Fine scale mapping of fish distribution from commercial catch and effort data using maximum entropy tomography *

M Vignaux†
National Institute of Water and Atmospheric Research
Wellington, NZ

G A Vignaux   S Lizamore
D Gresham
ISOR, Victoria University, Wellington, NZ

Abstract

We present a technique for mapping the spatial distribution of fish using commercial catch and effort data. Maximum Entropy tomography is a Bayesian method that allows the relative fish density to be estimated at scales smaller than the length of the unit of effort such as a tow. It can take advantage of the fact that the tows cross to give information about the density in those areas. This is a novel application of a well tested technique which has been used in other fields such as astronomical imaging. Its utility and robustness is demonstrated both on simulated data and on real data from the trawl fishery on spawning hoki off the west coast of the South Island of New Zealand.

1 Introduction

The fine-scale structure of fish distribution is difficult to measure using CPUE data although it is important for interpreting it.The spatial dimensions of effort (tens of kilometres) may be large compared with the dimensions of the spatial structure of interest (kilometres). To map the fishery,

*Published as: M Vignaux, G A Vignaux, S Lizamore, and D Gresham, "Fine-scale mapping of fish distribution from commercial catch and effort data using maximum entropy tomography",*Canadian Journal of Fisheries and Aquatic Sciences, v 55, no 5, 1998, pp 1220-1227.*

†Now at NZ Seafood Industry Council, Wellington, NZ

catch and effort data are typically aggregated into blocks of size comparable with or larger than the spatial dimensions of the effort in the fishery. However, there may be spatial correlations on scales smaller than this which can significantly affect catch rates. This is the case in the trawl fishery for New Zealand hoki, as shown in **?**.

Despite this, it is possible to infer information about the distribution of fish on a finer scale than the length of the tow (typically 10–20 n. miles). This is achieved because many individual tows cross and the catch data from such tows can be used to infer the density of the fish in the small region in common to them. Two tows that pass through an area of high fish density will share a higher than average catch. Transforming from information about catches to information about where the dense patches of fish might be is complicated. This transformation involves optimisation, since the catch rate data are typically very noisy, and one cannot solve the resulting equations exactly.

This problem of transforming data from one space into information or an image in another space (sometimes known as the Inverse Problem) is common to many fields of science, and the method of Maximum Entropy (known as MaxEnt) has been found to be a useful technique in many such applications. One familiar form of the Inverse Problem is the Fourier Transform, and MaxEnt has been used effectively for problems in spectral analysis (Buck and Macaulay 1990), nuclear magnetic resonance analysis (Sibisi 1983) and radio astronomy (Gull and Daniell 1978). An even closer analogy to the problem of fish density estimation is the general technique of tomography. In medical X-ray tomography for example, X-ray scans are passed through the body in different directions and the total absorption in each pass is measured. The density of absorption centres at different locations in the body are mapped by transforming the information from a large number of scans in many different directions.

Applications of MaxEnt in tomography problems in other fields include detecting cracks in rods (such as fuel pins) by X-ray,(Freiden and Zoltani 1985), using radio waves to investigate the electrical properties of rocks in situ, (Pendock and Wedepohl 1991) and astronomical imaging of binary stars, (Marsh and Horne 1989).

In the novel application described here, the area of the fishery corresponds to a body being scanned by X-rays, the mass of fish per unit volume correspond to the density of the bones and flesh within the body, the commercial fishing tows are equivalent to X-ray scans through it, and the fish catches are equivalent to the diminution of density of the X-ray film. The work is therefore a straightforward application of a well-known technique to a completely new type of data.

A few differences should be noted however. First, the scans in medical tomography have been chosen by the operator to give the most easily transformable and unambiguous data for preparing the image. For example, they may be chosen to follow a strict grid or along regularly spaced radii. In fishing, the tows are chosen by the fishers to maximise their catches, with, unfortunately, little regard to their information content for mapping purposes. Despite this, we have been able to obtain useful images.

A second difference is that the fish distribution is likely to change as fish move during the season. In this study, the season was divided into periods which were short compared with the persistence of the fish aggregates (Vignaux 1996a) and each period studied separately.

A third difference is that the catches in the tows from which we obtain our information are not simple measurements, and are highly variable. In particular, there are coarse scale differences between tows made by different vessels and at different times of day and these have effects on catch rates which are not related to the density distribution of the fish. It was important inn our studies to remove as much as possible of this uninteresting variation. The same linear model that was used for standardising the CPUE to estimate year to year changes in fish abundance was used to measure the size and importance of each of these effects. The catches in each tow were standardised using this fitted model, to calculate a standardised catch expected in each tow if the tows had been done under identical conditions.

Even when these systematic effects have been removed, we would still expect a large degree of variability between tows. Experience shows, however, that the algorithm used to find the Maximum Entropy image is relatively robust to such variability.

2 Method

Suppose we divide the ocean area of interest into n blocks and seek to determine the average density of fish, f_j, in each of the $j = 1, 2, \ldots, n$ blocks. Information is provided via the catches in each of p tows which pass through some of the blocks, where D_i is the catch in the ith tow.

Since each fish in D_i must have been caught in one of the blocks through which tow i passed, we can write the total catch as a sum of sub-catches in each of the blocks. Allowing for random error,

$$D_i = \sum_{j=1}^{n} R_{ij} f_j + \epsilon_i. \tag{1}$$

Here R_{ij} is the distance that tow i passed in block j (in nautical miles), in which fish were present at a relative linear density f_j kg per nautical mile. ϵ_i is a random error mean 0 for the ith tow. It is assumed to include a variety of unmeasurable effects such as variable catchability and observational errors.

Although we may not be able to invert this equation to find the values of f_j directly, we can use the D_i values to find the most plausible values of f_j that are consistent with the data D_i. This is done using a Bayesian method, in which the posterior probability of a particular set of f_j values (known as a candidate image) is based on both the prior probability - the plausibility of the image based on what we know about the fishery generally - and on the likelihood of obtaining the data if the trial image was in fact the true image.

The MaxEnt method uses the information entropy (also known as the configurational entropy) of a trial image to establish its prior probability. An image that varies a great deal is defined to have a lower prior probability than a smooth image. This is because in the absence of data to the contrary, we would expect the density distribution to be smooth, perhaps even uniform over the area of sea. This bias towards smoothness means that all structure seen in the final image must have come from the data and cannot be an artifact (See Daniell, 1991).

