

Competitive baseline methods set new standards for the NIPS 2003 feature selection benchmark

Isabelle Guyon^{a,*}, Jiwen Li^b, Theodor Mader^a, Patrick A. Pletscher^a,
Georg Schneider^a, Markus Uhr^a

^a *Pattern Analysis and Machine Learning Group, Institute of Computational Science, ETH Zurich, Universitaetstrasse 6, 8092 Zurich, Switzerland*

^b *Department of Informatics, University of Zurich, Binzmuhlestrasse 14, CH-8050 Zurich, Switzerland*

Received 21 November 2006

Available online 28 February 2007

Communicated by T.K. Ho

Abstract

We used the datasets of the NIPS 2003 challenge on feature selection as part of the practical work of an undergraduate course on feature extraction. The students were provided with a toolkit implemented in Matlab. Part of the course requirements was that they should outperform given baseline methods. The results were beyond expectations: the student matched or exceeded the performance of the best challenge entries and achieved very effective feature selection with simple methods. We make available to the community the results of this experiment and the corresponding teaching material [Anon. Feature extraction course, ETH WS 2005/2006. <http://clopinet.com/isabelle/Projects/ETH>]. These results also provide a new baseline for researchers in feature selection.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Feature selection; Matlab; Machine learning; Classification; Challenge; Benchmark; Curriculum

1. Introduction

In the recent years, it has been recognized by the machine learning and neural network communities that competitions are key to stimulate research and bring improvement. Several large conferences are now regularly organizing competitions. For NIPS 2003, a competition on the theme of feature selection, which attracted 75 participants, was organized (Guyon et al., 2004). The outcomes of that effort were compiled in a book including tutorial chapters and papers from the proceedings of that workshop (Guyon et al., 2006). The website of the challenge remains open

for post-challenge submissions (NIPS, 2003). Meanwhile, another challenge on the theme of model selection has taken place (WCCI, 2006), in which the participants were provided with a Matlab toolkit based on the Spider package (Weston et al., 2005). All this material constitute a great teaching resource that we have exploited in a course on feature extraction (Anon., 2005/2006). We are reporting on our teaching experience with several intentions:

- encouraging other teachers to use challenge platforms in their curricula;
- providing to graduate students simple competitive baseline methods to attack problems in machine learning;
- providing researchers in feature selection with baseline results supported by publicly available code.

The particular theme of the class is feature extraction, which we define as the combination of feature construction and feature selection. In the past few years, feature extrac-

* Corresponding author. Present address: Clopinet 955, Creston Road, Berkeley, CA 94708, USA. Tel./fax: +1 510 524 6211.

E-mail addresses: isabelle@clopinet.com (I. Guyon), li@ifi.unizh.ch (J. Li), tmader@student.ethz.ch (T. Mader), patrickp@student.ethz.ch (P.A. Pletscher), scgeorg@student.ethz.ch (G. Schneider), uhrm@student.ethz.ch (M. Uhr).

tion and space dimensionality reduction problems have drawn a lot of interest. More than a passing fancy, this trend in the research community is driven by applications: bioinformatics, chemistry (drug design, cheminformatics), text processing, pattern recognition, speech processing, and machine vision provide machine learning problems in very high dimensional spaces, but often with comparably few examples. A lot of attention was given in class to feature selection because many successful applications of machine learning have been built upon a large number of very low level features (e.g. the “bag-of-word” representation in text processing, gene expression coefficients in cancer diagnosis, and QSAR features in cheminformatics). Our teaching strategy is to make students gain hands on experience by working on large real world datasets (those of the NIPS 2003 challenge (Guyon, 2003)), rather than providing them with toy problems.

