An Approach to Seabed Classification from Multi-beam Bathymetric Sonar Data

[1]contact@arescon.com

1 Summary

Seabed classification, the inference of physical properties of the sea bottom from multi-beam bathymetric sonar, is generally attempted by analysis of backscatter amplitudes, a concept which is successfully applied to side-scan sonar data. Backscatter from multi-beam bathymetric sonars, however, involves several different physical scattering processes that depend on the angle of incidence of the sonar beam, and which have to be accurately modeled in order to obtain a normalized brightness image of the scattering surface. This is a classical inverse problem in which some a-priori knowledge of the physical properties of the surface and its orientation relative to the instrument platform is required in order to obtain a first order model.

Commercial systems are available that use a different approach: Multivariate statistical analysis of sonar returns, is taken in order to classify single-beam echo-sounder data

Such a system can perform what has been termed *un-supervised* classification by cluster analysis: Signals are grouped into classes according to their statistical properties. For very much the same reasons, the simple statistical approach fails with multi-beam data; returns from different angles of incidence, and thus different scattering domains, exhibit very different statistics.

The method presented in this report may be rightfully claimed by both the artificial neural-network (ANN) research community and those investigating generalized sub-space methods. Accordingly, the terminology varies with the perspective one wants to adopt. In ANN terminology the classification system proposed in this interim report employs a series of associative memories (linear NNs) which vote independently on a beams membership in one particular class. Votes are then subjected to a fuzzy decision rule which determines to which class, if any, the multi-beam return belongs.

With the limited test data available we have confirmed that the method :

- gathers returns from different scattering domains into one class.
- classifies beams, independent of their incidence angle.
- preserves the spatial resolution of the multi-beam dataset.
- is computationally feasible on a small workstation requiring minutes to load the memories with training-sets and fractions of a second to classify a beam.
- degrades in a gradual controlled way with increasing noise levels.

We now need to show, that this method actually classifies acoustic diversity related to the physical properties of the sea-bottom. A dataset collected over a

 contact@arescon.com
www.arescon.com

region with known sea-bottom geology and distinct geological features is needed to draw further conclusions.

2 Outline of the algorithm

2.1 Detection

Isolating the bottom return in each individual beams´s recording is not a trival task. Figure 1 shows full-waveform data from the 103 beams of a RESON SeaBat 8101 system and it is obvious that signal to noise ratios for far-angle beams are generally marginal (the nadir beam for an even bottom would be at zero). The full waveform was recovered from amplitude envelope and instantaneous phase data, which the SeaBat is able to record. The original 240 kHz carrier frequency is not present in the data, the frequency spectrum is shifted towards the origin, probably by subsampling.

To provide meaningful input to a classification algorithm, detection becomes a task of isolating *and selecting* returns which have sufficiently high signal to noise ratios. This is particularly important for the generation of training-sets, since noise-dominated signals will invariably bias the associative memories. Detection of a return can be viewed as a two step process. The noise level is estimated on a sample by sample basis, from the begin of the recording. Whenever a sample amplitude exceeds the adaptive noise threshold, a parametric spectrum at that particular instant in time is estimated. If sufficient power is present in the expected frequency band, the onset of a signal return is assumed. Relatively good control over signal quality can be exercised by introducing a factor $\delta > 1$ by which the signal amplitude has to be above the current noise level in order to qualify for a detection. Figure 2 shows a multi-beam profile with light-grey dots marking the footprints of available beams. The dark-grey dots show beams which were actually detected as “suitable” for further processing and it is no surprise that mostly wide-angle beams have been discarded.

2.2 Abstraction

Conventional algorithms capitalizing on statistical properties of sonar backscatter returns usually employ a layer of abstraction from the original time-domain data. The process of abstraction often involves the generation of a feature data-base. In a face recognition application, for example, one feature may be the distance of the eyes, another the relation of that distance to the length of the nose. Features of time domain data may use a subset of Fourier- or Wavelet-transform coefficients as part

