



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Roadknight, C;Rattadilok, P;Aickelin, U

Title:

Teaching Key Machine Learning Principles Using Anti-learning Datasets

Date:

2019-01-16

Citation:

Roadknight, C., Rattadilok, P. & Aickelin, U. (2019). Teaching Key Machine Learning Principles Using Anti-learning Datasets. Lee, MJW (Ed.) Nikolic, S (Ed.) Shen, J (Ed.) Lei, LCU (Ed.) Wong, GWK (Ed.) Venkatarayalu, N (Ed.) Proceedings of 2018 IEEE International Conference on Teaching, Assessment, and Learning for Engineering, TALE 2018, 00, pp.960-964. IEEE. <https://doi.org/10.1109/TALE.2018.8615252>.

Persistent Link:

<https://hdl.handle.net/11343/239237>

Teaching Key Machine Learning Principles Using Anti-Learning Datasets

Chris Roadknight
School of Computing
University of Nottingham
Ningbo, China

chris.roadknight@nottingham.edu.cn

Prapa Rattadilok
School of Computing
University of Nottingham
Ningbo, China

prapa.rattadilok@nottingham.edu.cn

Uwe Aickelin
School of Computing
University of Melbourne
Melbourne, Australia
uwe.aickelin@unimelb.edu.au

Abstract—Much of the teaching of machine learning focuses on iterative hill-climbing approaches and the use of local knowledge to gain information leading to local or global maxima. In this paper we advocate the teaching of alternative methods of generalising to the best possible solution, including a method called anti-learning. By using simple teaching methods students can achieve a deeper understanding of the importance of validation on data excluded from the training process and that each problem requires its own methods to solve. We also exemplify the requirement to train a model using sufficient data by showing that different granularities of cross-validation can yield very different results.

Keywords—Machine Learning, Anti-Learning, exclusive-or, Hadamard

I. INTRODUCTION

Much of the teaching of machine learning revolves around using local knowledge, in N dimensional space, to iterate towards a global solution. This relies upon the premise that there will be a fitness ‘hill’ to climb on the way to the fitness ‘peak’ [4]. There are many examples of algorithms reaching local minima and methods to reset from or avoid these fitness sub-peaks [eg. 5], but overall there is an assumption that for supervised, semi-supervised and unsupervised learning methodologies that use local information will eventually reach the best possible solution and therefore provide the best possible model of the data.

There exists a range of problems where the use of local information to iteratively improve the accuracy of a model fail entirely. In fact, they fail so badly that the resulting model predicts categories in an unseen validation set at a rate much worse than guessing, this is called Anti-Learning [6]. Data that exhibits this kind of behavior comes from a range of sources both real and synthetic. Problems in the areas of cancer biomarker modelling, yeast genomics and chemical regulation are all real-world examples where there exists anti-learnable datasets. Synthetic data exhibiting this phenomena includes some traditional matrices and various versions of exclusive-OR data. This paper explores why this phenomena exists and why it is a crucial part of any machine learning curriculum, given it elegantly explains some key features such as the requirement for validation and the use of high granularity cross-validation.

II. ANTI-LEARNING

The meaningfulness of nearest neighbour style clustering in highly dimensional data has been discussed previously [2] and it has been argued that there are serious problems with using this approach alone [3], but this is exactly what happens in most undergraduate modules on machine learning.

Anti-learning is the situation when, no matter how hard you optimise your algorithms, the performance of your machine learnt model, on data not used for its training, is reproducibly worse than the probability of guessing the answer. So for a 2 class problem, if your model is consistently less than 50% accurate on unseen data, the underlying model needs to be anti-learned. Most students, and even staff, are baffled by this scenario but its explanation ironically lies in one of the simplest, standard problems used to teach the need for multilayer perceptrons [1].

The exclusive-or problem is a trivial problem that is used by many teachers to explain why linear modelling solutions are insufficient to model non-linear problems [7]. It says that when inputs two inputs are both true or both false then the output should be false, while when the first and second input are different, the output should be true [17]. Figure 1 displays how the hyperplane needs to be separated (from [17])




Fig. 1. Visual representation of exclusive-or

This problem requires a curved hyperplane to be generalised to accurately represent the data, and as such is a non-linear problem. The main stumbling block here though is that ALL of the dataset is required to build the model, so therefore there will be no data left to test the model. If you leave one value out of the training set, the model built will ALWAYS give the incorrect answer for these excluded value. For the Exclusive-Or problem where we know we have all available data, modelling using all the data is acceptable, if trivial but for most problems we only model the existing data in an effort to be able to approximate or predict outcomes for unrepresented or future scenarios, so therefore to test if we have generalized we need to assess performance on data not used to train the model.

III. DATA AND TOOLS

Three datasets are used in this work.