2. Datasets and synopsis of the challenge

The NIPS 2003 challenge included five datasets (Table 1) from various application domains. All datasets are two-class classification problems. The data were split into three subsets: a training set, a validation set, and a test set. All three subsets were made available at the beginning of the challenge. The class labels for the validation set and the test set were withheld. The challenge participants could submit prediction results on the validation set and get their performance results and ranking on-line during a development period. The validation set labels were then revealed and the participants could make submissions of test set predictions, after having trained on both the training and the validation set. For details on the benchmark design (see Guyon et al., 2004).

The identity of the datasets and of the features (some of which were random features artificially generated) were kept secret during the challenge, but have been revealed since then. The datasets were chosen to span a variety of domains and difficulties (the input variables are continuous or binary, sparse or dense; one dataset has unbalanced classes). One dataset (MADELON) was artificially constructed to illustrate a particular difficulty: selecting a feature set when no feature is informative by itself. To facilitate the assessment of feature selection methods, a number of artificial features called *probes* were drawn at random from a distribution resembling that of the real features, but carrying no

information about the class labels. A good feature selection algorithm should eliminate most of the probes. The details of data preparation can be found in a technical memorandum (Guyon, 2003).

The distribution of the results of the challenge participants for the various datasets are represented in Fig. 1. Using these graphs, we introduced in class the datasets in order of task difficulty. We provided the students with baseline methods approximately in the best tenth percentile of the challenge entries. For the class, the students had access to the training and validation set labels, but not the test labels. The test set labels have not been released to the public to keep an on-going benchmark. The students could thus made post-challenge submissions to the web site of the challenge (NIPS, 2003) to obtain their performance on the test set. The students were in similar conditions as the challenge competitors in that they did not have access to the test labels and they had to make fewer than five submissions (the limit set by the rules of the challenge). We asked students to try to outperform the baseline method and gave them extra credit for outperforming the best challenge entry (see Section 4).

3. Learning object package

The machine learning package we used for the class called CLOP (Challenge Learning Object Package) is available from the website of the “performance prediction challenge” (WCCI, 2006). For an introduction, see the QuickStart guide (Amir Reza Saffari Azar Alamdari, 2006). We present in this section a high level overview of the package.

3.1. Data and algorithm objects

The Spider package (Weston et al., 2005) on top of which CLOP is built, uses Matlab[®] objects. Two simple abstractions are used:

- *data*: Data objects include two members X and Y , X being the input matrix (patterns in lines and features in columns), Y being the target matrix (i.e. one column of \pm for binary classification problems);
- *algorithms*: Algorithm objects represent learning machines (e.g. neural networks, kernel methods, decision trees) or preprocessors (for feature construction,

Table 1
NIPS 2003 challenge datasets

Dataset	Domain	T	N_{feat}	F_{probe}^+	#Tr	#Va	#Te	$\%_{\text{[+]}}$	Tr/F
ARCENE	Mass spectrometry	d	10^4	0.30	100	100	700	44	0.01
DEXTER	Text classification	s	$2 \cdot 10^4$	0.50	300	300	2000	50	0.015
DOROTHEA	Drug discovery	sb	10^5	0.50	800	350	800	10	0.008
GISETTE	Digit recognition	d	5000	0.50	6000	1000	6500	50	1.2
MADELON	Artificial	d	500	0.96	2000	600	1800	50	4

For each dataset we show its domain of application, its type T (d = dense, s = sparse, or sb = sparse binary), the total number of features N_{feat} , the fraction of probes in the original feature set F_{probe}^+ , the number of examples in the training, validation, and test sets, the fraction of examples in the positive class $\%_{\text{[+]}}$, and the ratio number of training examples to number of features Tr/F. All problems are two-class classification problems.