In MaxEnt (e.g. Buck and Macaulay, 1990), the prior probability of a candidate image $\mathbf{f}'$ is:

$$P(\mathbf{f}'|m) \propto exp(\alpha S(\mathbf{f}', \mathbf{m})) \tag{2}$$

where $\mathbf{f}' = \{f'_1, f'_2, \dots f'_n\}$ is the vector of f'_j values, $\mathbf{m}$ is the vector of m_i, representing what is called the null model for the density distribution in the absence of data (taken in this case to be uniform), α is a constant, and $S(\mathbf{f}', \mathbf{m})$ is the information entropy of the trial image $\mathbf{f}'$ relative to the null model $\mathbf{m}$:

$$S(\mathbf{f}', \mathbf{m}) = \sum_i \left[f'_i - m_i - f'_i \log(\frac{f'_i}{m_i}) \right]. \tag{3}$$

The likelihood of obtaining the dataset $\mathbf{D}$ if the true image was $\mathbf{f}'$ and assuming normal errors is

$$P(\mathbf{D}|\mathbf{f}') \propto \prod_i \exp(-\frac{(F'_i - D_i)^2}{2\sigma^2}) = \exp(-L(\mathbf{f}')), \tag{4}$$

where the

$$F'_i = \sum_{j=1}^{n} R_{ij} f'_j \tag{5}$$

are the predicted catches given the trial image, and σ is an estimate of the variance of the expected noise. The likelihood therefore measures the size of the discrepancy between the set of predicted catches given the trial image, and the true set of catches. If the trial image was very good (i.e., it fitted the data well) the likelihood would be large. Hence, using Bayes' theorem, the posterior probability of the trial image $\mathbf{f}'$ is

$$P(\mathbf{f}'|\mathbf{D}) \propto P(\mathbf{D}|\mathbf{f}') \times P(\mathbf{f}'|\mathbf{m}). \tag{6}$$

To find the most plausible image (i.e. to maximise the Bayesian posterior), we begin with a trial image $\mathbf{f}'$ and in a series of iterations move towards the $\mathbf{f}'$ which maximises $\exp(\alpha S(\mathbf{f}', \mathbf{m}) - L(\mathbf{f}'))$. The algorithm for this is as described in Skilling and Bryan (1984) and a particularly clear exposition of the method is given by Sivia (1996).

The assumed noise parameter σ of equation (4) is provided to the algorithm by the user. If σ is set too low, the algorithm is expecting the catches D to be fitted so precisely that it will never find an acceptable solution. If σ is set too high, the algorithm will allow quite a poor fit between D and F', and may terminate abruptly without trying a number of solutions. For values between these the algorithm will iterate to find the most probable solution. Since σ is unknown in this application, the algorithm was run a number of times, to find the minimum value of σ for which convergence is still possible.

The null model, $\mathbf{m}$, which is also known as the Lebesque Measure (Sivia, 1996), expresses the idea of a pre-prior density. In most cases this will be uniform over the region so that all m_i are the same constant but this choice depends on the details of the application. Theoretically $\mathbf{m}$ should be normalised but normalisation will in any case occur naturally during the fitting process.

The parameter α describes the trade-off between optimising the entropy of the prior, and optimising the fit of the image given the data. We have used the usual approach of setting it so that the misfit of the image (the logarithm of the likelihood or $\sum_i (F_i' - D_i)^2/2\sigma^2$), is equal to the number of records in the dataset, since that is approximately the number of degrees of freedom in the problem. The reasons for this choice are described in Sivia (1996).

All data manipulation and the production of figures and images in this study were carried out using S-PLUS functions which extracted the data sets, set up the parameter files for the final inversion, called the optimisation programs, read back the solution files and plotted the graphs and images.

The actual optimisation of the trial image f' was done with separate C programs based on Macaulay's, 1992 modifications of the original published algorithm by Skilling and Bryan (1984).

3 Testing the Method

The MaxEnt tomography method has been used and tested in other areas of application but the differences here were sufficient to persuade us to test the method with randomly distributed tomography lines and to investigate its sensitivity to the choice of other parameters. A 10 by 10 square artificial pattern of varying density between 1 and 9 in the figure of a face was generated (Figure 1a). The mean level of the image is 3.94 (standard deviation 2.36). Then 150 lines, representing tows, were passed across the image from random start to random finish points within the image (Figure 1b). The mean line length was 5.44 units (standard deviation 5.50). Because all the points were generated within the image the edge blocks of the figure have fewer lines passing through them and we would expect a poorer fit of the image there. To simulate the "catch" in each tow the density was integrated along each line giving a mean of 22.30 (standard deviation 11.57). 64% of this variance would be generated even with a constant density image because of the variation in line lengths.

The simulated "catches" and locations of the lines were passed to the maximum entropy tomography algorithm which was able to regenerate the figure of the face (Figure 1c). A number of runs were made, varying the control parameter, σ, to find the smallest value that would still produce a solution. The algorithm would be unable to terminate with a value of σ smaller than this. This choice of σ should therefore give an image with maximum contrast. The sum of squares of the densities of the original image was 2106. After optimisation, the sum of squares of differences between the fitted image and the original was 51, explaining 97.58% of the sum of squares.

A number of further runs were made, in each case adding normal variate noise (mean 0) to the catches. Figures 1d to 1g show the images for noise standard deviations of 0.1, 0.3, 1.0 and 3.0, respectively. As noise is added, the image degenerates, becoming less and less intense as the added noise increases until it relapses into a uniform field. This demonstrates one of the advantages of the method: that it does not invent patterns that do not exist. The resulting residual sum of squares (out of an original 2106) are shown in Table 1. Further testing could have been done with non-normally distributed noise.

