Visual Learning of Weight from Shape Using Support Vector Machines

*F. Odone**, E. Trucco*, A.Verri***

* Department of Computing and Electrical Engineering
Heriot-Watt University, Edinburgh, UK
`mtc@cee.hw.ac.uk`

** Dipartimento di Informatica e Scienze dell'Informazione
Università di Genova, Genova, Italia
`{odone,verri}@disi.unige.it`

Abstract

We investigate the automatic estimation of fish weight from sets of morphometric measurements. Our solution combines a vision system with a robust regression method, the Support Vector Machine (SVM). Measurements are taken automatically from two binarised views of each fish in a training sample, then fed to a quadratic SVM along with approximate weight estimates. The SVM learns the law linking weight to shape *directly* (without computing volume) and compensates for several inaccuracies in the training measurements. We suggest a methodology identifying optimal shape measurements for the task, and report results obtained with a sample of 99 trouts between 300 and 600g, showing good accuracy and reliability, and better performance with respect to length-weight relations adopted commonly in fisheries science.

1 Introduction

This work explores a new way of estimating fish weight from shape using computer vision. The relation between weight and shape is important both for fish biology [4, 5, 7, 11] and fisheries applications. Weight measures from shape are approximate as indirect (density may not be uniform across individuals), but no technology currently allows to weight directly live, moving fish. The stress induced by constraining or anaesthetizing is strongly discouraged for most species, and may lead to serious damage for others (e.g., sea bass).

Many studies in fish biology and farming are based on equations predicting weight from a single shape measure, traditionally length [5, 13], or width in most mechanical fish graders. In particular, it is common practice to model weight as a function of length through the allometric equation

$$W = KL^b, \tag{1}$$

introduced by Fulton [5], where W is the weight, L the length, K is the *condition factor* (sometimes used as an index of fish health), and b a coefficient close to 3. The inadequacy of formulae like (1) for weight prediction was pointed out by Fulton himself [5], and, more recently, by various researchers [1, 2, 13]. The main inadequacies include the facts that weight depends really on several shape parameters, that the exponential is in itself an approximation (especially if isometric growth is assumed, i.e., $b \approx 3$), and that such laws assume invariance to time, populations and grading conditions, or make hypothesis on the fish (e.g., sex, age) or on the environment (e.g., temperature, season). The practical advantage of single-measurement relations are that they are simple, and can be implemented easily in mechanical devices (e.g., by graded holes) to support fast fish grading.

Vision systems can take *multiple* measurements in real time, as fish swims in front of imaging sensors. Therefore, they are well-placed to provide better accuracy of weight estimates than mechanical systems, at comparable speed. Indeed experimental and commercial vision systems exist for biomass estimation, counting, species classification, grading on shape, weight and colour. The crux of existing weight-based graders is that weight is computed through volume, which in turn is estimated using approximations and fixed formulae which are assumed invariant against time, populations and grading conditions.

Most of these assumptions would be unnecessary if the system *learned* the weight-shape relation for each fish batch considered, which is the scenario investigated by our study. An automated learning algorithm, the Support Vector Machine (SVM), learns the specific relation between weight and a vector of morphometric measurements for *each* fish batch under consideration. Measurements are extracted automatically from two binarised views of each fish. The basic architecture is illustrated in Figure 1. Instead of using fixed, approximate formulae to estimate fish weight from one shape parameter, the system tunes its estimates to the *specific* batch, with a resulting increase in accuracy given a sufficiently large training set. The law linking shape and weight is found *directly*, without going through volume at all. There is no need for assumptions on the fish density, notoriously a variable quantity, nor to assume density constant across batches, as a specific law is learnt for each batch.

This paper is organised as follows. Section 2 reviews briefly the main features of SVMs. Section 3 describes the prototype vision system acquiring morphometric measurements. Section 4 illustrates the learning process, and Section 5 reports the results of our experimental assessment of the system with a batch of 99 trouts. Finally, Section 6 discusses our findings and draws some conclusions.

