

Parameterizing Random Test Data According to Equivalence Classes

Christian Murphy
Dept. of Computer Science
Columbia University
New York, NY
cmurphy@cs.columbia.edu

Gail Kaiser
Dept. of Computer Science
Columbia University
New York, NY
kaiser@cs.columbia.edu

Marta Arias
Center for Computational Learning
Systems
Columbia University
New York, NY
marta@ccls.columbia.edu

ABSTRACT

We are concerned with the problem of detecting bugs in machine learning applications. In the absence of sufficient real-world data, creating suitably large data sets for testing can be a difficult task. To address this problem, we have developed an approach to creating data sets called “parameterized random data generation”. Our data generation framework allows us to isolate or combine different equivalence classes as desired, and then randomly generate large data sets using the properties of those equivalence classes as parameters. This allows us to take advantage of randomness but still have control over test case selection at the system testing level. We present our findings from using the approach to test two different machine learning ranking applications.

Categories and Subject Descriptors

D.2.4 [Software Engineering]: Software/Program Verification.
D.2.5 [Software Engineering]: Testing and Debugging.

General Terms: Reliability, Verification.

Keywords: Software testing, Random test data generation.

1. INTRODUCTION

During our investigation of the software testing of machine learning (ML) applications, we immediately encountered the problem of the availability of testing data. Real-world data sets are not always accessible and, even in the cases when they are, may not necessarily contain all the equivalence classes that proper testing demands. Hand-generation of data is an option but is only useful for small tests; however, these applications are used in the real world on extremely large data sets. Random testing [5, 7] could, of course, be an option in these situations, but is difficult given that machine learning applications fall under the category of “non-testable programs” [21]; that is, there is no reliable test oracle to ensure that the output is correct for the given random input. In fact, Hamlet [7] even points out that “Random testing cannot be attempted without an effective oracle. A vast number of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RT’07, November 6, 2007, Atlanta, GA, USA
Copyright 2007 ACM ISBN 978-1-59593-881-7/07/11...\$5.00

test points are required, and they cannot be trivialized to make things easier for a human oracle.” Thus pure random testing may not be the best approach in these cases, either.

Our approach, then, has been to combine partition testing with what we call “parameterized random test data generation”. In order to obtain data sets that provide the different combinations of equivalence classes, or the desired separation and isolation of equivalence classes, it is necessary to automatically generate random data sets, but parameterized to control the range and characteristics of those random values. We present a hybrid testing approach that couples the benefits of using randomness with the necessary control over the properties of the testing data.

In this paper, we describe the manner in which we generate data for testing ML applications that implement ranking algorithms (a requirement of a real-world problem domain). We have developed a tool that, given certain parameters, will create randomly-generated data sets that are targeted at testing specific equivalence classes. We use randomness in the test data generation to reveal bugs and inconsistencies in the ML applications that could not otherwise have been shown by real-world or hand-crafted data.

2. BACKGROUND

We are concerned with testing a decision support system that uses ranking of susceptibility to failure of some large number of electrical devices in a real-world industrial environment where preventative maintenance can potentially be guided by such susceptibility. Classification in the binary sense (“will fail” vs. “will not fail”) is not sufficient because, after enough time, every device will eventually fail. Instead, a ranking of the propensity of failure with respect to all other devices is more useful. The prototype application uses implementations of both the MartiRank [12] and SVM [19] algorithms to produce rankings; the dependability of the application has real-world industrial implications, rather than just academic interest. Our focus is only on the ML component of this application. We do not discuss the full application further in this paper; see [6] and [1] for details.

2.1 Data Sets in the Problem Domain

In general, data sets used in ML (both ranking and classification) consist of a collection of *examples*, each of which has a number of *attribute* values and a *label*. The examples can be thought of as rows in a table, each of which represents one item from which to learn (in our case, electrical devices). The attributes are the columns of the table, and the label indicates how the example is categorized; for us, it is an indication of how many times the device failed in a given time period. In some phases of supervised

ML learning there is a label, and in some there is none; without loss of generality, here we only discuss the former.