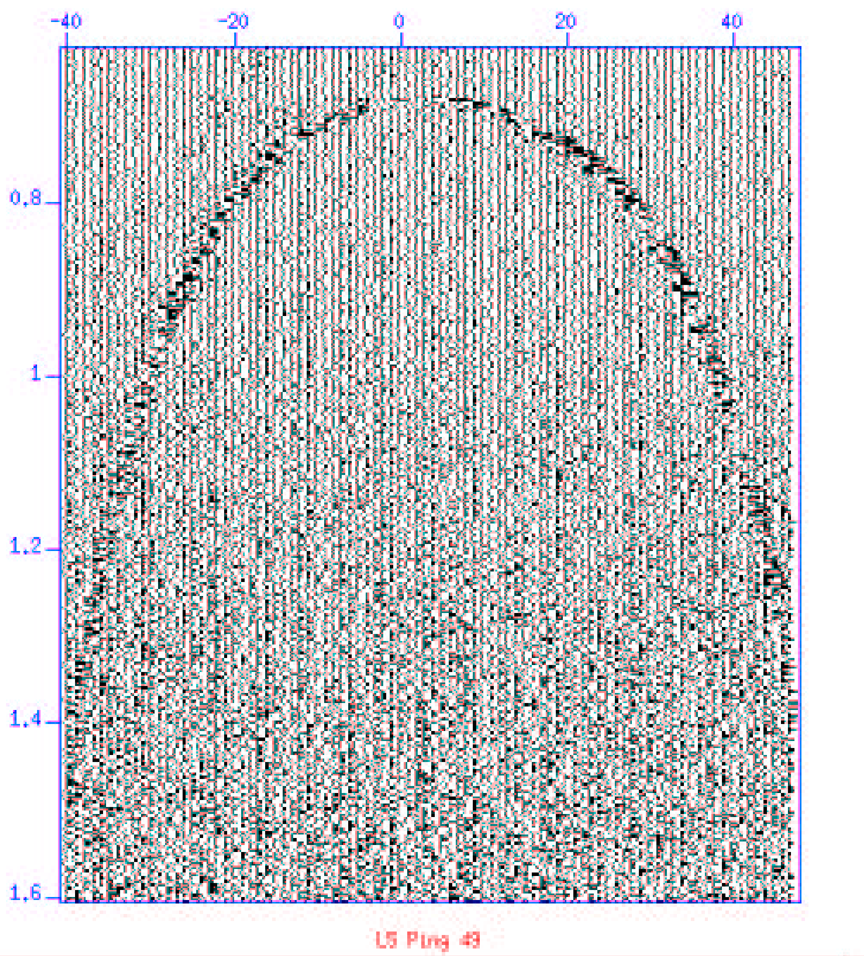

Figure 1: Full-waveform data from a Reson SeaBat Ping. Beam number zero is nadir beam. Times are in seconds x 10.