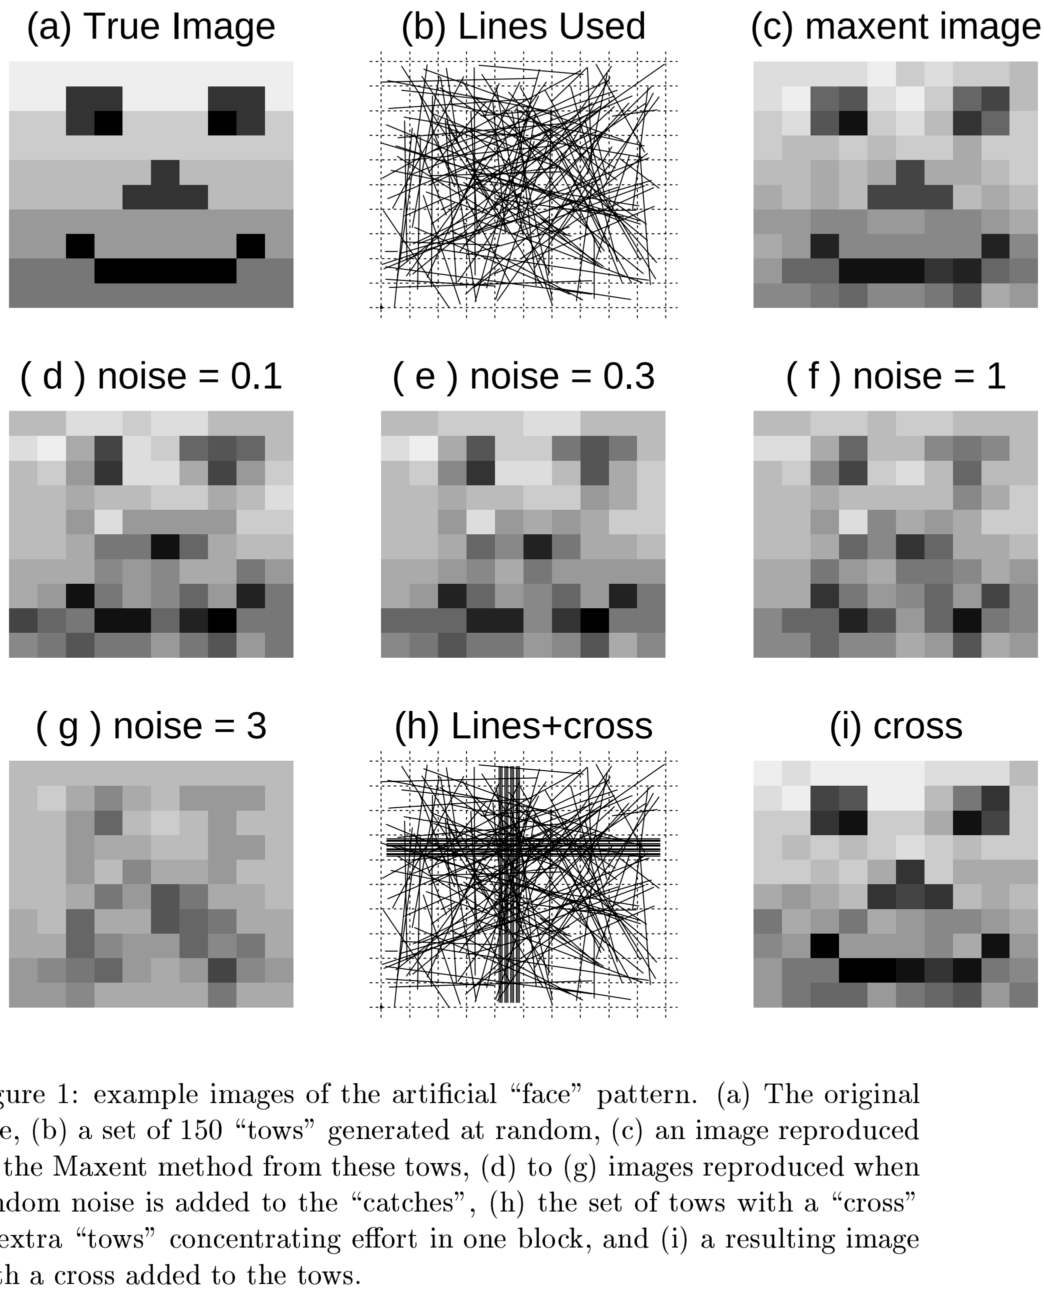

Figure 1: example images of the artificial "face" pattern. (a) The original face, (b) a set of 150 "tows" generated at random, (c) an image reproduced by the Maxent method from these tows, (d) to (g) images reproduced when random noise is added to the "catches", (h) the set of tows with a "cross" of extra "tows" concentrating effort in one block, and (i) a resulting image with a cross added to the tows.

Noise s. dev	Residual SSQ	Propn of SSQ Fitted
0.0	51	0.9758
0.1	157	0.9255
0.3	191	0.9093
1.0	252	0.8803
3.0	396	0.8120

Table 1: The esidual sums of squares (SSQ) when fitted with different levels of noise added to the catches. The last column shows the proportion of SSQ explained by the fit. As the noise increases, the fit gets poorer.

As the fishery data have the additional disadvantage that the locations of the tows are very uneven, we tested the behaviour of the method when there were many more tows through certain squares than through others. We wanted to ensure that extra effort in a block would not, in itself, cause the system to conclude that there was a greater density in that block. We added a cross of 12 vertical and 12 horizontal lines passing through one square as shown in Figure 1h. The added lines should even improve the fit of the resulting image as extra information is being provided to the algorithm. Figure 1i shows this. The final fitted image is improved with a residual sum of squares of 36.7 (Fitted correlation of 0.9826) compared with 51 (0.9758) for the image without the cross (Figure 1c). Examining the resulting image closely shows that densities are indeed closer to the original for those rows and columns containing the cross.

To emphasise this point, the density difference between the true image and the fitted image for the 100 blocks in the face image were plotted against the total length of tomography lines in the square, corresponding to the effort in that region (Figure 2). The cross was not used in this example. While the variance decreases as the length increases (as the density in the square becomes better defined with more data), those blocks with more lines running through them have the same zero average difference in density as those with few lines.

4 Data standardisation

As stated previously, the quantity of fish caught in a tow depends not only on the density of fish but also on the size of the net, characteristics of the

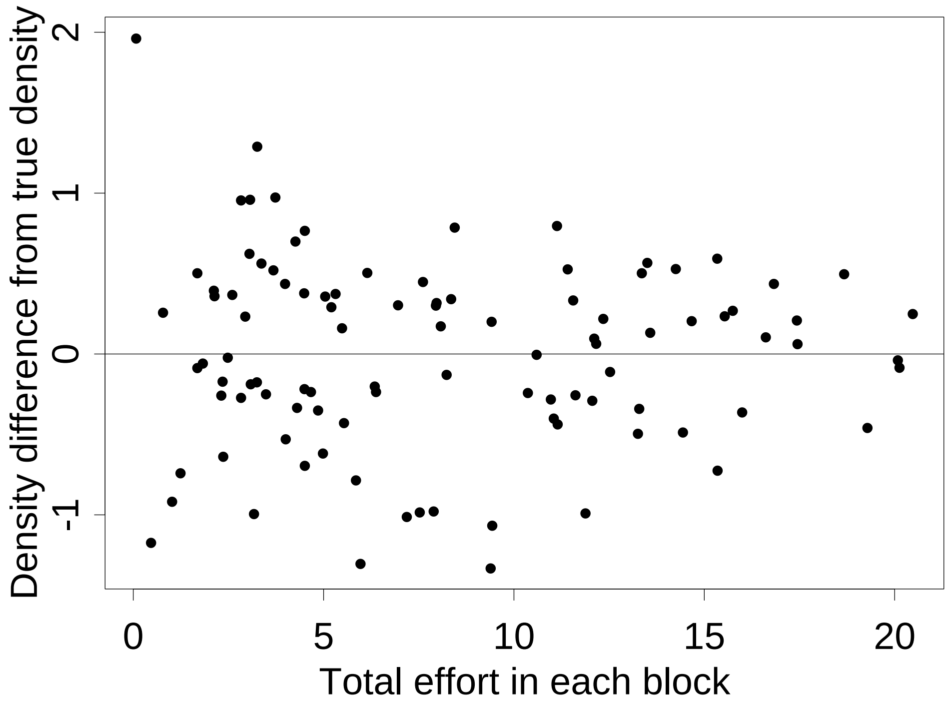

Figure 2: Density differences between the true face image blocks and a Maxent reproduction versus the total length of "tows" passing through each block (the effort). The horizontal line indicates zero differences between original and image blocks.