2 Support Vector Machines

A detailed explanation of SVMs would take us far beyond the scope of this paper, so we limit ourselves to a brief summary with references. SVMs implement a learning technique useful for solving both classification and regression problems [3, 9, 10, 18]. The main idea behind SVMs is the minimisation of *structural risk* as opposed to *empirical risk*. Intuitively, given a finite number of sample points and a class C of approximating functions, *empirical risk minimisation* determines the function $\hat{f} \in C$ which best approximates f by minimising a certain cost function

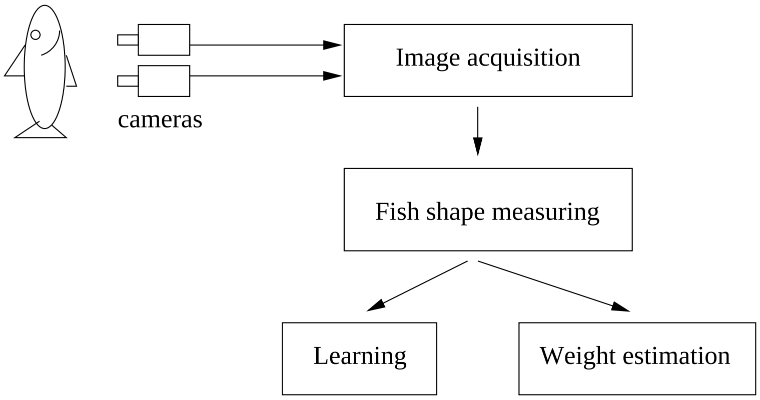

Figure 1: Essential architecture of the prototype.

(e.g., in the least squares case, the sum of the squared distances of the sample points from $\hat{f}$). Apart from robustness considerations, empirical risk minimisation is a sound minimisation principle only in the presence of a large number of sample points (here, *large* means that all the theoretical estimations are valid in the asymptotic case). *Structural risk minimisation*, instead, aims at minimising an upper bound on the approximation error and works also in the case of small number of samples. This bound consists of a term which accounts for how well the function $\hat{f}$ approximates f plus a term depending monotonically on the *VC-dimension* of the class C of approximating functions (roughly a measure of the approximating power of the given function class [17]). Therefore, SVMs trade accuracy of the approximation at the sample points for smoothness of the approximating function $\hat{f}$. As a result, SVMs are able to determine the optimal approximating function, which depend on, say, m parameters, even in the presence of a number of sample points far smaller than m.

3 Acquiring shape measurements

Although our case study focused on harvested fish, we took into account the main requirements of real fish grading environments: high speed, reliable measurements, low cost. To achieve fast processing and reliable shape measurements, we worked with silhouettes acquired by illuminating the fish from the back, and thresholding the raw images. Measurements were computed from two views: the fish was rested on the translucent top of a box containing neon tubes (Figure 2), first flat (side view), then vertically using a transparent support (top view). All fish appear approximately horizontal[1].

From the side view, the system computed area, perimeter, length, the ratios area/length and area/perimeter, and the minimum and maximum widths. The

[1] This simulates a design suitable for live fish moving through a mechanical grader, which is currently covered by confidential agreements.

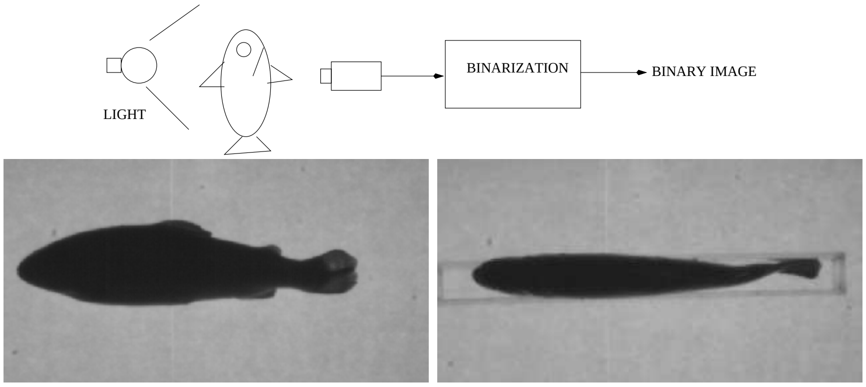

Figure 2: Schematic of the acquisition system and examples of top and side view before binarisation. Notice the image of the transparent support used to stand the fish under the camera (top view), which is eliminated by binarisation.

