BRIEF COMMUNICATION

Automatic oestrus detection from milking data - a preliminary investigation

R. A. SHERLOCK[1], L. A. SMITH[2] AND R. S. MITCHELL[2]

[1] Dairying Research Corporation Limited, Private Bag 3123, Hamilton, NZ

[2] Department of Computer Science, University of Waikato, Private Bag 3105, Hamilton, NZ.

Keywords: **Machine learning, oestrus detection, dairy cow.**

Efficient oestrus detection is important economically, particularly with seasonal-calving herds employing artificial insemination, as every missed oestrus effectively costs 21 days milk production for that cow in that season which is about 8% of the total. Traditionally, oestrus detection relies mainly on visual observation of animal behaviour - a cow in oestrus will stand and allow herself to be mounted by herd-mates. Such events are readily noted when they take place amongst assembled cows, such as at milking times, but may be missed entirely in herds of free-grazing animals which are only brought in for milking twice a day, as is usual in New Zealand style systems. A common low-cost aid to oestrus detection is the use of "tail-paint" (Macmillan and Curnow, 1977), and although this can be very effective when properly executed the regular visual inspection and repainting is labour-intensive.

With the introduction of milking systems that automatically record individual animal production (and possibly other variables) at every milking, there is potentially an opportunity to use these data to assist in identifying oestrus. It is well known by farmers that some cows exhibit a characteristic pattern at oestrus of a significantly reduced milk volume followed by a compensatory higher volume at the following milking. Also, in some cases, behavioural changes occur in that cows which usually have a well-defined position in the milking order will present themselves for milking significantly out of that sequence. Such effects are only clearly visible in the raw data for between 5 and 10% of animals, but this gives some grounds for hope that a machine recognition method may be able to increase this figure to a more useful value; say > 50%. Such a method could be run automatically on a suitable herd database, thus not involving additional farmer time. Any resulting increase in the success rate for identifying cows in oestrus would make it a useful adjunct to traditional techniques.

The application of two established machine learning methods, C4.5 (Quinlan, 1992) and FOIL (Quinlan, 1990), to this task was investigated. Both are examples of “similarity-based” learning schemes which are “trained” using a set of example data that have been partitioned into two or more classes - oestrus and non-oestrus in our case.

The C4.5 scheme is a system for inducing production rules and decision trees using a “top-down” or “divide-and-conquer” approach. The set of training examples is first partitioned into two or more subsets on the outcome of a test based on the value of a single attribute; the particular test being chosen by an information-theoretic heuristic that generally gives a nearly optimal partitioning (Quinlan,1992). This procedure is repeated on each of the new subsets and continues until either a subset contains only examples of a single class or the partitioning tree has reached a predetermined maximum depth. The tree can then be used to classify a test case by starting at the top node of the tree and following the branches corresponding to the attribute values until a leaf node is reached. C4.5 incorporates a criterion for determining the best partitioning of the examples at each decision tree node, “prunes” the decision tree to reduce the chance of overfitting the data, and has the ability to derive production rules from the unpruned decision tree. These derived rules are comparable in accuracy to the pruned decision tree, but are more easily interpreted by people.

In contrast to C4.5, FOIL is a “bottom-up” classifier: rules are induced for each class, one class at a time, by dividing the training examples into “positive” - those from the class in question, and “negative”- those from all other classes. FOIL then attempts to find a set of logical clauses that cover or match some positive examples but no negative examples. The matching positive examples are removed, and the process repeated until no positive examples remain. A similar procedure is followed for each class. The end result is a set of relational rules for each class in the data. An important feature of FOIL is its ability to express relationships between the attributes in an example whereas so-called “zero-order classifiers” such as C4.5

can only compare attribute values with constant numbers or symbols, not other attributes. We expected that this feature would be useful when classifying examples containing time-series data.

These particular methods were chosen because they are well understood among machine learning researchers and are also readily available as part of the WEKA (Waikato Environment for Knowledge Analysis) machine learning workbench (McQueen *et al*, 1994). WEKA integrates a wide range of machine learning algorithms and support tools into a single interactive package, allowing data to be analysed using different learning systems and the results to be evaluated in a consistent manner.