ship towing it, and other factors. A linear multiplicative model was fitted to standardise the catches so that data from tows would be more comparable:

$$\frac{D'_i}{R_i} = k_{1,i}k_{2,i}k_{3,i}\ldots\exp(e_i). \tag{7}$$

Here D'_i is the observed catch in tow i, R_i is the length of that tow, $k_{1,i}$, $k_{2,i}$, $k_{3,i}$, etc. are fitted regression terms reflecting the effect of other factors on the catch. e_i is the residual from the fitted model and represents fluctuations in catch rate that cannot be explained by other factors. Some of this residual may be explained by spatial variation in fish density on a local scale. The expression is made additive by a logarithmic transformation.

The model was fitted using stepwise regression. The variables included were: time within the season, vessel nationality (a surrogate for many behavioural factors), time of day, year, vessel length, latitude and headline height. The most significant variable is a seasonal effect, modelled as a cubic continuous variable on the day of the season. The second most important variable is nationality, then come the time of day and year. The last is a surrogate for changes in total biomass from year to year. Vessel size, net size and tow latitude are minor variables. All the variables explain only about a quarter of the total variance. Apart from any variation due to spatial fluctuations in fish density there will also be a great deal of noise, and variations due to other undocumented changes in fishing efficiency such as net configuration and water currents. The model is described in more detail by **?**.

A standardised catch is defined as:

$$D_i = \frac{D'_i}{k_{1,i}k_{2,i}k_{3,i}\ldots} = R_i\exp(e_i) = \sum_{j=1}^{n} R_{ij}f_j + \epsilon_i. \tag{8}$$

This is the equivalent of Eq.1.

While all factors should be fitted in the regression model, it is not necessary to use all factors $(k_{1,i}k_{2,i}k_{3,i}\ldots)$ in the standardisation of equation 8. For example, variability due to seasonal and spatial effects in the D_i should not removed by standardisation.

5 Background to the fishery

Hoki (*Macruronus novaezelandiae*) is a gadid species widely distributed in New Zealand waters. The largest stock, known as the western stock, is thought to have feeding grounds in the Sub-Antarctic region south of New

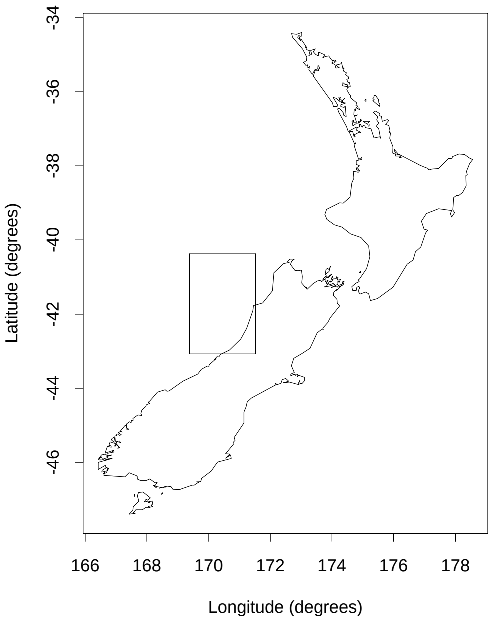

Figure 3: The location of the New Zealand hoki spawning fishery. The box corresponds to the study area used in figures later in the paper.

Zealand, and to migrate to the West Coast of the South Island (Figure 3) to spawn in July and August each year (Sullivan et al. 1995). The 1994 stock assessment estimated the biomass in 1995 of this stock to be 1.54 million tonnes, with a virgin biomass of 1.69 million tonnes and a long term average yield of about 185 000 $t.yr^{-1}$ (Sullivan et al. 1995).

There is a concentrated trawl fishery on the spawning aggregations, consisting mainly of large foreign vessels chartered by New Zealand companies. In 1993 there were 61 vessels in the fishery, in an area about 130 n. mile long by 30 n. mile wide. The vessels range from 25 to 105 m overall length, and come from a number of countries, including Japan, Korea, and Norway. They use mid-water or bottom trawl nets in depths of 300–700m. All vessels fishing in New Zealand waters are required to report both their catch and effort on forms submitted to the Ministry of Agriculture and Fisheries. Most of the larger vessels must include details about each tow.

6 Results

We applied the method to CPUE residuals data from hoki seasons 1991 through 1993. The area of the fishery chosen for analysis was a 96 by 160 n. mile rectangle with the lower left-hand corner at point 169° 22.5'E, 43 ° 4.5'S. It was divided into rectangular blocks for each of which we would estimate a fish density. Several block sizes were tested (see Lizamore MSc thesis, 1995)) and a block size of 8 by 8 n.miles was finally used for the analyses reported here. This accords well with Vignaux, 1996a, who gave evidence for spatial correlation in this fishery up to 11 n.miles. Apart from choosing an appropriate block size we assumed no spatial correlation between blocks though one would hope for better fitted images if such information was incorporated into the model. Work is progressing to incorporate spatial correlation as well as other information into the model.

The data for a season was divided into time periods, each period analysed independently. One would choose the shortest period possible balanced by the need for sufficient trawls to generate a useful image. Periods of a day, a week and a month were all tested and, for this paper, 7-day periods were used. Vignaux, 1996, gave evidence of temporal correlations up to 14 days. It was therefore assumed that within a 7-day period the pattern of fish distribution would be relatively stable. No further temporal correlation was included in the model. Each period was analysed independently with the possible loss of some quality of image generation.

The null model, $\mathbf{m}$, was set to be a uniform distribution over the area, that is, the pre-prior probability of fish in each of the squares was assumed to be the same. As the area chosen includes land in the South-Eastern corner (as is shown in Figure 3) we might have set the null model to zero in these squares, reflecting the low probability of catching fish there. Based on prior knowledge of the fishery, one could also argue that there should be an enhanced prior probability in the Hokitika Canyon which forms a segment of the analysis area and that the m_i in that area should be higher than elsewhere. For simplicity, we chose not to vary m_i in this way.

Figures 4 to 6 shows a sequence of density plots for 1-week periods in 8 $\times$ 8 n. mile squares for the whole of the 1991 season.