same shape parameters were computed from the top view, excluding the minimum width, which proved unstable as too dependent on the pose of the fish. Area is given simply by the number of black pixels, and perimeter obtained by edge detection followed by chaining and edge following. The ratios area/perimeter and area/length indicate elongation. Length and width are calculated following [15]: we obtain a first, rough approximation of length by calculating the number of image columns intersecting the fish silhouette. Then, we use a subset of these columns (e.g., one every 10; the frequency depends on the accuracy desired) to compute an approximate skeleton by joining the midpoints of the column portions falling within the fish silhouette. We then compute a piecewise linear approximation of the resulting curve. The fish width at each linear segment is computed by tracing the normal to each segment; the length is the length of a linear piecewise skeleton along the normals to the width segments (Figure 3). The minimum width (side view only) corresponds to the thickness of the fish in the narrowest part between body and tail.

4 Learning the shape-weight relation

Each fish is associated with a vector $\mathbf{p}$ of n measures ($n \leq 13$ in our experiments) taken from the two views. By showing a number of training points $(\mathbf{p}_j, w_j)$, where w_j is a weight estimate for the j-th fish, the SVM determines the best approximating hypersurface $w = f(\mathbf{p})$, that is, a relation linking weight and shape parameters. It is important to notice that, in the case of live fish, the weight estimates w_j used in the training stages may include rather significant errors and even outliers, but these are compensated for by the robust approximation performed by SVM in the training stage.

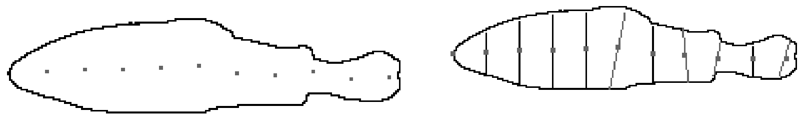

Figure 3: Illustration of length/width estimation: (a) initial skeleton points (column midpoints), (b) normal sections and final skeleton points.

5 Using the shape-weight relation: results

We tested the system with a sample of 99 harvested trouts between 300 and 600g approximately[2] (see Figure 4). Weights and lengths used for training were

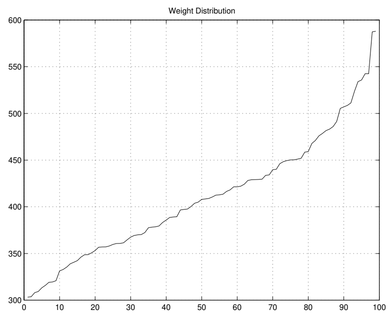

Figure 4: Weights of the trout sample, sorted by weight and plotted against fish index.

measured manually. The weights were read from a set of electric scales for shop use, nominally accurate to 5g in our weight range. Lengths were measured using a standard graduated board, with an uncertainty of about 1cm (as per EU fisheries guidelines).

Testing equation 1. We checked the consistence of our sample with equation (1), with coefficients computed by linear regression (based on $\log_e(W) = \log_e(K) + b\log_e(L)$) [4, 5, 7, 11]. Using the average of the lengths in the top and side views for L, we got $b = 2.177$, $\alpha = \log_e K = -1.605$, with standard deviations $\sigma_\alpha = 0.52$ and $\sigma_b = 0.36$. The results, summarised in Figure 5 and Table 1, suggest that the hypothesized relation between length and weight is unsatisfactory [11, 4]. Figure 5 also suggests that the assumption of isometric growth does not hold satisfactorily for this sample; in fact, the estimated b from the weight-length regression was 2.177, not 3; Fulton's K decreases with increasing length instead of being almost constant [11].

[2] We report only one real experiment as organising tests with large numbers of fish is a lengthy process, and this study spanned a few months only.