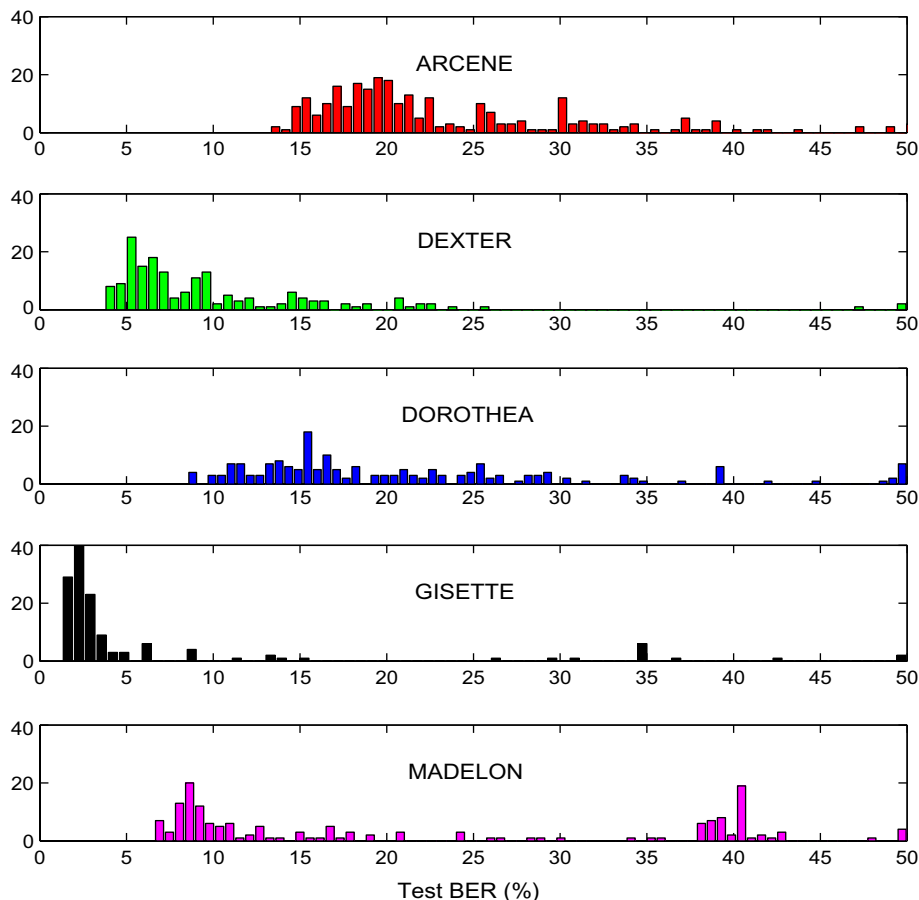


Fig. 1. Distribution of the challenge participant results. We show histograms of the balanced error rate (BER) for the five tasks.

data normalization or feature selection). They are constructed from a set of hyper-parameters and have at least two methods: train and test. The train method adjusts the parameters of the model. The test method processes data using a trained model.

Trained models are simply saved and reloaded with the `save/load` Matlab commands. This feature is convenient to verify results and enforce reproducibility.

The Spider (with some CLOP extensions) provides ways of building more complex “compound” models from the basic algorithms with two abstractions: *chains* and *ensembles*. Chains combine models serially while ensembles combine them in parallel. The Spider provides several objects for cross-validation and other model selection methods.

3.2. Learning objects provided for the class

It is easy to be overwhelmed when starting to use a machine learning package. CLOP is an extremely simplified package, limited to a few key methods, each having just a few hyper-parameter values with good default values. This makes it suitable for teaching a machine learning class. More advanced students can venture to using other methods provided in the Spider package, on top of which CLOP is built.

The CLOP modules correspond to methods having performed well in the feature selection challenge (Guyon et al., 2004). There are five classifier objects (we reference the implementations used): Kernel ridge regression (Guyon, 2005), naive Bayes (Guyon, 2005), neural networks (Bishop, 1996), Random Forest (Leo Breiman, 2001), and Support Vector Machines (Chang and Lin, 2001). We limited the number of hyperparameters of the algorithms to simplify model selection. For instance, the kernel methods use a single kernel: $k(\mathbf{x}, \mathbf{x}') = (\text{coef}_0 + \mathbf{x} \cdot \mathbf{x}')^d \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2)$ with three hyperparameters coef_0 , d , and γ . The CLOP modules also include five preprocessors performing standard normalization and feature extraction steps, and five feature selection methods.