Figure 4 shows the tows recorded for each period. Anomalous tows appear in some of the periods. The effect of their removal is described in the sensitivity studies outlined later in this section (Figures 7 and 8). Figure 5 shows the results of the tomographic fit in the form of densities in 8 $\times$ 8 n. mile squares. The greatest density of fish is represented by the blackest squares. Each period is analysed and normalised independently. Figure 6 shows this same result in the form of density contours which smooth the density changes at the boundaries of squares. This presentation is appealing but may misrepresent the densities as the darker areas in this plot correspond to the steeper gradients rather than the highest density levels.

Where there are no tows both the squares and the contours (apart from smoothing) are blank as no information about density is provided where there is no effort. This highlights the fact that the data come from the activities of the commercial fishers who are hoping to concentrate their effort in high density locations and avoid low density areas. However, Vignaux (1996b) has produced evidence that fishers are not entirely successful at this, and often do not put greatest effort into the areas with greatest fish density. This is probably because they do not have enough information to know where these areas are.

Even though each picture is obtained from a completely independent set of trawl data, the density patterns seem to change in a smooth way from one week to the next, suggesting that interpreting the dense areas as patches of fish would not be unreasonable. Figure 6 suggests a high density conglomeration of fish in the south at the start of the season which grows, spreads north and then splits into two. The southern concentration of fish remains throughout the season, but seems to contract.

If patches of fish move on time scales of weeks, we might expect that from one day to the next that there be relatively little movement. This would suggest that the images taken from two days should be relatively similar. However, on a daily scale there will be few tows in many of the squares so we

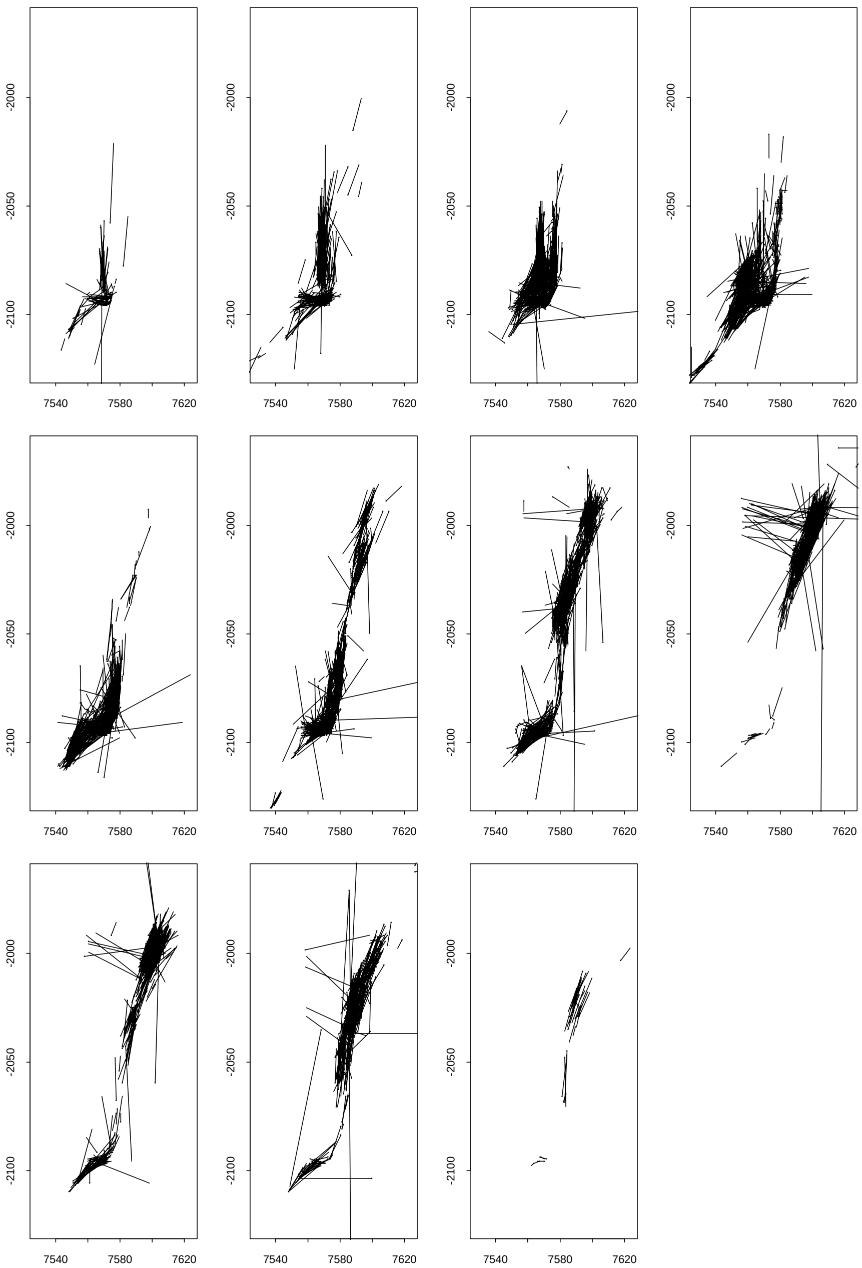

Figure 4: Tows undertaken in each week of the 1991 season in the area corresponding to that outlined in Figure 3. Axes are labelled in nautical miles relative to the point 0°E, 0°S.