In previous work [14], we describe how we determined equivalence classes and performed test case selection for implementations of ML algorithms. The first part of our testing approach was to consider the problem domain and try to determine equivalence classes based on the properties of real-world data sets. We particularly looked for traits that may not have been considered by the algorithm designers, such as data set size, and the potential ranges (or even the existence) of values.

The data sets of interest are very large, both in terms of the number of attributes (hundreds) and the number of examples (tens of thousands). The label could be any non-negative integer, although it is typically a 0 (indicating that there was no device failure) or 1 (indicating that there was). As would be expected in any large data set, many attributes have repeated values; however, in the real-world data, many values may be missing for various reasons, raising the issues of breaking “ties” and handling unknowns during sorting.

Though much of the real-world data of interest consists of numerical values – including floating point decimals, dates and integers – some of the data are instead alphanumeric. ML ranking algorithms rely on sorting, and while in principle lexicographic sorts could be employed, non-numerical sorts do not seem intuitively appealing as ML predictors. To solve this problem, the data sets use categorical data. Categorical data refers to attributes in which there are K different distinct (non-numeric) values, but there is no sorting order that would be appropriate for the ranking algorithm. In these cases, a given attribute with K distinct values is expanded to K different attributes, each with two possible values: a 1 if the example has the corresponding attribute value, and a 0 if it does not. That is, amongst the K attributes, each example should have exactly one 1 and $K-1$ 0’s.

2.2 Machine Learning Ranking Algorithms

We investigated two ranking algorithms. The first, MartiRank [6], was specifically designed with the device failure application in mind. The algorithm looks for the combination of sorting and segmenting the data to try to achieve the best “quality”. Our testing concerned three implementations of the MartiRank algorithm, developed independently by three different developers: the original, written in Perl; a C version, written to improve performance (speed) and introduce some experimental options to try to improve quality; and another implementation also written in C, designed to minimize the costly overhead of repeatedly sorting the attribute values. The other algorithm, SVM [19], belongs to the “linear classifier” family of ML algorithms that attempt to find a hyperplane that separates examples from different classes. The goal is to find the hyperplane that maximizes the distance to its closest points. For our testing, we investigated SVM-Light [9].

2.3 Related Work

There has been investigation of applying machine learning techniques to software engineering in general and software testing in particular; however, our research does not address this. Rather, we are taking the opposite approach: we are applying software testing techniques to ML applications.

There has been much research into the generation of test data sets [4, 10], though whereas much of the early work in random test data generation [2, 8] started in the area of compilers, we are looking at a way of creating parameterized random test data specifically for ML algorithms. Of course, even purely random test data is somehow “parameterized”, but this generally refers to specifying the data type or range of acceptable values. The term “parameterized random testing” appears in circuit design literature [11] but refers to parameterizing the distribution of input values, and does not address parameterizing according to equivalence classes or partitions, which we present here.

Wichmann [22] proposes something similar to our approach in his recommendations to the British Computer Society Specialist Group in Software Testing. He notes the role that randomization can have even within the “limitations” of partition testing when it comes to randomly selecting testing data for a given equivalence class. However, his work in this area has only focused on software components and not system-level testing.

Our work is also similar to that of Thénevod-Fosse *et al.* [18], who labeled their approach “structural statistical testing” in that random input are selected according to given criteria, particularly related to path selection. Our approach differs, though, in that we are focused on the equivalence classes of the input data for black-box system testing, and not for coverage testing. Our work also differs from what they call “uniform statistical testing” because although we do select random data over a uniform distribution, we parameterize it according to equivalence classes.

Repositories of “reusable” ML real-world data sets have been collected (e.g., the UCI Machine Learning Repository [16]) for the purpose of comparing result quality, but not for testing in the software engineering sense. Of course, from an ML perspective, random test data is not generally as useful as real-world or hand-crafted data in testing the predictive power of an algorithm because it will not exhibit any trends, and therefore there is nothing for the algorithm to learn; however, random data can be used for testing the correctness of the results or for detecting bugs.

Lastly, Mayer *et al.* [13] have investigated the use of random testing with applications that have no test oracle. However, their work is focused on randomized software, which is different from the case of the ML applications we investigate, which are deterministic but, of course, the results cannot be known *a priori*.

