be positioned with respect to the school of fish in the net. Further, if a mass estimation is to be based on these images it is necessary to ensure that any bias in the masses of the fish which present themselves to the cameras or in the images which are automatically interpreted, is recognised and quantified.

Some optical techniques for improving the quality of the collected images could also be investigated. This might include the selective use of illumination or filtering. However, since illumination would make the system more intrusive and possibly of create a bias in the sampling, its use should be avoided, if possible.

Although all the work reported in this article has concerned Atlantic Salmon, similar techniques could be used on a range of other fish. Morphological models for the deep-bodied type fish Tilapia and the flat fish Halibut have been investigated. Results suggest that the mass of both these fish types could be successfully estimated using a similar morphological model (Ross et al., 1998).

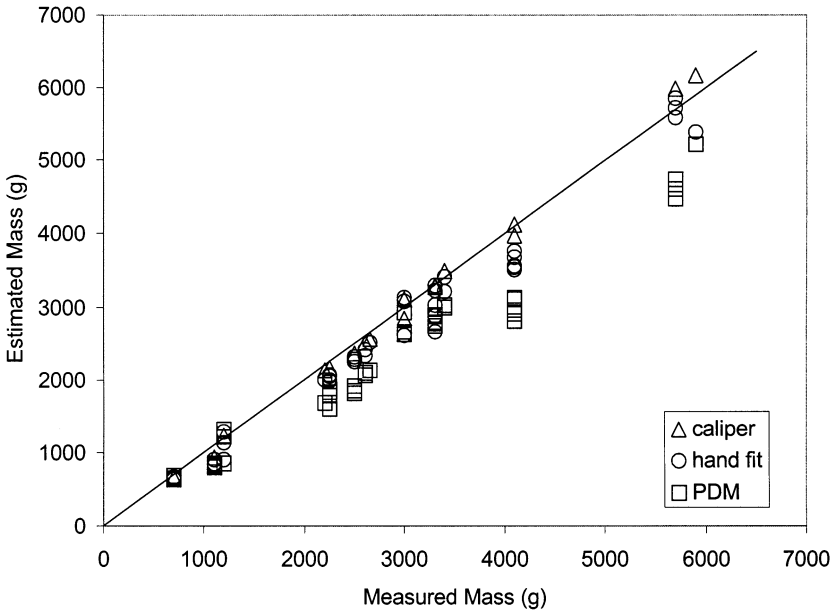

Fig. 7. Comparison of fish mass with estimated mass based on B6 model using calliper measurements (one set per fish), landmark points manually identified from images (several images of each fish) and automatically fitted point distribution model (several images of each fish).

5. Conclusion

The basic components of an automatic image-based mass estimation system have been created. These components enable images likely to contain a useable image of a salmon to be identified; enable the self-occluding boundary of the salmon to be identified in 3D space; and enable the mass of the salmon to be estimated from information about this boundary.

These components have been tested on salmon in a tank and have been shown to produce estimates of fish mass, which are well-correlated with the measured mass. The image analysis components have also been shown to be capable of working on some of the complex images typical of a sea cage.

Further improvements in the accuracy, robustness and error detection will be sought, using a combination of generic and ad-hoc techniques. These are required before an adequate level of system reliability can be demonstrated. This development and testing is dependant on the collection of an expanded database of sea cage images. Optimal mass models will be identified from those, which have been developed. The choice will be in part determined by the accuracy with which the various truss lengths can be estimated from the images.

Acknowledgements

This work was funded by the Biotechnology and Biological Sciences Research Council, UK. Field experiments were carried out with the assistance of Marine Harvest McConnell.

References

Beddow, T.A., Ross, L.G., 1996. Predicting biomass of Atlantic salmon from morphometric lateral measurements. Journal of Fish Biology 49, 469–482.

Beddow, T.A., Ross, L.G., Marchant, J.A., 1996. Predicting salmon biomass remotely using a digital technique. Aquaculture 146, 189–203.

Chan, D., Hockaday, S., Tillett, R.D., Ross, L.G., 1990. A trainable n-tuple pattern classifier and its application for monitoring fish underwater. In: Proceedings of the IEE conference on Image Processing and its Applications, 13–15 July, 1999, Conference publication no. 465, IEE, London. ISBN 0 85296 717 9, pp. 255–259.

Chan, D., Hockaday, S., Tillett, R.D., Ross, L.G., 2000. Automatic initiation of a model fitting algorithm by using an n-tuple classifier for monitoring fish underwater. In: Fourth Asian Conference on Computer vision (ACCV 2000) Taipei, Taiwan, 8–11 January, 2000, IEEE (Asia).

Chan, D., McFarlane, N., Hockaday, S., Tillett, R.D., Ross, L.G., 1998.Image processing for underwater measurement of salmon biomass. In: Proceedings of the IEE Colloquium on Image Processing in Underwater Applications, 25 March, 1998, no. 217, IEE Colloquium (Digest), London, pp. 12/1–12/6.

Hockaday, S., Ross, L.G., Beddow, T.A., 2000. A comparison of models built to estimate the mass of different strains of Atlantic salmon, Salmo salar L, using morphometric techniques. Aquaculture (in press).

Hockaday, S., Ross, L.G., Tillett, R.D., 1997. Using stereo image pairs to measure mass in strains of Atlantic salmon (Salmo salar L). In: Paper presented at "Sensors and their applications VIII", Section A Environmental and Biomedical Sensors, 7–10 September, 1997, Institute of Physics publishing, Bristol, UK, pp. 21–26.

Humphries, J.M., Bookstein, F.L., Chernoff, B., Smith, G.R., Elder, R.L., Poss, S.G., 1981. Multivariate discrimination by shape in relation to size. Systematic Zoology 30 (3), 291–308.

McFarlane, N., Tillett, R.D., 1997. Fitting 3D point distribution models of fish to stereo images. In: Proceedings of the British Machine Vision Conference, BMVC 1997, University of Essex, UK, 8–11 September, 1997, vol. 1, BMVA Publishing, Malvern, UK. ISBN 0952 189887, pp. 330–339.

Petrell, R.J., Shi, X., Ward, R.K., Naiberg, A., Savage, C.R., 1997. Determining fish size and swimming speed in cages and tanks using simple video techniques. Aquacultural Engineering 16, 63–84.

Ross, L.G., Hockaday, S., Tillett, R.D., Chan, D., 1998. Remote weighing of fish: myth and reality. In: Proceedings of the British Trout Farming Conference 1998. Published by Sparsholt College/CEFAS, Winchester, Hampshire, England.

Stonham, I.J., 1986. Practical face recognition and verification with WISARD. Aspects of face processing. Martinus Nighoff, Dordrecht.

Tillett, R.D., McFarlane, N., Lines, J.A., 2000. Estimating dimensions of free-swimming fish using 3D point distribution models. Computer Vision and Image Understanding 79, 123–141.

Tsai, R.Y., 1986. An efficient and accurate camera calibration technique for 3D vision. In: Proceedings of the IEEE Conference on computer vision and pattern recognition, Miami Beach, USA, 22–26 June, 1986. IEEE, New York.