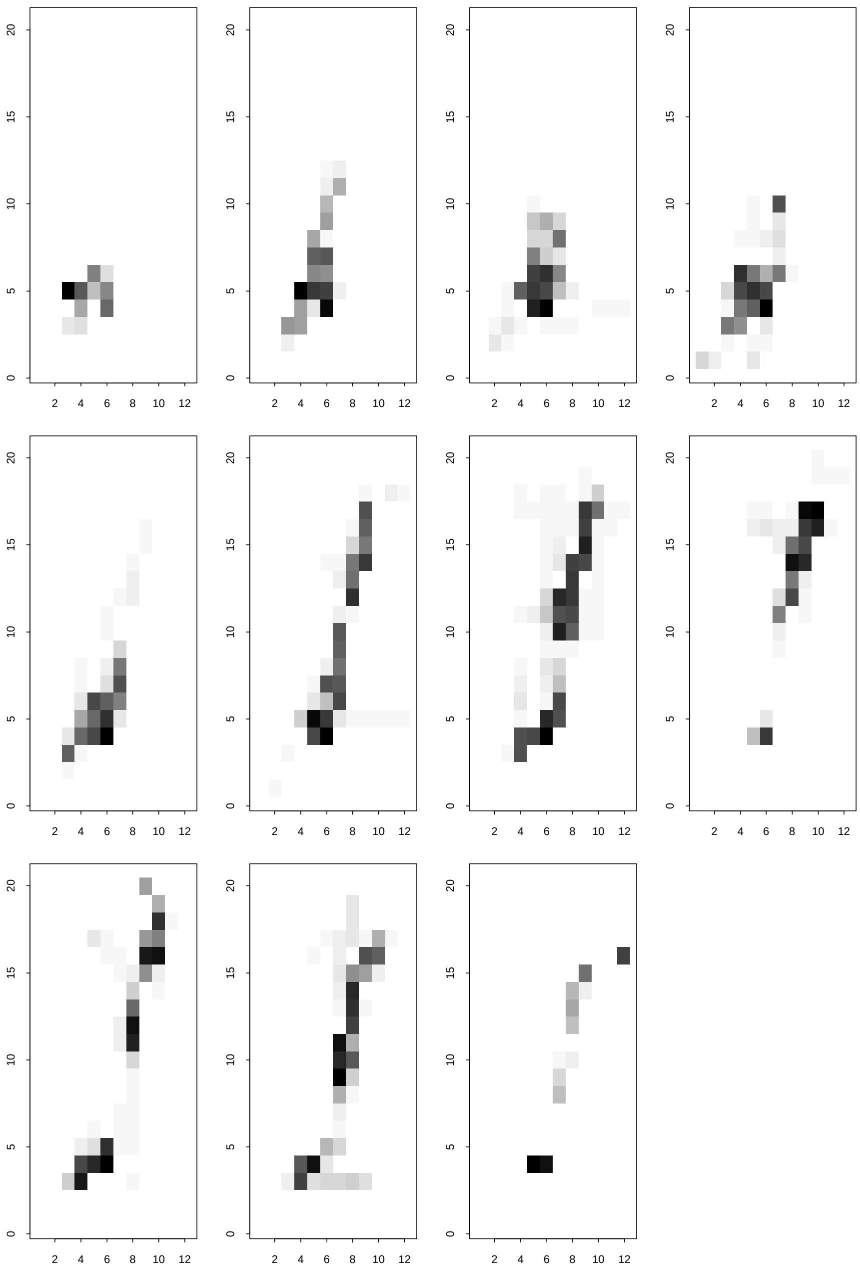

Figure 5: Image densities for each week of the 1991 season. Each small figure shows the same area as in Figures 3 and 4. Each is divided into 8×8 n. mile blocks with the lower left-hand corner being the point 169° 22.5'E, 43 ° 4.5'S. In this and the next figure the axes show the block locations.

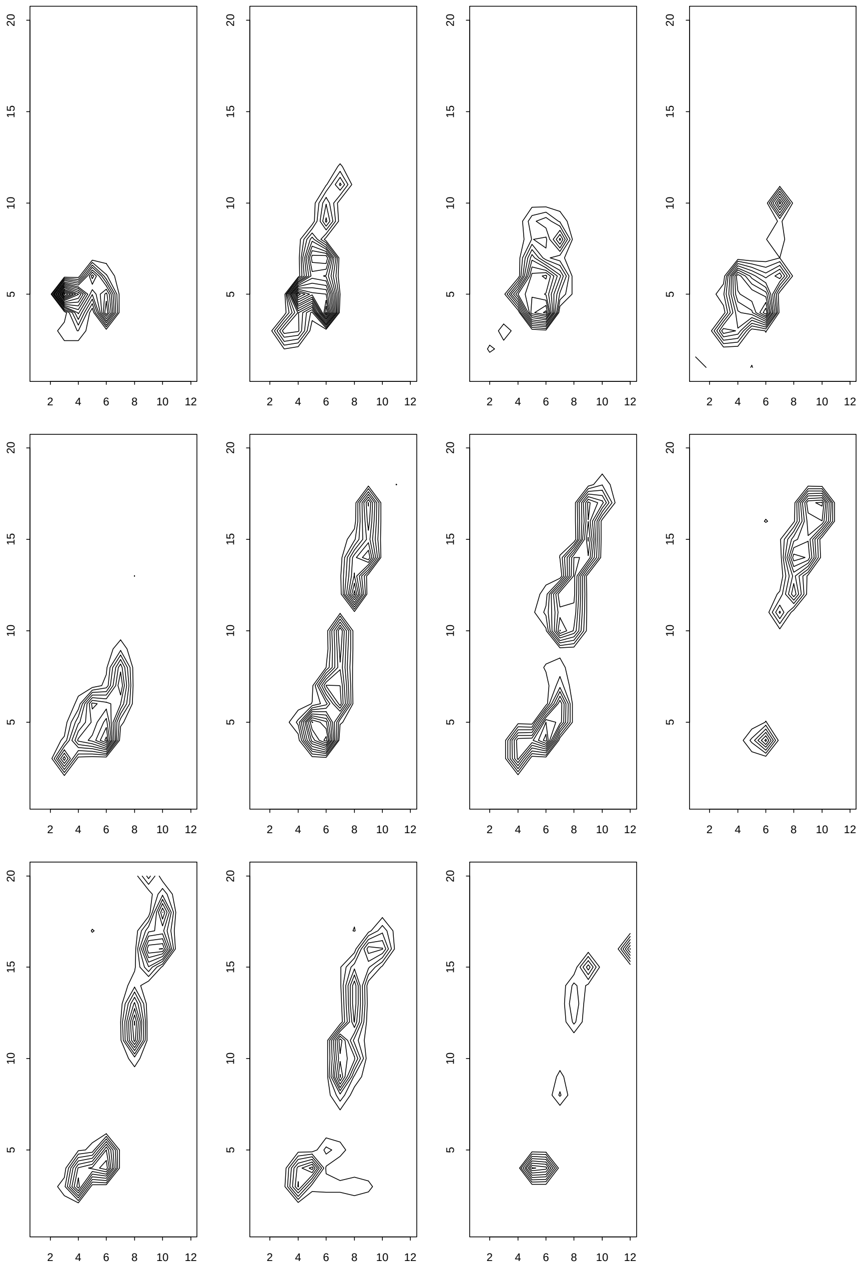

Figure 6: Image contours for each week of the 1991 season. Each small figure shows the same area as in Figures 3, 4, and 5.

might expect poor quality images of density. We extracted 14 consecutive 1-day data sets from the 1993 season to investigate this. We took all squares where there was data available on two consecutive days, and calculated the correlation in the densities. A high correlation would indicate that a square that had high image density on one day is more likely to have a high image density on the next and vice versa. Although the images on this level display a large degree of variation, a small but statistically significant correlation was found between the data on consecutive days. The non-parametric rank order correlation coefficient was 16.4%, which with 435 degrees of freedom has a t value of 3.48 (p=0.0003) showing that this result is significant at the 1% level. However, pairs two days apart have a correlation coefficient of only 6.9%, which with 363 degrees of freedom has a t value of only 1.31. The probability of obtaining a t value as large as this under the null hypothesis that the two patterns are independent is 9.5%. That is, after two days the daily data are not good enough to detect any correlation (at a 5% significance level). There is too much noise and there are too few records when only a day's data is used.