contact@arescon.com
www.arescon.com

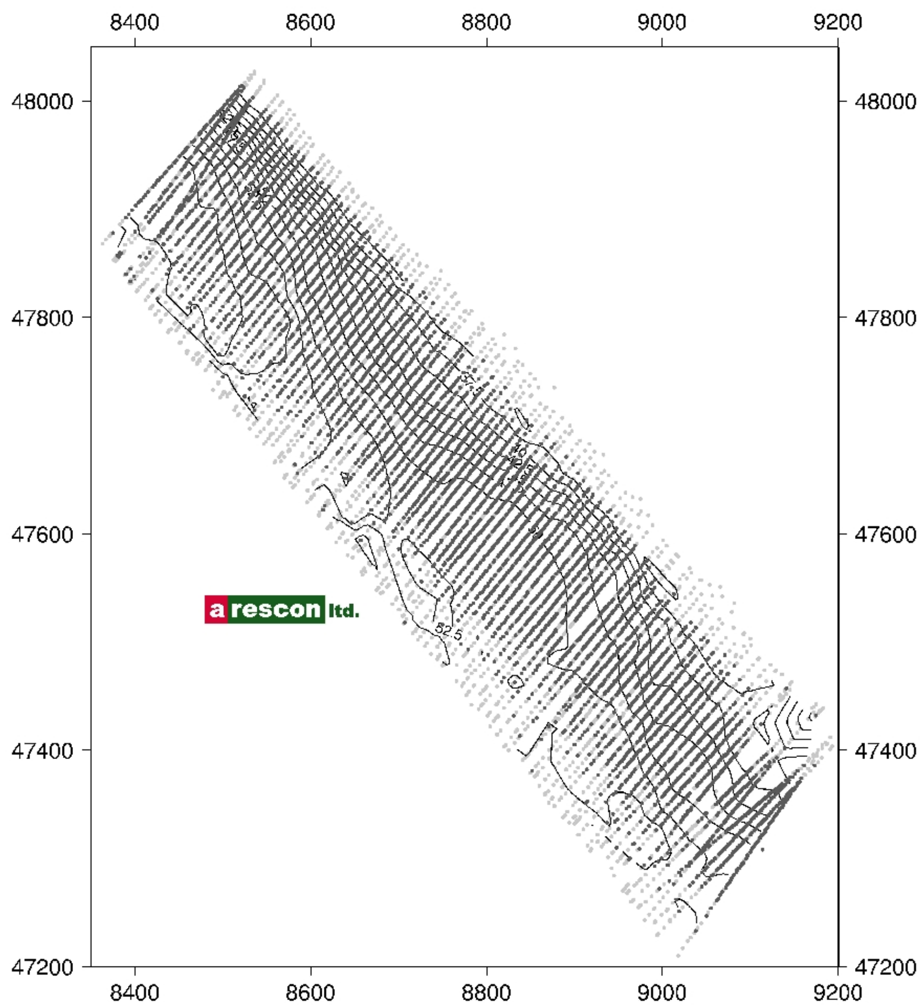

Figure 2: Beam footprints of a survey line. Light-grey dots mark available beams, dark-grey dots mark beams of satisfactory S/N. Depth contours are generated from travel-times of dark-grey dotted beams.

of an abstract representation of the data. Commonly associated with the generation of abstract features is a reduction of the data and their variance in a statistical sense. This abstraction process depends on good knowledge with respect to which features are relevant in a classification task and which are not.

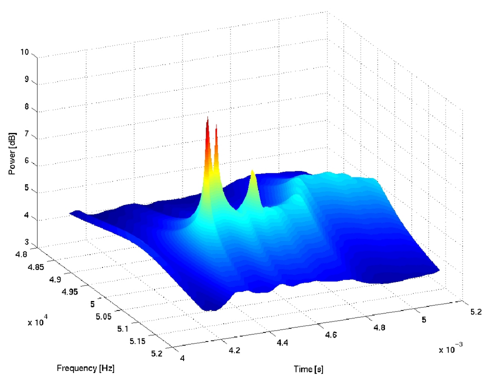

Figure 3: Time-Frequency representation of a sonar return

Since relatively little is known about the statistics of backscattering, this approach is simply not feasible for multi-beam sonar data. In fact, the proposed method seeks to increase the data variance, trusting, that the associative memories will lock on to the most salient views of the data. The capability of associative memories to generate and retain salient views has been very well demonstrated with object recognition applications.

Each individual return is transformed into a high-resolution time-frequency image, a three-dimensional representation of the evolution of the signal´s spectrum over the timespan it exists. An example is shown in figure 3. Those time-frequency images are the entities which are then employed in the classification algorithm.

3 Network architecture

In sea-bed classification the number of different acoustic classes mapped during a survey will generally be very limited. A system that is capable of discriminating 10 bottom-types in any particular survey would be satisfactory for all practical purposes. With only a small number of classes, it is possible to employ one associative memory for each individual class.

An echo-return is then presented to every one of those associative memories and the response would merely indicate whether this return is similar to the returns the memory was trained with or not. Votes from all memories are then presented to a fuzzy decision rule, which ultimately assigns class membership of the echo and a level confidence for that decision. Figure 4 shows a schematic diagram of the classification network.

4 Training and Classification

Associative memories do not require a large number of training samples to load.

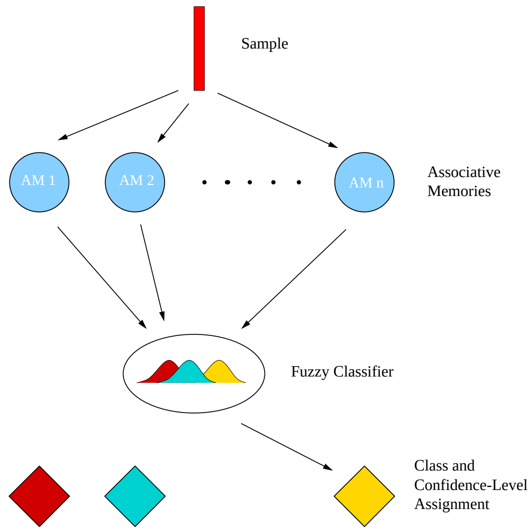

Figure 4: Schematic of classification algorithm

In the example given below, about 300 beams were used in the training sets, but less than 100 are actually required to load each class memory. Associative memories are different from other types of artificial neural networks in that there is no iterative error-back-propagation in the training procedure for example. Presenting the memory with a new training sample basically results in an update process quite similar to subspace tracking methods employed in adaptive array processing. The computational complexity is roughly of $O(N^2)$, where $N = p \times q$, with p and q being the dimensions of the time-frequency image submitted for each beam.

