

December 1998

User Perceptions of Machine Learning

Robert McQueen
University of Waikato

Geoffrey Holmes
University of Waikato

Follow this and additional works at: <http://aisel.aisnet.org/amcis1998>

Recommended Citation

McQueen, Robert and Holmes, Geoffrey, "User Perceptions of Machine Learning" (1998). *AMCIS 1998 Proceedings*. 63.
<http://aisel.aisnet.org/amcis1998/63>

This material is brought to you by the Americas Conference on Information Systems (AMCIS) at AIS Electronic Library (AISEL). It has been accepted for inclusion in AMCIS 1998 Proceedings by an authorized administrator of AIS Electronic Library (AISEL). For more information, please contact elibrary@aisnet.org.

User Perceptions of Machine Learning

Robert J. McQueen
Geoffrey Holmes
University of Waikato

Abstract

Machine learning has potential use in the understanding of information hidden in large datasets, but little is known about user's perceptions about the use of the technology. In this study, a number of datasets were solicited from agricultural researchers and processed using a machine learning workbench. The results were reported to the researchers, and then interviews were conducted with some of them to determine their perceptions about the use of machine learning as an additional analysis technique to traditional statistical analysis. A number of themes about their satisfaction with this technique were constructed from the interview transcripts, which generally indicate that machine learning may be able to contribute to analysis and understanding of these kinds of datasets.

Introduction

Machine learning is an information analysis tool that has potential for aiding the understanding of data collected from research experiments, but is seldom used in this environment in comparison to traditionally accepted methods of statistical analysis. In part, this may be a result of a low awareness of the characteristics of machine learning by these scientists, as well as the lack of a strong motivation to move away from the traditional, and well accepted statistical techniques. Successful fielded implementations of rules and expert systems derived from machine learning rule induction have been reported (Langley and Simon, 1996), but often only successful examples are given. There seems to be no parallel evidence on machine learning systems that have been attempted, and discarded, or the reasons for their success or failure.

For the purposes of this study we define machine learning as the acquisition of structural descriptions from examples, and consider the most commonly performed task in machine learning—classification. The examples comprise a set of attributes and a classification. Structural descriptions are “learned” from tables containing attributes and classifications in a variety of ways. The oldest and arguably the most popular are decision trees (Quinlan, 1986) but other forms of representation such as rule sets are also possible. A decision tree is a set of nodes and branches. Each node is labelled with an attribute name and each branch leading out is labelled with one of the possible values of that attribute. Leaves are labelled with values of the classification attribute. These trees are similar to classification and regression trees which are used in the statistical community (Breiman et al, 1984), and many statistical techniques are commonly incorporated into machine learning algorithms, but in practice, machine learning is still a relatively unknown data analysis technique, especially among research scientists.

This study reports on the user satisfaction of scientists with machine learning analysis of their experimental data in the domain of agricultural research. The WEKA (Waikato Environment for Knowledge Analysis) machine learning workbench was used for this study, and provides an integrated environment which gives easy access to a variety of machine learning techniques through an interactive interface (Holmes, Donkin and Witten, 1994), as well as a variety of software tools to support other aspects of data analysis such as attribute editing and data visualisation.

Methodology

About 100 New Zealand agricultural research scientists were given a short explanation of machine learning, and asked whether they had any datasets of past research work that they would be willing to provide for analysis. About 30 of these responded positively, and 14 useable datasets were eventually received. The number of attributes per instance in these datasets ranged from 6 to 32, and the number of instances varied from 19 to 21,448. The machine learning technical staff who prepared and processed the datasets through the workbench had little domain knowledge about the meaning and significance of the dataset attributes (other than the attribute name) and no knowledge about expected results. A short report was prepared containing the results and decision trees and was posted back to the researcher. A summary of the results was compiled (Thomson & McQueen, 1996).

Interviews were arranged with eight of these researchers, selected mainly through availability and interest in the study, and conducted either face-to-face or by telephone. The interviews were recorded and transcribed. The interviews were semi-structured, with a set of questions asking about the nature of the original investigation, original expected results, perceptions about the machine learning analysis results presented to them, initial and present perceptions about machine learning, and perceptions about machine learning as a possible future analytical tool for their research.