1. Hadamard matrix
2. An pyramidically aggregated Exclusive-Or dataset
3. A randomly aggregated Exclusive-Or dataset

A. Hadamard matrix

A Hadamard matrix is a square matrix where each entry is either 1 or -1, and each pair of rows/columns has matching entries in exactly half of their columns and mismatched entries in the remaining columns/rows. The following is an example for hadamard matrices of size 2 and 4.

$$H_2 = \begin{matrix} 1, 1 \\ 1, -1 \end{matrix}$$

$$H_4 = \begin{matrix} 1, 1, 1, 1 \\ 1, -1, 1, -1 \\ 1, 1, -1, -1 \\ 1, -1, -1, 1 \end{matrix}$$

Larger scale hadamard matrices can be displayed visually (fig 2).




Fig. 2. Visual representation of a Hadamard Matrix

For the purposes of this work, we will use the proceeding $n-1$ values for each row as variables by which machine learning methods may attempt to learn dependent variable n (the final variable in the row).

For the purposes of teaching, this dataset is easily generated using a number of tools. For instance the function `Hadamard(n)` in the `phangorn` package [13] produces a hadamard square matrix of size n .

B. Pyramidically aggregated Exclusive-Or

A larger version of the standard exclusive-or logic can be generated by aggregating multiple XOR variable pairs. In this work we take 8 binary variables, $a-h$ and calculate the XOR of the sequential 4 pairs, and then repeat this with the results. This is shown in equation 1.

$$y = \text{xor}(\text{xor}(\text{xor}(a,b), \text{xor}(c,d)), \text{xor}(\text{xor}(e,f), \text{xor}(g,h))) \quad (1)$$

This gives 256 training patterns that are exactly 50% True and 50% False.

Students can generate this dataset themselves using a variety of tools, including excel. The first few rows are shown below in table 1:

a	b	c	d	e	f	g	h	out
0	0	0	0	0	0	0	0	FALSE
0	0	0	0	0	0	0	1	TRUE
0	0	0	0	0	0	1	0	TRUE
0	0	0	0	0	0	1	1	FALSE
0	0	0	0	0	1	0	0	TRUE
0	0	0	0	0	1	1	0	FALSE
0	0	0	0	0	1	1	1	FALSE
0	0	0	0	0	1	1	1	TRUE

Table 1. First 8 rows of the aggregated ‘pyramid’ exclusive-or dataset

C. Randomly aggregated Exclusive-Or

A more unbalanced version of aggregated XOR can be achieved by using a random variable structure. This has the benefit of allowing variables to be used multiple (or zero) times and is therefore a more heterogeneous structure. The random allocation we used is shown in equation 2.

$$y = \text{xor}(\text{xor}(\text{xor}(d,g), \text{xor}(a,d)), \text{xor}(\text{xor}(h,f), \text{xor}(d,b))) \quad (2)$$

This gives 256 training patterns that are exactly 50% True and 50% False. Again this dataset is trivial to produce using excel.

D. Machine Learning

A multitude of off the shelf machine learning tools are available such as H2o [9], the caret [10] or e1071 [11] packages in R and WEKA [12]. A large range of algorithms have been used for anti-learning in the past [8] but here we just use Naïve Bayes, Support Vector Machine and a Multi-layer Perceptron.

IV. RESULTS

We carried out performance tests on the three datasets, using three different machine learning algorithms and using 7 different levels of cross-validation granularity (4,8,16,32,64,128,256 folds). 256 fold cross-validation is equivalent to ‘leave-one-out’ (LOOCV) cross-validation. Figures 2, 3 and 4 show predictive accuracy on validation data. It is important to note that in all cases predictive accuracy on training data was always above 50% (usually above 75%).

The Naïve Bayes [16] approach (Figure 3) never predicts validation outcome at greater than 50% for any of the datasets at any of the cross-validation granularity. Moreover, at LOOCV it achieves 0% accuracy on all three datasets. Therefore by inverting the outcome Naïve Bayes can predict outcome on unseen data at 100% accuracy. This anti-learning result is better than any learning approach. The Multi-layer perceptron approach shown in figure 4 does achieve a high (>90%) accuracy on unseen data at low cross-validation granularity, but this falls swiftly to well below 50% once more cross-validation folds are introduced. This is a key result whereby the importance of cross-validation methodology is highlighted. The support vector machine [15] approach similarly tends towards an accuracy of well below 50% (figure 5) as the number of folds increases, though in this case the

accuracy never goes above 50%, so is always in the anti-learning state.