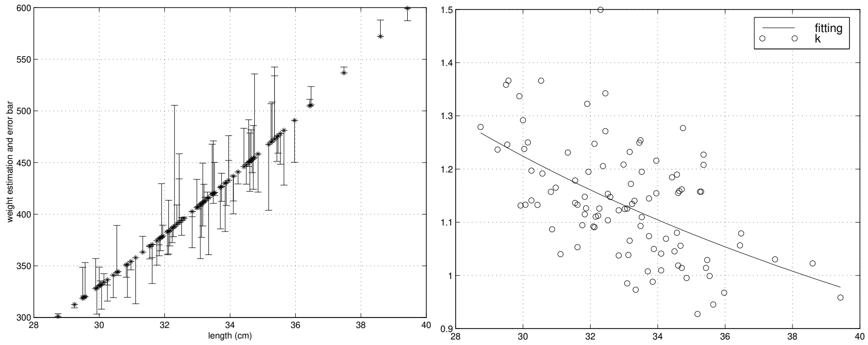

Figure 5: Left: result of fitting equation 1 to our data ($\alpha = \log_e K = -1.605$, $b = 2.177$). Right: regression of Fulton's condition factor. Our sample violates the assumption of isometric growth ($b \approx 3$); in fact, K decreases with increasing length instead of being almost constant.

Testing SVM regression. We then trained a SVM with quadratic kernel using different subsets of the 13 shape measurements available and 98 trouts at a time, each time estimating the weight $\hat{W}$ of the trout left out. The quadratic kernel is simple, has a limited number of coefficients, and seems appropriate for the limited weight-length ranges envisaged for batch grading on fish farms. We recorded the average and standard deviation of the percentage error, $(\hat{W} - W_t)/W_t$, where W_t are the weights measured by hand. The whole test was run for several different subsets of shape parameters, containing increasing numbers of features. The results are summarised in Table 2. We notice immediately a better accuracy than equation 1 (0% average error and 3% standard deviation without outliers, less than half of the 7% standard deviation from equation 1), with and without outliers.

	err	*st.dev*	*outlrs number*	*err (0 outlrs)*	*st.dev.(0 outlrs)*
fixed b	-11 %	8 %	7	-10 %	7 %
regression	0 %	8 %	2	1 %	7 %

Table 1: Error analysis for weight estimation. First row: assuming isometric growth ($b = 3$, fixed, and $log_e K = 4.212$). Second row: K, b obtained by linear regression.

Optimal sets of shape parameters. It is evident from Table 2 that not all feature sets are equally reliable for weight estimation, but produce varying accuracies on the same data. Indeed, Table 2 allows one to identify the *optimal feature sets* (as well as the most important individual features) for weight estimation, that is, the sets leading to minimum average error and error spreads. The practical significance is apparent: given a fish batch, one can identify the sets of shape measurements leading to minimum errors for practical tasks like sorting, grading and so on. Notice that some shape parameters worsen the weight estimates because they cannot always be measured reliably: for instance, the values of area and perimeter would be sometimes skewed by half-transparent fins, and the minimum-width section could not always be located correctly in the thin top view. Figure 6 visualises the errors of the experiments of the first four lines of Table 2, compared with the errors of equation 1. Using only the lengths (features 4 and 11) measured in the two views, the errors are comparable, as expected; adding two unreliable shape parameters (perimeters: 2, 9) the SVM errors increase; adding reliable parameters (ratios area/length and maximum width in the top view: 5, 7, 12) the SVM errors decrease.

feat.s in set	*feat.s used*	*err*	*st.dev.*	*err (0 outls)*	*st.dev. (0 outls)*
2	4 11	1 %	8 %	1 %	7 %
4	2 4 9 11	1 %	8 %	0 %	7 %
4	4 5 11 12	0 %	6 %	0 %	3 %
5	4 5 7 11 12	0 %	5 %	0 %	3 %
6	4 5 7 11 12 13	0 %	4 %	0 %	3 %
7	3 4 5 7 11 12 13	0 %	5 %	0 %	4 %
7	4 5 6 7 10 11 13	0 %	4 %	0 %	4 %
7	1 4 5 7 11 12 13	0 %	4 %	0 %	3 %
8	3 4 5 7 10 11 12 13	0%	6 %	0 %	4 %
8	2 3 4 5 9 10 11 12	-1%	7 %	0 %	4 %
13	all features	0 %	9 %	0 %	7 %

Table 2: Error analysis with SVM estimation using a quadratic kernel and different feature sets. One outlier was present in all experiments. Side view: 1=area, 2=perimeter, 3=area/perimeter, 4=length, 5=area/length, 6=minimum width, 7=maximum width. Top view: 8=area, 9=perimeter, 10=area/perimeter, 11=length, 12=area/length, 13=maximum width.