3. DATA GENERATION FRAMEWORK

Although real-world data sets were available in abundance, we required more control over the data so that we could ensure that we addressed all equivalence classes. We thus created a tool that randomly generates values and puts them in the data set according to certain parameters. This allowed us to separately test different equivalence classes and then create a suite of tests that addressed those classes, focusing on boundaries. The parameters include the number of examples, the number of attributes, and the names of the output test data set files. Example data sets are available in [15], as is the algorithm for generating data.

The data generation tool can be run with a flag that ensures that no attribute values are repeated within the data set. This option was motivated by the need to run simple tests in which all values are different, so that sorting would necessarily be deterministic (no “ties”). It works as follows: for M attributes and N examples,

generate a list of integers from 1 to $M \cdot N$ and then randomly permute them. The numbers are then placed into the data set. If the flag is not used, then each value in the data set is simply a random integer between 1 and $M \cdot N$; the tool ensures at least one set of repeating numbers. The tool currently only generates positive integers.

The utility is also given the percentage of “positive examples” to include in the data set; positive examples have a label of 1, and negative examples have a label of 0. Similarly, a parameter specifies the percentage of missing values. Our data generation framework has been designed to guarantee that the number of positive examples and the number of missing values come out to be the right number, even though the values are randomly placed (or omitted) throughout the data set.

Parameters could be provided for generating categorical data (with K distinct values expanded to K attributes as described above). For creating categorical data, the input parameter to the data generation utility is of the format $(a_1, a_2, \dots, a_{K-1}, a_K, b)$, where a_1 through a_K represent the percentage distribution of those values for the categorical attribute, and b is the percentage of unknown values. The utility also allows for having multiple categorical attributes, or for having none at all.

Figure 1 shows two data sets created by the data generation tool, using the same parameters (10 examples, 10 attributes, 40% positive examples, 20% missing, repeats allowed); question marks indicate missing values, and the last number in each row is the label. Although the parameters for generating the data are the same, randomness comes into play in the placement of missing values, deciding which examples have labels of 1, etc.

27,81,88,59,?,16,88,?,41,?,0	35,3,20,41,91,?,32,11,43,?,1
15,70,91,41,?,3,?,?,?,64,0	19,50,11,57,36,94,?,96,7,23,1
82,?,51,47,?,4,1,99,?,51,0	24,36,36,79,78,33,34,?,32,?,0
22,72,11,?,96,24,44,92,?,11,1	?,15,?,19,65,80,17,78,43,?,0
57,77,?,86,89,77,61,76,96,98,1	40,31,89,50,83,55,25,?,?,45,1
76,11,4,51,43,?,79,21,28,?,0	52,?,?,?,?,39,79,82,94,?,0
6,33,?,?,52,63,94,75,8,26,0	86,45,?,?,74,68,13,66,42,56,0
77,36,91,?,47,3,85,71,35,45,1	?,53,91,23,11,?,47,61,79,8,0
?,17,15,2,90,70,?,7,41,42,0	77,11,34,44,92,?,63,62,51,51,1
8,58,42,41,74,87,68,68,1,15,1	21,1,70,14,16,40,63,94,69,83,0

Figure 1. Two data sets generated with the same parameters.

In addition to the data generation tool, we created a complementary utility that randomly permutes the order of the examples (rows) and attributes (columns) of the same data set.

4. FINDINGS

It is important to note that, in the general case, without a reliable test oracle, it is only possible to test for obvious and egregious errors, such as core dumps and runtime errors, but it is not possible to detect minor issues or even tell if the output is correct. However, because we had three implementations of the MartiRank algorithm, we were able to use them as “pseudo-oracles” [3] for each other to see if we were getting the correct output; this led to most of our more interesting findings. In the case of SVM, though, for which we only tested one implementation, we looked for consistency more than correctness. More details and case studies about the testing of MartiRank can be found in our tech report [15]; only the findings most related to parameterized random testing are presented here.