We gave to the students other examples of learning objects which can be found in the homework instructions (Anonymous reference, 2005/2006). These include methods of feature extraction for image data (e.g. Fourier transform, and two-dimensional convolution), and statistical tests for feature selection. Details are found in a longer technical report (Guyon, 2006).

We provided the students with baseline methods, which are compound models combining the basic learning objects with “chains” and “ensembles” (see Guyon, 2006 for details). We chose the simplest possible baseline methods,

Table 2
Performance comparison

Dataset	Baseline		Challenge best		Student best		
	BER^0	F_{feat}^0	$\text{BER}^* \pm \delta_{\text{BER}}$	F_{feat}^*	BER	F_{feat}	Training time (s)
ARCENE	0.1470	0.11	0.1073 ± 0.0117	1.00	0.1048	0.14	3.8
DEXTER	0.0500	0.02	0.0330 ± 0.0040	0.19	0.0325	0.23	1.2
DOROTHEA	0.1237	0.01	0.0854 ± 0.0099	1.00	0.0930	0.01	0.7
GISETTE	0.0180	0.20	0.0126 ± 0.0014	1.00	0.0111	0.20	11.2
GISETTE (pixels)	0.0106	1.00	NA	NA	0.0078	1.00	127.9
MADÉLON	0.0733	0.04	0.0622 ± 0.0057	1.00	0.0622	0.04	7.8

The table shows the balanced error rate and the corresponding fraction of features used for the baseline method, the best challenge entry, and the best student method. The students earned one point for $\text{BER} < \text{BER}^0$ or $\{\text{BER} = \text{BER}^0 \text{ and } F_{\text{feat}} < F_{\text{feat}}^0\}$. They earned 2 points for $\text{BER} < \text{BER}^* + \delta_{\text{BER}}$. We also show the training time of the best student models on a 1.5 GHz Pentium.

which attain performances approximately in the best tenth percentile of the challenge entries (Table 2).

4. Performance assessment and class requisites

As part of the class requirements, the student had to submit results on all five datasets of the challenge and include their corresponding CLOP models. Performances were assessed using several metrics:

- **BER**: The balanced error rate, that is the average of the error rate of the positive class and the error rate of the negative class. This metric was used rather than the simple error rate because some datasets (particularly DOROTHEA) are unbalanced in class cardinality. Unless otherwise stated, the BER is computed with test examples.
- F_{feat} : The fraction of features selected.
- F_{probe} : The fraction of probes found in the feature set selected.¹ By normalizing F_{probe} by F_{probe}^+ , the fraction of probes in the whole feature set, we obtain an estimate of the False Discovery Rate (FDR). The FDR is indicative of Type I errors (irrelevant features wrongly selected).
- Z_{BER} : To assess the effectiveness of feature selection, we can compare a reduced model built with the selected features with the full model built using all the features. A large balanced error rate of the reduced model BER, compared to BER^+ obtained using all the features, is indicative of Type II errors (false negative, i.e. relevant features wrongly discarded). To assess the incidence of Type II errors we use the statistic $Z_{\text{BER}} = (\text{BER} - \text{BER}^+)/\sigma_{\Delta}$, where σ_{Δ} is the standard deviation of $\text{BER} - \text{BER}^+$.²

¹ We remind the reader that “probes” are meaningless features purposely added to the feature set to assess the effectiveness of feature selection.

² For small training sets, the benefit of reducing the dimensionality may outweigh the loss of information by discarding useful features, making Z_{BER} an imperfect indicator of Type II errors. The BER may actually be better than BER^+ , in which case Z_{BER} will be negative.