The classification phase differs from training mainly in the fact, that no update is performed on the memories. It is however possible to interweave classification and training in order to implement a form of re-enforcement learning. Samples presented for classification, which qualify with high confidence values for a particular class, could be used to update and re-enforce the memory associated with that class. The effects of memory re-enforcement have not been investigated at this time.

The architecture of the classifier presented in this report is inherently parallel and real-time classification would be feasible on a multi-processor system. On a single processor AMD K6-400 equipped computer under LINUX, classification of a single beam return takes about $200ms$. Updating a class memory with a single beam return during the training phase can take up to 2 seconds on the same computer.

5 An Example

Three training sets were selected from more or less arbitrary contiguous regions of the survey line in figure 2. Those beam subsets are shown in figure 5 with red, blue or green colored dots respectively. Training sets “green” and “blue” reach across the whole swath and contain 4 pings each. Training set “red” contains off-nadir beams only. No background information about the sea-bottom geology was available, the red, blue and green areas are most likely totally unrelated to variations in the sea-floor and do not represent a true acoustic class. This somewhat artificial setup serves only one purpose: Any classification algorithm should recognize and classifiy members of the training-sets correctly within a larger data set.

The classification result is shown in figure 6. With the exception of a few beams, all data from the training sets have been classified correctly. The exceptions may serve to point out that the proposed classification scheme is a competitive one. Class memories virtually compete for a sample, which is presented to them. Whenever classes overlap in their statistical distribution, the fuzzy classifier assigns the most likely class membership. Consequently a training set member may end up in a different class whenever its statistics resemble the foreign class more closely.

The apparent dominance of class “red” in the classified profile is to some degree an illusion created by overlapping dots:red dots were plotted last, on top of the blue and green ones. A final processing step will remove that artifact.

5.1 Spatial relationships

So far, the classification scheme is spatially “un-aware”, i.e. the spatial relationships of neighboring beams are not considered. Beams may be classified in any random order. Since it can be expected that sea-bottom geology is uniform over contiguous regions, that a-priory knowledge should be represented in a classification scheme. This is accomplished in a very straight-forward manner. The georeferenced and classified field of beams is spatially filtered. The filter weights are the confidence values accompanying each class assignment. The class membership of an individual beam thus becomes dependent on the class membership of its neighbors weighted by the respective confidence values. The spatial filter can at the same time be used for a reduction of the data, as it has been done in the example shown in figure 7.