6 Discussion

We have addressed the estimation of fish weight from shape measurement using a vision system incorporating SVM-based learning. Equation (1), commonly used in fisheries applications, makes weight a function of length only, and depends on several restricting assumptions (fish density constant and invariant to fish batch, temperature, season, time, grading conditions, and sometimes isometric growth). These facts limit the accuracy of weight predictions, both with fixed coefficients K, b and with K, b specialised for a given fish batch (found by linear regression) [4, 7, 13]. This conclusion has indeed been confirmed by our experience.

Our technique does not make any assumptions on the fish batch; its only limit is the fact that estimating weight from shape is an indirect process depending on an unobservable quantity, density. However, within this inevitable restriction, the technique has many advantages: fast, binary vision allows one to acquire many shape measurements, not one, and preserve real-time performance; learning the shape-weight relation for each fish batch allows one to adapt predictions to specific batches; robust SVM regression tolerates and compensates for errors in the approximate weight estimates (and shape measurements) of the training set.

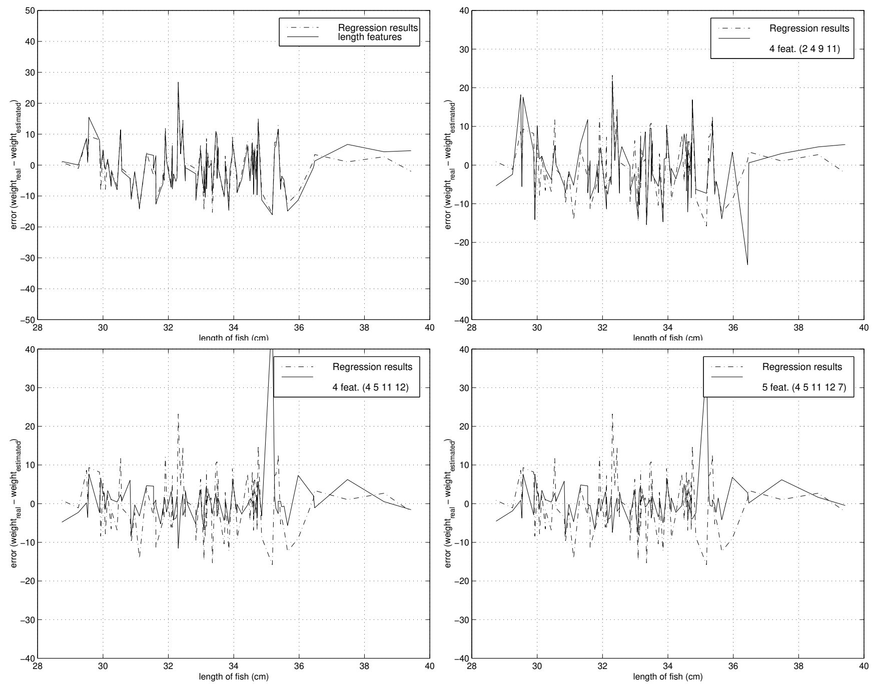

Figure 6: Average errors of weight estimates using the SVM-computed law (solid lines) and Equation 1 (dashed-dotted lines). Each graph refers to an experiment with a different set of shape parameters. Top left: 4 and 11 (first line in Table 2). Top right: 2,4,9 and 11 (second line). Bottom left: 4,5,11 and 12 (third line). Bottom right: 4, 5, 7, 11 and 12 (fourth line).

Moreover, thanks to the ability of SVMs to generalise from limited numbers of examples, reliable laws can be learnt from limited percentages of fish from a rather homogenous fish batch, such as most of those those routinely graded on fish farms.