Perceptions about Machine Learning from the Interviews

A number of themes about the effectiveness of machine learning were developed from the interview transcripts, and a few exemplars are reproduced here. It is not intended that these themes can be generalised and applied outside of the boundaries of this study, but they may provide some insight into what arm's length users of machine learning perceive to be strengths and weaknesses of the technology. These themes may also prove useful to providing a starting point for similar studies in other application domains.

For machine learning, there may be a high return from a small amount of time put into preparing the data for machine learning analysis, especially when compared to the large amount of time necessary to prepare data for conventional statistical analysis

There is often a large front end time commitment in preparing experimental data for statistical processing. Part of this time may be in narrowing the statistical analysis to a small area of potential interest, running the analysis, and then shifting to the next area of interest. Machine learning can usually take all of the data in one pass, and produce results, with relatively little front end preparation and selection. The issue then becomes the increase of effective understanding of the data resulting from statistical or machine learning analysis, versus the time spent preparing for that analysis. One interviewee put it as follows:

Well, we didn't really expect [machine learning] to actually match so closely with what we had done with our statistical analysis because we'd put a lot of time into [the statistical analysis while] the machine learning was a very minor input on our behalf So in terms of effective use of time [machine learning] was really very very effective ... we hadn't expected it to link so closely with what we came up with [from the statistical analysis].

Machine learning output (decision trees) can help a researcher visualise the “shape” of the data.

The output of machine learning schemes often take the form of decision trees, graphs or rule lists, which can help give a visual and mental structure to the data being processed. This can be especially helpful with large data sets. One interviewee commented:

if you have a large data set I think the machine learning has got a lot of potential as a first cut and maybe its got potential further than that ...

Machine learning may not find obvious strong relationships in complex data

One of the interviewees perceived that machine learning had not found some of the relationships in the data that she had found from the previously done statistical analysis.

Because machine learning is not broadly used, publication of research results might be difficult for other researchers to understand

As the interviewees were research scientists, publication of results in peer reviewed journals and conferences is an important part of their activities. There was concern expressed that publication of machine learning based results would be more difficult, and might have reduced status, because of the traditional convention of using statistical analysis to present research results.

Some researchers would be interested in using machine learning in future studies, provided they could learn more about how it should be used.

This first experience with machine learning had generated interest in using machine learning in subsequent studies. When asked whether he would be interested in using machine learning again, most of the interviewees said they would like to try it again. One said...

Well, we're really keen to use it again. We would consider designing research [to make use of machine learning]. I think it could be more appropriate than more traditional techniques for some particular purposes so if we learn more about it we may well definitely consider doing that.

Conclusions

This paper has examined the perceptions of agricultural research scientists who participated in a study on the application of machine learning analysis techniques to datasets from previously undertaken research experiments. There were both positive and negative perceptions about how effective machine learning would be in analysing experimental data, and presenting the outcomes in presentations and publications. This study involved “arms length” interaction between the machine learning technician and the research scientist, and future studies of this kind might be directed toward more closely tied, interactive collaborations between machine learning expert and research domain expert to see if some of the shortcomings identified here may be overcome.

Acknowledgements

This work was supported in part by the New Zealand Foundation for Research, Science and Technology.

References

- Breiman, L., Friedman, J.H., Olshen, R.A., and Stone, C.J. *Classification and Regression trees*, Wadsworth Press, 1984
- Holmes, G., Donkin, A. R., and Witten, I.H. (1994) "WEKA: A Machine Learning Workbench." *Tech Report 94/9*, Department of Computer Science, University of Waikato, New Zealand, 1994.
- Langley, P. And Simon, H.A.. "Applications of machine learning and rule induction", *Communications of the ACM*, V. 38, No. 11, 1995, pp. 55-64.
- Quinlan, J.R. "Induction of decision trees." *Machine Learning 1*, 1986, pp 81-106.
- Thomson, K. and McQueen, R.J. "Machine learning applied to fourteen agricultural datasets" Working paper 96/18, Dept of Computer Science, University of Waikato, 1996. Available at <http://www.cs.waikato.ac.nz/~ml/publications.html>