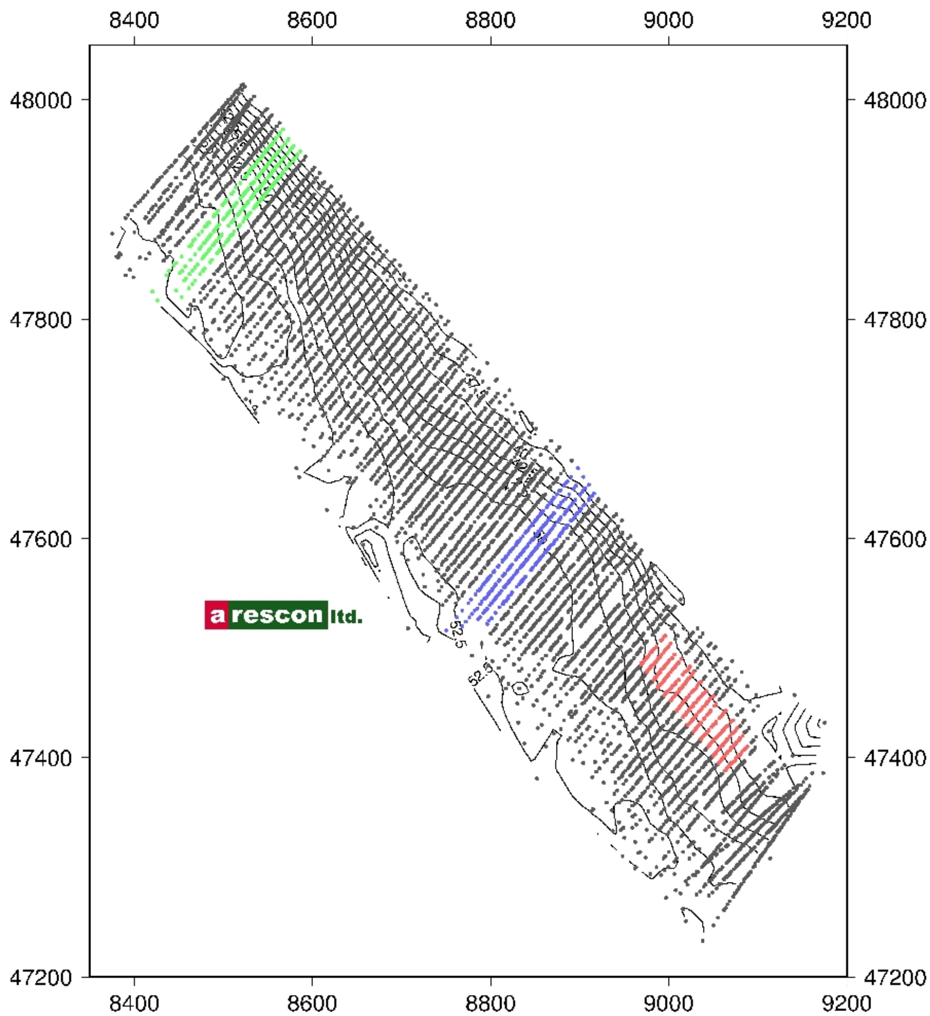

Figure 5: Location of the training sets “red”, “blue” and “green”.

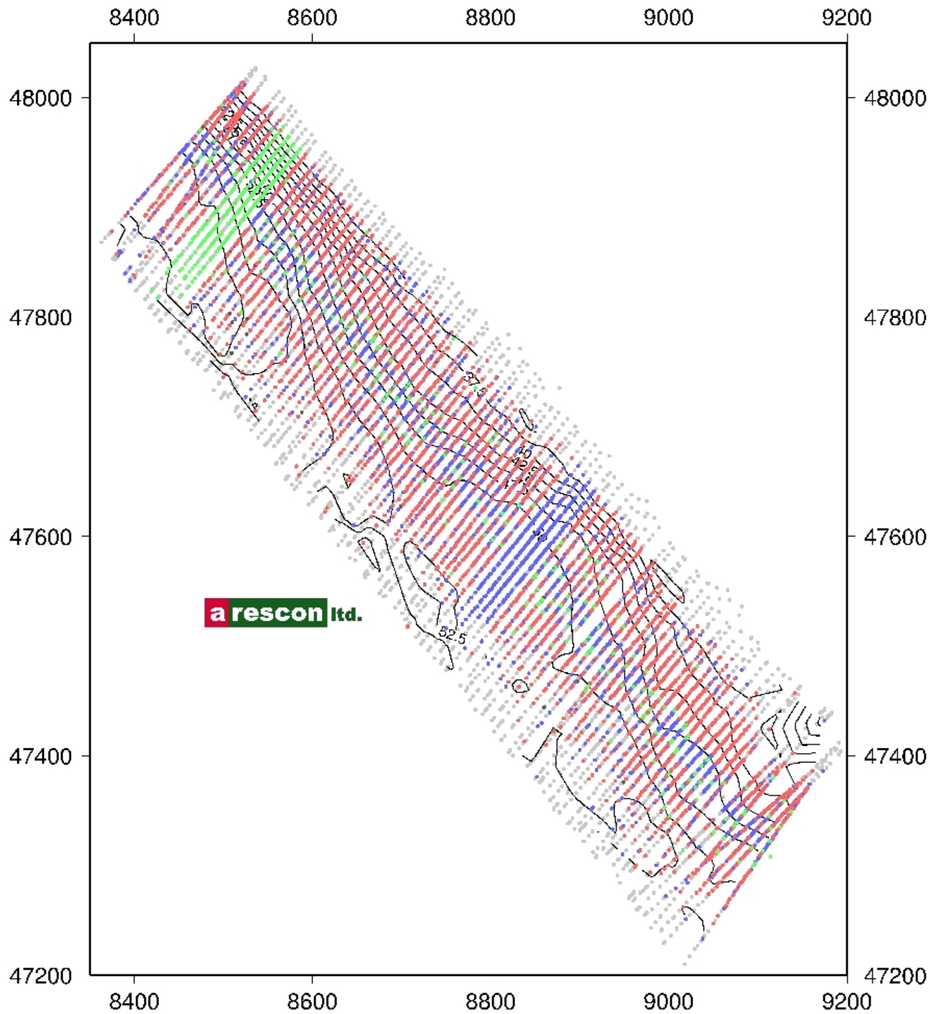

Figure 6: The line after all beams have been classified. The training sets are classified correctly. Any remaining black dots designate "no class", their confidence level was to small to warrant a classification.