The model assumes that the tows are straight lines between the start and end positions entered on the fisher's form. In fact, descriptive records by Scientific Observers posted on some of the vessels suggested that of the (not necessarily random) sample of tows observed in the 1993 season, 73% were straight, and of the rest, most involved a single U turn, coming back on the same path.

Another way of checking whether a tow is straight is to compare the displacement between the start and end positions of the tow with the distance that a vessel could travel given the recorded speed and start and end times. As seen in Figure 7 most tows go directly from start to end position (within the errors associated with the start and end positions, start and end times and particularly with the speed of the tow). But a significant number of tows achieve less than half the displacement possible. There is also a smaller number of tows which apparently travel much longer distances than are possible in the time. These are clearly reporting errors.

The two images in Figure 8 show the effect of removing those tows outside the boundary lines (i.e. those whose displacement to distance ratio is less than 0.7 or greater than 1/0.7). This left 623 of an original 1007 tows. The figures also differentiate between the blocks where no tows passed (white) and those where there was at least one tow but no inferred density (light grey). The densities have been adjusted by a square root scaling before plotting to produce a better contrast between the light grey (no density) areas and those where a density is inferred by the algorithm. The density

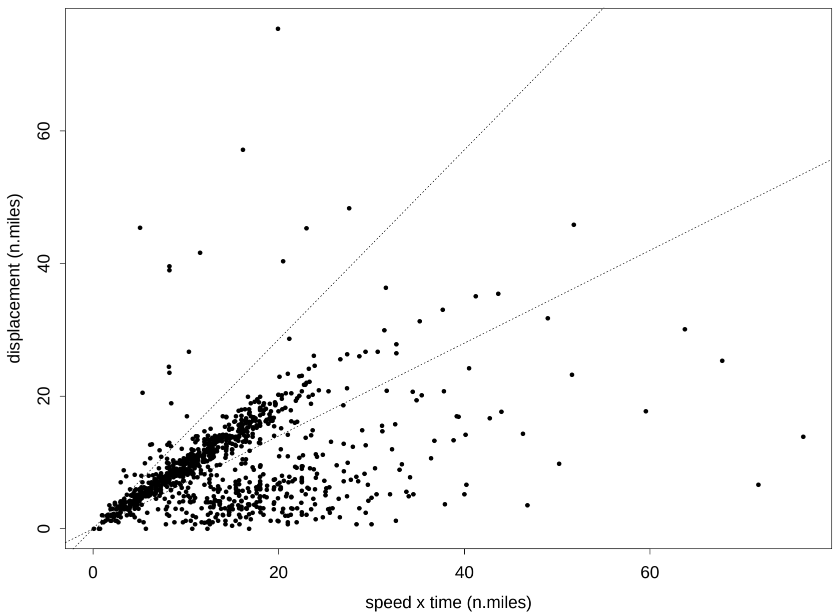

Figure 7: A scatterplot from 1993 week 6 of the distance travelled between start and end positions of a tow versus the product of speed and duration of tow (the maximum possible distance travelled if the tow was straight). The lines indicate the boundaries of those considered to be "straight".

pattern is similar in the two diagrams but some of the outlying dark areas have disappeared. The correlation between the two images is 0.62 (using the 77 blocks in common).

To test the robustness of the images produced we randomly removed tows from the dataset, starting from the set of 623 "straight" tows, and recalculated the image. We then calculated the correlation between the old and the new dataset using only blocks where there was effort in both datasets. This process was repeated ten times at each level of removals. The results are shown in Figure 9 with a boxplot at each number of removals. The average observed correlation falls slowly as the tows are removed and does not drop below 0.9 until 100 are removed (i.e. about 16% of the original 623 tows). This shows that the main features of the image are robust, although details may be sensitive to some features of the data. The ultimate safeguard is that, as the information is reduces, the image becomes less and less intense asymptotically reducing to uniformly grey. This is demonstrated by the simulation experiments.

7 Conclusions

Tomography is a well known technique in a variety of different fields, and the MaxEnt algorithm for solving the tomography equations is well tested. It has not (to our knowledge) been previously applied in this novel way. However, there is an elegant analogy between X-ray beams probing the density distribution inside a body and tows probing the density distribution of a fishery. This means that the mathematical structure of the fishing situation maps neatly into the mathematical structure of the tomography problem.

Thus MaxEnt tomography is a useful technique which allows fine scale maps of fish density distribution to be generated from CPUE data. The tomography method requires data consisting of totals of some quantity accumulated over a distance. Trawl, long-line or dredge CPUE are examples of datasets which would be appropriate for this purpose. It is desirable that the data be dense, with numbers of tomography lines passing through the same squares, so that the fish density in each square can be determined in a robust way from a large number of data values. Data from an intense fishery, such as a fishery on a spawning stock is therefore ideal. The inherent noisiness of CPUE data and the non-uniformity and non-randomness of the tow locations do not seem to have prevented the method from producing valid images (as shown by the significant correlations of the images from