Our experiments with a batch of 99 trouts indicate a better accuracy for the polynomial SMVs than for equation 1. Notice that the weights of the fish used (Figure 4) span a rather wide range; our experience suggests that errors in weight estimates improve significantly with limited weight ranges, which is the case with farmed fish (fish grown in a same pool or cage are kept approximately uniform in size and weight). Importantly, our technique can reveal the optimal subset of shape parameters, that is, the one guaranteeing the best accuracy for the batch, obviously a crucial information in applications. For instance, in our experiments area, perimeter and minimum width proved the most unreliable measurements, as very sensitive to the fish position; lengths and maximum widths (in both views) proved the most critical, i.e., their presence in the feature set was crucial to obtain good results, followed by the ratios between area and length. A corollary is that larger sets of shape parameters, even if plausible intuitively, do not imply necessarily better accuracies.

Acknowledgments

Thanks to Nick Yonge of Gala Fish Farming Ltd. for invaluable input, and to Stephen Cameron of Scots Trout Ltd. for providing the fish batch. This work was supported by a EU CRAFT Stage-1 contract (FAIR-CT97-8208).

References

[1] T. A. Beddow and L. G. Ross, *Predicting Biomass of Atlantic Salmon from Morphometric Lateral Measurements*, Journal of Fish Biology, vol 49, 1996, 469-482.

[2] T. A. Beddow, L. G. Ross and J. A. Merchant, *Predicting Salmon Biomass Remotely Using a Digital Stereo-Imaging Technique*, Aquaculture, vol 146, 1996, 189-203.

[3] C. Cortes and V. Vapnik *Support-Vector Networks*, Machine Learning, vol 20, 1995, 273-297.

[4] J. Doyon, J. A. Downing and E. Magnin *Variation in the Condition of Northern Pike, Esox Lucius*, Can. J. Fish. Aquat. Sci., vol 45, 1988, 479-483.

[5] T.W.Fulton *The Rate of Growth of Fishes*, Fisheries Board of Scotland Annual Report 22 (part 3), Edinburgh, 1904, 141-241.

[6] R.M. Haralick and L.G.Shapiro *Computer and Robot Vision* Addison-Wesley, 1992.

[7] B. Knights *Body Dimentions of Farmed Eels (Anguilla Anguilla L.) in Relation to Condition Factor, Grading, Sex and Feeding*, Aquacultural Engineering, vol 1, 1982, 297-310.

[8] N. J. B. McFarlane and R. D. Tillett, *Fitting 3-D Point Distribution Models of Fish to Stereo Images*, British Machine Vision Conference BMVC97, 1997, 330-339.

[9] E. Osuna, R. Freund and F. Girosi *Training Support Vector Machines: an Application to Face Detection*, Computer Vision and Pattern Recognition CVPR97, 1997, 130-136.

[10] M. Pontil, A.Verri *Support Vector Machines for 3D Object Recognition*, PAMI, vol. 20, 1998.

[11] W.E.Ricker *Computation and Interpretation of Biological Statistics of Fish Populations*, Bulletin of the Fish. Res. Board of Canada, vol 191, 1975.

[12] B. Scholkopf, K. Sung, C. Burges, F. Girosi, P. Niyogi, T. Poggio, V. Vapnik *Comparing Support Vector Machines with Gaussian Kernels to Radial Basis Function Classifiers* AI Memo No. 1599, CBCL Paper No. 142, MIT, Cambridge, 1996.

[13] R. Scott Cone, *The Need to Reconsider the Use of Condition Indices in Fisheries Science*, Trans. of the Amer. Fisheries Society, vol 118, 1989, 510-514.

[14] N. J. C. Strachan, P. Nesvadba and A. R. Allen, *Fish Species Recognition by Shape Analysis of Images*, Pattern Recognition, vol 25, 1990, 539-544.

[15] N. J. C. Strachan, *Sea Trials of a Computer Vision Based Fish Species Sorting and Size Grading Machine*, Mechatronics, vol 4, 1994, 773-783.

[16] N. Sumpter, R. D. Boyle and R. D. Tillett, *Modelling Collective Animal Behaviour Using Extended Point Distribution Models*, British Machine Vision Conference BMVC97, 1997, 242-251.

[17] V. N. Vapnik, A. Y. Chervonenkis *On the Uniform Convergence of Relative Frequences of Events to their Probabilities* Th. Prob. and its Applications, 1971, 17(2) 264-280.

[18] V. N. Vapnik *The Nature of Statistical Learning Theory*, Springer-Verlag New York Inc., 1995.