4.1 Testing with repeating values

Since any ML ranking algorithm will use sorting, our test case selection criteria led us to look at data sets with repeating values, to see how the software would deal with ties that may come from sorting examples with the same attribute values. In the real-world data sets, attributes like voltage level and activation date involve many repeating values. Our initial testing with small hand-crafted data sets showed that two of the MartiRank implementations were producing the same results, but a third was not. We then used the parameterized random data generation tool to automatically create larger data sets with repeating attribute values, and confirmed our intuition that the two implementations that agreed were using sorting routines that were “stable” (i.e., they maintain the relative order of the examples from the previous step when the values are the same), whereas the other was using a faster sorting algorithm that was not a stable sort. The MartiRank algorithm [12] does not address the specific implementation issue of which sorting approach to use, so different interpretations led to different results. This does not matter with respect to formal proofs, but does with respect to consistency.

4.2 Testing with sparse data sets

Our selection criteria then pointed us towards test data sets that had missing values. We used the data generation framework to create large, randomly-generated (but non-repeating) data sets, this time with the percent of missing values as a parameter. In these tests, all three MartiRank implementations were initially generating different results, and there was no way to know which was “correct” since the MartiRank algorithm does not dictate how to handle missing values. Consulting with the ML researchers, we decided that the sorting should be “stable” with respect to missing values in that examples with a missing attribute value should remain in the same position, with the other examples (with known values) sorted “around” them. Other deterministic options for handling this case (such as putting all missing values at the end) were considered, but this was deemed to be most in the MartiRank spirit since it kept the examples with missing values consistent with respect to any previous sort order.

4.3 Testing with categorical data

Because categorical data provides a combination of necessarily repeating (all 0s or 1s) and sometimes missing values, we created data sets with categorical attributes as part of our test data. We used a data set that included categorical data to discover a bug in the Perl implementation by the incorrect use of a global variable in a crucial calculation. More importantly, though, this bug did not surface when testing only with repeating values or only with missing values; it was the data sets that combined these two equivalence classes that allowed the bug to be revealed.

4.4 Testing effects of permuting input data

Among our findings with respect to the SVM implementation we tested, the most relevant is that randomly permuting the order of the examples in the training data caused it to generate different results. The practical implication is that the order in which the data happens to be assembled can have an effect on the final outcome. The SVM algorithm theoretically should produce the same result regardless of the input data order; however, an ML researcher familiar with SVM-Light told us that because it is inefficient to run the quadratic optimization algorithm on the full data set all at once, the implementation performs “chunking”

whereby the optimization algorithm runs on subsets of the data and then merges the results [17]. Numerical methods and heuristics are used to quickly converge toward the optimum; however, the optimum is not necessarily achieved, but instead this process stops after some threshold of improvement. Here the implementation deviates from the specification.

5. EVALUATION

Our contribution is to have demonstrated how random testing and partition testing can work together, instead of as alternatives, and also to show that this hybrid approach can be useful in applying the principles of random testing to “non-testable programs”.

The testing framework facilitated our work by aiding us in the creation, execution and analysis of the test cases. By combining parameterization and randomness, we gained the ability to control the properties of very large data sets, which was critical for limiting the scope of individual tests and for pinpointing specific issues in how the code was handling different equivalence classes. The data generation tool proved to be preferable to alternative approaches we considered, such as culling real-world data.

Additionally, the tool could be used for the testing of any ML ranking algorithm, not just MartiRank or SVM; it could also be used for supervised ML classification algorithms. The data generator supports plug-replaceable modules for creating data set files in whatever format is needed. Two such modules are currently implemented, one for MartiRank implementations (csv files) and the other for SVM-Light (a “sparse” attribute-value pair representation for data sets with a high ratio of missing values).

Our work does have some limitations. We have not yet chosen formal metrics for measuring the effectiveness compared to other approaches, such as looking at statement and path coverage, seeding the software with bugs to see whether they are detected, or using “pure” random testing by disabling parameterization. In addition, in order to create test cases reminiscent of real-world data, the framework should be extended to generate data sets that exhibit the same correlations among attributes and between attributes and labels as do real-world data, building upon [20]. Ideally it could also be extended to generate arbitrarily large data sets with repeating, missing and/or categorical data such that an arbitrary ML ranking algorithm could definitively enable a “predictable” ranking, i.e. where there is a clear-cut “correct” output. But this may be impossible in the general case (we have noted in [14] that data sets that yield predictable rankings in MartiRank do not necessarily yield the same ranking in SVM).