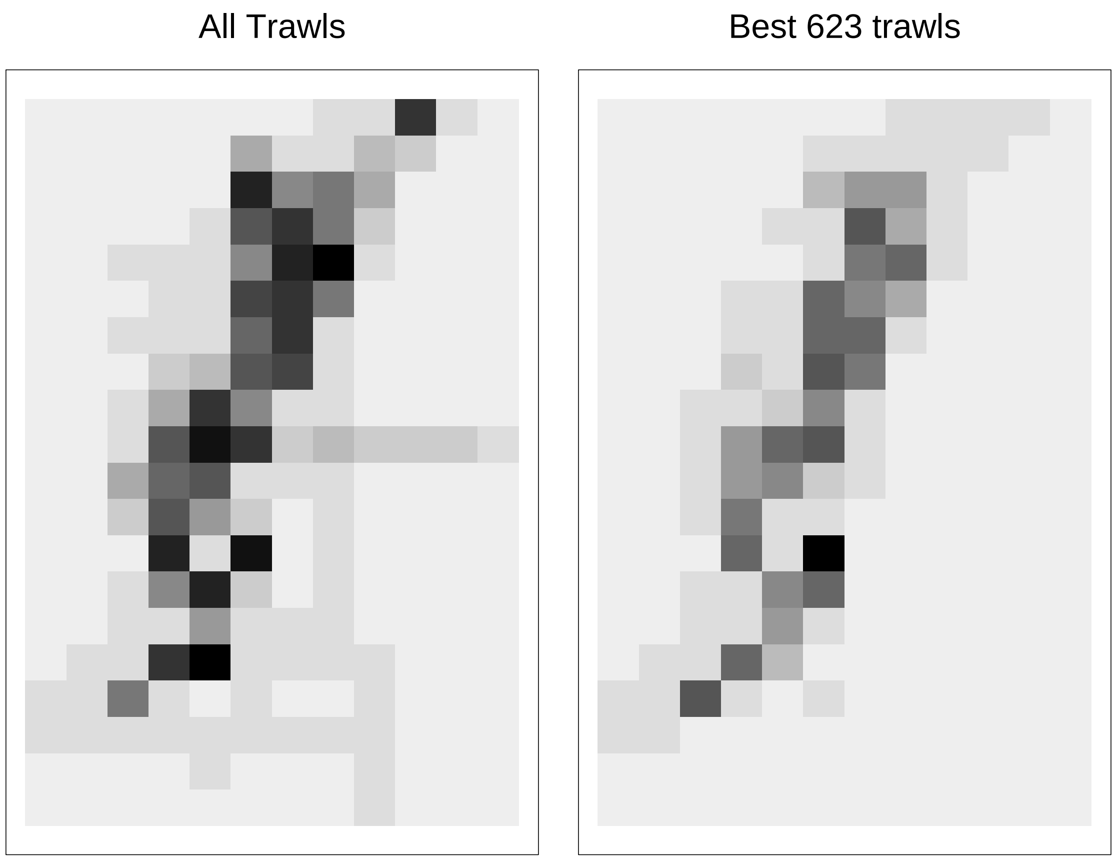

Figure 8: Images of the fish density for the week of 19-25 July, 1993 (Week 6 of the season). The left-hand figure shows the density with all reported tows, the right-hand picture shows it with all tows with the straightness ratio outside the range 0.7 to 1/0.7 removed. The white area indicates those blocks for which there is no information. The densities are on a common density range but the original densities are scaled using a square-root scaling to bring out the detail of the images and the zone of no information.

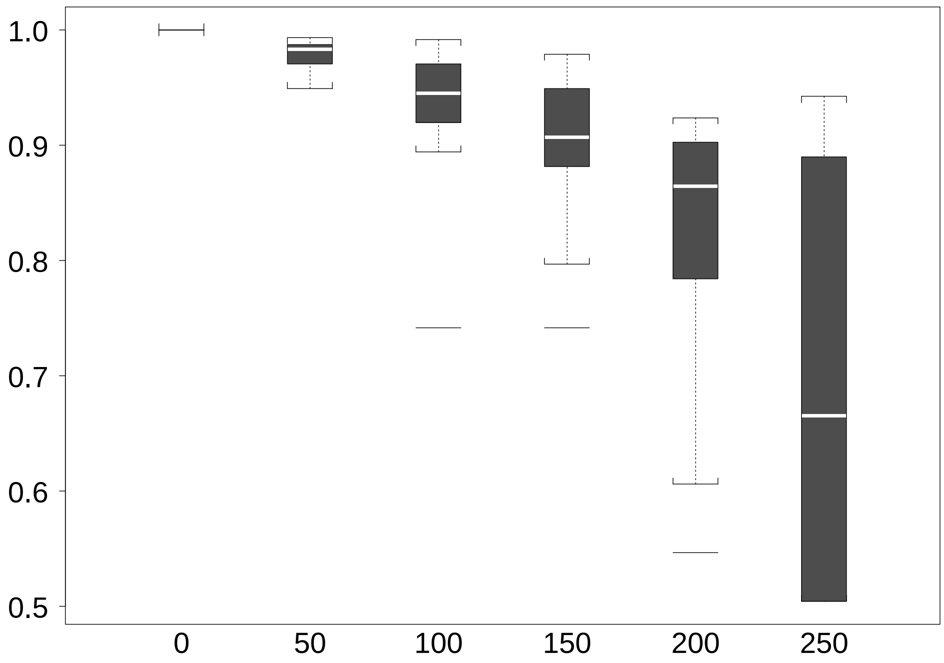

Figure 9: Correlation between the original image and a series of images recalculated with 50, 100, ... 250 tows removed from the dataset at random. The original dataset (giving the base image) is for 1993 week 6 and contains the 623 straightest tows.

consecutive days, based on independent data sets).

The method allows mapping of fish distributions from commercial data where there are many vessels collecting data. This has the advantage that there is a great deal of data coming in from many areas at the same time, so you do not need to assume that the fish distribution is static over the time scale that it takes a single vessel to sample the whole area adequately. It can take significant time, even using a rapid method such as acoustics, to cover the whole fishery area, and in this time the fish distribution may have changed. The images that are produced from CPUE data are, in contrast, largely synoptic. However, because it relies on commercial data, produced by fishers who are unlikely to search the entire extent that the fish may occupy, the picture will have gaps in those areas where fishers did not operate. Also, because the data originate from many different vessels and under different conditions, the method requires that the catches be standardised, using a linear or GLM model to ensure comparability.

Once this is done, the tomography method allows fish distribution to be mapped at finer spatial scales than the size of a tow - producing higher resolution maps than you would expect were possible. **?** showed that it is possible to detect such fine scale density distributions using CPUE data, but this tomography method is necessary to produce images at the fine scale. The resolution of the image produced by this method is only limited by the number of tows and the noise in the data. Another advantage of the general method is that data from different sources such as trawls and acoustic surveys can be combined to produce a single density map.

Mapping fish distribution is an essential first step in any study of a fishery (Caddy and Garcia 1986; Hilborn and Walters 1992). We have shown that images during the season can be generated, so that fish distribution and movements can be visualised. These images give insight into the fishery that cannot be obtained simply by looking at crude maps of CPUE or maps of the effort distribution.

But determining the spatial distribution is also essential for understanding the relationship between CPUE and abundance (Cooke, 1985, and Hilborn and Walters, 1987). **?** point out that when fish distribution is heterogeneous, skippers can manipulate their CPUE by appropriate choice of fishing area, altering the relationship between CPUE and abundance. Even fine scale spatial distribution can be important in the relationship between CPUE and abundance, because even on a fine scale skippers are making decisions about where to fish (Vignaux 1996b).

Of course, the location of the fish is of great interest to the fishers themselves. In principle, if you built a linear model to standardise the CPUE data based on a large data set from a number of years, you would not expect the model or the parameter estimates to change greatly with the addition of a few extra records. You could therefore use the model to rapidly standardise the data as it came in, allowing images to be constructed week by week. This would allow visualisation of the fish distribution during the season itself for anyone who had immediate access to tow by tow data from a large number of vessels.

Our results are preliminary but show promise of development. Work is proceeding on testing the effects of including spatial correlation, uncertainty of trawl tracks and combining trawl data with acoustic survey data (see Gresham 1997).

8 Acknowledgements

We thank Chris Francis and two referees for useful comments and Valerie Chan for her contributions.

References