contact@arescon.com
www.arescon.com

The highest confidence values (biggest circles) occur within the training sets. This just confirms the earlier conclusion, that the network is well suited to recognize patterns it has been trained with. Outside the training sets the result is less conclusive. Obviously class “green” and class “blue” have very similar properties. Beam by beam classification as shown in figure 6 results in interspersed green and blue dots. When confidence values and neighboring beams are considered as in figure 7, class “blue” wins over “green” almost anywhere outside the training sets, due to higher confidence values. The filtered result in figure 7 now also shows a clear preference of “blue” for nadir and near-nadir beams, while “red” dominates the beams with greater incident angles. The “red” training set did not contain any nadir exemplars and it is rather a surprise that in the north-west and south-east parts of the line, some of the nadir beams nevertheless classify “red”.

However, the general behavior, that, with statistically very in-homogenous, arbitrary classes, which most likely overlap, the system locks on to whatever salient differences remain, is rather promising.

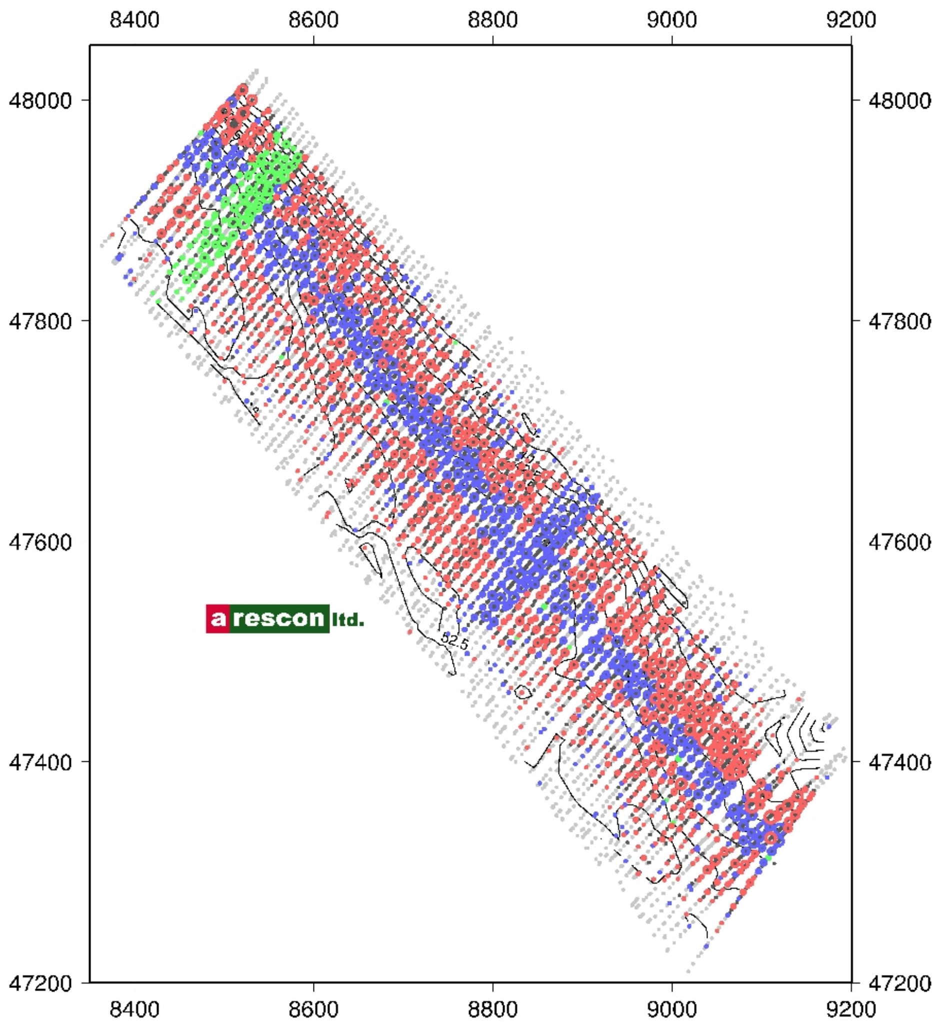

Figure 7: Spatially "aware" classification. Neighbors have a vote, weighted by their own confidence value, to decide a beams class membership. Circle size represents confidence level. Class "green" looses competition with class "blue" outside the training set.