For each dataset, the students would earn one point if they obtained a better BER than the provided baseline method, or a smaller F_{feat} for the same BER. They would earn two points if they matched the BER of the best challenge entry (within the statistical error bar). The overall grade also included points for a poster presentation of their results and for the presentation in class of a research paper from the book (Guyon et al., 2006). The other performance metrics (F_{probe} and Z_{BER}) are used in our analysis (Section 5), but were not used for grading.

5. Student work

During the curriculum, the datasets were introduced one at a time to illustrate topics addressed in class, based on progressive difficulty. A description of the homework assignments associated with each dataset is provided in (Guyon, 2006). Briefly: GISETTE (the handwritten digit dataset) was used to illustrate feature construction. The student could experiment with transforms such as Fourier transform and convolution using the original pixel representation and create their own feature extraction learning objects. DEXTER (the text classification dataset) was used to learn about univariate filter methods and how to optimize the number of features with statistical tests or cross-validation. MADÉLON (the artificial dataset) gave an opportunity to the students to experiment with multivariate filters such as Relief, or try wrappers or embedded methods with non-linear multivariate classifiers. ARCENE (the mass-spectrometry dataset) included two difficulties: a small training set and heterogeneous data (coming from three sources). Therefore, it was a good dataset to further study multivariate feature selection algorithms, and experiment with mixtures of experts or ensemble methods. DOROTHEA (the drug discovery dataset) is the hardest because of its 100,000 features for only 800 training examples, with only ten percent of positive examples. The students learned about problems of overfitting and unbalanced data.

We report the results of all five students who returned results to get credit for the class. Each student made fewer than five entries (the limit set by the rules of the challenge). The best entries are shown in Table 3. The training set and

Table 3
Methods employed by the best student entries

Dataset	Code of the methods employed
ARCENE	<code>my_svc=svc({'coef0=2','degree=3','gamma=0','shrinkage=0.1'}); my_model=chain({'relief('f_max=1400')',normalize,my_svc})</code>
DEXTER	<code>my_svc=svc({'coef0=1','degree=1','gamma=0','shrinkage=0.5'}); my_model=chain({'s2n('f_max=4500')',normalize,my_svc})</code>
DOROTHEA	<code>my_model=chain({'TP('f_max=15000')',normalize,... relief('f_max=700')',naive,bias});</code>
GISETTE	<code>my_svc=svc({'coef0=1','degree=3','gamma=0','shrinkage=1'}); my_model=chain({'s2n('f_max=1000')',normalize,my_svc})</code>
GISETTE (pixels)	<code>my_svc=svc({'coef0=1','degree=4','gamma=0','shrinkage=0.1'}); my_model=chain({'convolve(gauss_ker({'dim1=5','dim2=5'})),... normalize,my_svc})</code>
MADOLON	<code>my_svc=svc({'coef0=1','degree=0','gamma=0.3','shrinkage=0.3'}); my_model=chain({'relief('f_max=20')',standardize,my_svc})</code>

validation set of the challenge were merged for training. The hyper-parameters were adjusted by cross-validation (usually 5-fold cross-validation). The final model is trained on all the available labeled data. As indicated in Table 2, the training times are modest (of the order of 1–10 s per model, except for GISETTE when the pixel representation is used, because of a time consuming preprocessing). Most students matched the performances of the best challenge entries (within the statistical error bar) and therefore earned the maximum number of points (Table 2 and Fig. 2). Some student entries even exceeded the performance of the best challenge entries. We think that the achievements of the students are remarkable. They exceeded our expectations. Of course, it can be argued that they had several advantages over the competitors: knowledge of the application domain, of the nature of the features and fraction of probes, and access to test




Fig. 2. Result comparison. The student performances matched (within the statistical error bar) or exceeded the performances of the best challengers, using comparable or smaller feature set sizes (see Table 2).

performance. However, this does not diminish significantly their achievement since many other post-challenge entries were made since the end of the challenge and hardly any matched or outperformed the best results of the challengers.³