Fig. 3. Performace of a naïve Bayes approach on the three synthetic datasets for different granularities of cross-validation




Fig. 4. Performace of a multi-layer perceptron approach on the three synthetic datasets for different granularities of cross-validation




Fig. 5. Performace of a support vector machine approach on the three synthetic datasets for different granularities of cross-validation

V. DISCUSSION

So why should anti-learning be included in undergraduate modules in machine learning. Most modules include supervised, unsupervised and semi-supervised learning

paradigms in their curriculum. In this paper we suggest that anti-learning should be added to this. The reasons are as follows:

1. A small subset of real world and synthetic data are best modelled using an approach that exploits the anti-learning phenomena [8]. These real world problems are of high significance and complexity.
2. It elegantly teaches the importance of focusing on the performance of models on data that was excluded from the training process. Building a model that accurately reflects the data it is trained with is trivial but as is shown in this paper, largely meaningless if performance on unseen data (eg. future data) is a priority.
3. It highlights the importance that should be given to the validation process, in particular the degree of cross-validation. If it is not computationally impossible, 'leave one out cross-validation' is the safest approach, using the largest amount of data to train the models while still validating across the entire dataset. If the LOOCV results are very different to results using fewer folds, it must be suspected that an insufficient number of samples were used.
4. It teaches pragmatism. Some anti-learning datasets are highly dimensional with insufficient samples to represent the feature space completely but the data is still important and a timely model for subjects such as cancer prognosis is imperative. The best model turns out to be an anti-learned model in some instances [14]. It may be possible to revisit the models at a later date, when more data is available.
5. It exemplifies the differences shown but different modelling techniques on the same dataset. In some instances choosing the most sophisticated, tunable modelling tool may not always be the best solution. Naïve Bayes is one of the oldest and simplest machine learning methods yet yields the most accurate anti-learning models for all three datasets.

Overall, introducing datasets that exhibit the anti-learning phenomena, whether it be because they are highly non-linear or they are highly dimensional, can only improve a student's understanding of appropriate methods of tackling all kinds of datasets.

REFERENCES

- [1] P. Sankar K., and S Mitra. "Multilayer Perceptron, Fuzzy Sets, Classification." (1992).
- [2] C. Roadknight, U. Aickelin, A. Ladas, D. Soria, J. Scholefield, and L. Durrant. "Biomarker clustering of colorectal cancer data to complement clinical classification." In *Computer Science and Information Systems (FedCSIS), 2012 Federated Conference on*, pp.187-191. IEEE, 2012.
- [3] K. Beyer, J. Goldstein, R. Ramakrishnan, U. Shaft, When is nearest neighbor meaningful? in: *Proc. Int. Conf. Database Theory*, 1999, pp. 217-235.
- [4] H. Kaizhu, H. Yang, I. King, and M. Lyu. "Global Learning vs. Local Learning." *Machine Learning: Modeling Data Locally and Globally* (2008): 13-27.
- [5] D. Gorse, A. Shepherd, and J. G. Taylor. "A classical algorithm for avoiding local minima." In *Proceedings of the World Congress on Neural Networks*, pp. 364-369. 1994.
- [6] A. Kowalczyk, and O. Chapelle. "An analysis of the anti-learning phenomenon for the class symmetric polyhedron." In *International Conference on Algorithmic Learning Theory*, pp. 78-91. Springer, Berlin, Heidelberg, 2005.
- [7] Z. Yanling, D. Bimin, and W. Zhanrong. "Analysis and study of perceptron to solve XOR problem." In *Autonomous Decentralized System, 2002. The 2nd International Workshop on*, pp. 168-173. IEEE, 2002.
- [8] A. Kowalczyk. "Classification of anti-learnable biological and synthetic data." In *European Conference on Principles of Data Mining and Knowledge Discovery*, pp. 176-187. Springer, Berlin, Heidelberg, 2007.
- [9] H2O. <http://h2o.ai/>
- [10] Caret. <https://cran.r-project.org/web/packages/caret/vignettes/caret.html>
- [11] e1071. <https://cran.r-project.org/web/packages/e1071/e1071.pdf>
- [12] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. "The WEKA data mining software: an update." *ACM SIGKDD explorations newsletter* 11, no. 1 (2009): 10-18.
- [13] Phangorn. <https://cran.r-project.org/web/packages/phangorn/index.html>
- [14] C. Roadknight, U. Aickelin, G. Qiu, J. Scholefield, and L. Durrant. "Supervised learning and anti-learning of colorectal cancer classes and survival rates from cellular biology parameters." In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, pp. 797-802. IEEE, 2012.
- [15] J. Platt. "Fast training of support vector machines using sequential minimal optimization." *Advances in kernel methods* (1999): 185-208.
- [16] G. John, and P. Langley. "Estimating continuous distributions in Bayesian classifiers." In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pp. 338-345. Morgan Kaufmann Publishers Inc., 1995.
- [17] M. Lewin. "All about XOR". In *Overload Journal #109 - June 2012. Programming Topics*. 2012. <https://accu.org/index.php/journals/1915>