Milking performance data was collected from a Dairying Research Corporation herd of 130 cows over the 1993–94 milking season using PC-based MilkMAID software polling a Ruakura Milk Harvester machine (Sherlock and Woolford, 1992). Attributes for the machine learning algorithm were chosen to minimise effects of global differences between animals and between successive milkings. For example, all milking variables were expressed as a percentage of herd means to eliminate most of the effects of paddock and climatic variation. Several attribute schemes were investigated; the most successful (on which the Table 1 results are based) expressed milk yields in 10-milking time series segments (TSS's) of normalised differences from that cow's running mean. The normalisation expressed the differences as a percentage of the running variance (Mitchell *et al.*, 1996) in an attempt to provide the learner with a consistent level of statistical significance in the data from different animals. Other attributes were a similarly normalised milking-order TSS, and also the time since last oestrus. Only if the last milking in the TSS coincided with a behaviourally observed oestrus recorded in the MilkMAID database was that example classified as an oestrus event.

[TABLE 1 IN HERE]

The results presented in Table 1 are averages over 10 learning runs, each using a different split of the database into training (66%) and test (33%) sets. The "overall correct" classification is not a very useful measure; the high values arise because true oestrus events make up only around 5% of the data - a "naive learner" could achieve 95% overall accuracy by classifying all events as negative! Although the average values in Table 1 are not directly applicable to a single case, an illustrative example of the performance of the best scheme (C4.5 rules) is provided by considering a milking of a 500 cow herd at which there were 25 cows in oestrus. Typically 17 (69% of 25) correct identifications would be made, and hence 8 true oestrus cases would be missed. In addition, around 48 cows would be incorrectly identified as being in oestrus, giving 48/(48+17) = 74% false positives. Of these deficiencies the missed cases are of greater econmic significance than the false positives, although a reduction of both is clearly desirable.

While this performance is definitely not yet at a level suitable for a practical application, it is encouraging that C4.5 does much better than human “experts” who would be unlikely to identify more than 10% of true positives from the same data. It should be possible to improve machine learning accuracy by better training, in particular by replacing the error-prone human observation of oestrus with an absolute physical measure such as milk progesterone level. Reducing the amount of noise in the data - by increasing the measurement accuracy of the milking machinery - would certainly help, although there is no simple way of quantifying the resultant improvement in classification accuracy. Additional automatically measured variables, such as the electrical conductivity of the milk, may also be included to see if they assist the classification. Further study is required to assess the relative contributions of the different variables to the learning process - an analysis of the rules and decision trees produced by the learning schemes would be useful in this respect. Further work is also needed to investigate the nature of the distribution of cow performance variations around oestrus events - we may have to abandon hope of a single learning strategy that can be applied to the entire herd. There is almost certainly not enough data to learn individual cow behaviours, but there may be a possibility of grouping together animals with similar behaviours.

In summary, we consider that Machine Learning systems have the potential to become a useful practical aid to oestrus detection, and emphasise that on-farm implementations would require no additional labour (all of the necessary data being collected and analysed automatically) or running costs, given the presence of a suitable milking system.

REFERENCES

Macmillan, K.L. and Curnow, R.J. (1977) Tail painting: A simple form of oestrus detection in New Zealand dairy herds. New Zealand Journal of Experimental Agriculture, 5: 357-361.

McQueen, R.J., Garner, S.R., Nevill-Manning, C.G. and Witten, I.H. (1994) Applying machine learning to agricultural data. Computers and Electronics in Agriculture, 12: 275–293.

Mitchell, R.S., Sherlock, R.A. and Smith, L.A. (1996) An investigation into the use of machine learning for determining oestrus in cows. Computers and Electronics in Agriculture, in press.

Quinlan, J.R. (1990) Learning logical relations from definitions. Machine Learning, 5: 239–266.

Quinlan, J.R. (1992) C4.5: Programs for Machine Learning. Morgan Kaufmann Publishers, Inc., San Mateo, CA.

Sherlock, R.A. and Woolford, M.W. (1992) Automated cow and machine performance monitoring in the Ruakura Milk Harvester. In: A.H. Ipema, A.C. Lippus, J.H.M. Metz and W. Rossing (Editors), Proceedings of the International Symposium on Prospects for Automatic Milking, November 1992, Wageningen, The Netherlands. Pudoc Scientific Publishers.

TABLE 1. Summary of classification results

Dataset / Algorithm	**Overall correct %**	**% Oestrus identified**	**% False positives**
Training set:			
FOIL rules	99.6	94.9	0.0
C4.5 pruned trees	91.6	64.2	35.7
C4.5 production rules	85.3	66.3	48.8
Test set:			
FOIL rules	96.3	20.0	68.0
C4.5 pruned trees	87.8	48.8	76.9
C4.5 production rules	82.2	68.7	73.7