6. ACKNOWLEDGMENTS

We would particularly like to thank David Waltz, Hila Becker, Wei Chu, Philip Gross, Bert Huang, David Lee, Phil Long and Rocco Servedio for their assistance and encouragement. Murphy and Kaiser are members of the Programming Systems Lab, funded in part by NSF CNS-0717544, CNS-0627473, CNS-0426623 and EIA-0202063, NIH 1 U54 CA121852-01A1, and are also affiliated with the Center for Computational Learning Systems (CCLS). Arias is fully supported by CCLS, with funding in part by Consolidated Edison Company of New York.

7. REFERENCES

[1] H. Becker and M. Arias, “Real-time Ranking with Concept Drift Using Expert Advice”, *Proc. of KDD 2007*, Aug 2007.

[2] D. Bird and C. Munoz, “Automatic generation of random self-checking test cases”, *IBM Systems Journal* vol. 22 no. 3, 1983, 229-245.

[3] M.D. Davis and E.J. Weyuker, “Pseudo-Oracles for Non-Testable Programs”, *ACM '81 Conference*, 1981, 254-257.

[4] R. A. DeMillo and A. J. Offutt, “Constraint-Based Automated Test Data Generation”, *IEEE Trans. on Soft. Eng.* vol 17, no. 9, 1991, 900-910.

[5] J. Duran and S. Ntafos, “An Evaluation of Random Testing”, *IEEE Trans. on Soft. Eng.* vol 10, 1984, 438-444.

[6] P. Gross *et al.*, “Predicting Electricity Distribution Feeder Failures Using Machine Learning Susceptibility Analysis”, *Proc. of the 18th Conference on Innovative Applications in Artificial Intelligence*, Boston MA, 2006.

[7] D. Hamlet, “Random Testing”, in *Encyclopedia of Software Engineering*, Wiley, New York, 1994, 970-978.

[8] K.V. Hanford, “Automatic Generation of Test Cases”, *IBM Systems Journal* vol. 9 no. 4, 1970, 242-257.

[9] T. Joachims, *Making Large-Scale SVM Learning Practical. Advances in Kernel Methods - Support Vector Learning*, B. Schölkopf, C. Burges and A. Smola (ed.), MIT-Press, 1999.

[10] B. Korel, “Automated Software Test Data Generation”, *IEEE Trans. on Soft. Eng.* vol.16 no.8, August 1990, 870-879.

[11] K.J. Lieberherr, “Parameterized Random Testing”, *Proc. of the 21st Design Automation Conference*, 1984.

[12] P. Long and R. Servedio, “Martingale Boosting”, *Eighteenth Annual Conference on Computational Learning Theory*, Bertinoro, Italy, 2005, 79-94.

[13] J. Mayer, R. Guderlei, “Test Oracles Using Statistical Methods”, *Proc. of the First International Workshop on Software Quality*, 2004, 179-189.

[14] C. Murphy, G. Kaiser, M. Arias, “An Approach to Quality Assurance of Machine Learning Applications”, *Proc. of SEKE 2007*, Boston, July 2007, 167-172.

[15] C. Murphy, G. Kaiser, M. Arias, Columbia University Dept. of Computer Science tech report cucs-034-06, Sept 2006.

[16] D.J. Newman, S. Hettich, C.L. Blake and C.J. Merz, *UCI Repository of machine learning databases*, UC Irvine Dept. of Information and Computer Science, Irvine CA, 1998.

[17] R. Servedio, personal communication, 2006.

[18] P. Thévenod-Fosse, H. Waeselynck and Y. Crouzet, “An Experimental Study on Software Structural Testing”, *Proc. of the Twenty-First International Symposium on Fault-Tolerant Computing*, Montreal, June 1991, 410- 417.

[19] V.N. Vapnik, *The Nature of Statistical Learning Theory*. Springer, 1995.

[20] E. Walton, *Data Generation for Machine Learning Techniques*, University of Bristol, 2001.

[21] E.J. Weyuker, “On Testing Non-Testable Programs”, *Computer Journal* vol.25 no.4, November 1982, 465-470.

[22] B.A. Wichmann, “Some Remarks About Random Testing”, [http://www.npl.co.uk/scientific_software/publications/validation/random_testing.pdf]