General observations can be made about the best models. All of them use simple classifiers linear in their parameters (`naive` = naïve Bayes (Guyon, 2005) and `svc` = Support Vector Classifier (Boser et al., 1992)). Both the naïve Bayes and SVC are known to be robust against overfitting. Feature selection is performed with only three simple filters: `TP`, a filter useful for highly unbalanced classes like DOROTHEA to pre-select “true positive” features (Weston et al., 2003), `s2n`, a simple univariate filter, which ranks features according to the “signal-to-noise” (`s2n`) ratio (Golub et al., 1999), and `relief` a filter which ranks features according to their separating power “in the context of other features” (Kira and Rendell, 1992). The strategy adopted is to reduce space dimensionality while retaining as many relevant features as possible, with no attempt to remove feature redundancy. Additional regularization is obtained by tuning the classifier “shrinkage” hyperparameters.

By analyzing the results of the student work, stronger conclusions can be drawn about the effectiveness of feature selection than by analyzing the results of the challenge itself, because for four datasets out of five the best challenge entries used 100% of the features. In contrast, the student entries use a small subset of features, yet outperform the best challenge entries or match their performances within the statistical error bar.

³ We noted that, except in the case of GISETTE, where they were asked to take advantage of the knowledge of the features to outperform the challengers, the students did not take advantage of the knowledge of the nature of the features. Since they did not have access to the test labels, they could only submit results to the web site to get feed-back on their performance, which limited the possibility of tuning their method to the test set.

Table 4
Effectiveness of feature selection

Dataset	Full model		Reduced model			FDR	Z_{BER}
	F_{probe}^+	BER^+	F_{feat}	F_{probe}	BER		
ARCENE	0.30	0.1186	0.14	0.04	0.1048	0.14	-1.09
DEXTER	0.50	0.0410	0.23	0.55	0.0325	1.1	-1.75
DOROTHEA	0.50	0.3094	0.01	0.03	0.0930	0.05	-5.82
GISETTE	0.50	0.0180	0.20	0.00	0.0111	0	-3.49
MADÉLON	0.96	0.4872	0.04	0.00	0.0622	0	-35.53

The table shows the fraction of probes in the original data F_{probe}^+ and the BER of the full model BER^+ , the fraction of selected features F_{feat} , the fraction of probes in the selected feature sets F_{probe} , and the BER of the model built with the reduced feature set BER, the false discovery rate $FDR \simeq F_{\text{probe}}/F_{\text{probe}}^+$, and the Z statistic $Z_{\text{BER}} = (\text{BER} - \text{BER}^+)/\sigma_{\Delta}$.

To further evaluate the impact of feature selection, we reran the best models without feature selection, keeping the same hyper-parameters. We summarize the results in Table 4. In all cases, the reduced model (built with the selected features) has a smaller test BER than the full model (built with all the features) and consequently Z_{BER} is negative: there is no significant loss of information by discarding useful features and/or the benefit of space dimensionality reduction outweighs that loss. For the three last datasets, the reduced model is significantly better than the full model at the 1% risk level ($Z_{\text{BER}} < -2.33$) and the false discovery rate (FDR) is zero or near zero. This means that we have both very few type I and type II errors (few irrelevant features wrongly selected and few useful features discarded). At the 5% risk level ($Z_{\text{BER}} < -1.65$), feature selection yields significant performance improvement for one more dataset (DEXTER), even though the FDR of the selected features is very high. This shows that space dimensionality reduction plays an important role in performance improvement, regardless of the effectiveness of the method to filter out irrelevant features. For ARCENE, the performances of the reduced model and the full model are not significantly different. Yet, the feature set is significantly reduced without performance degradation and the FDR is relatively low.

Although it is difficult to predict ahead of time which method will work best on a given dataset, in retrospect, some justifications can be given. A simple rule-of-thumb for the choice of a classifier is that its complexity should be proportional to the ratio Tr/F of number of training examples over the number of features. If we sort the classifiers according to that ratio, we obtain: DOROTHEA < ARCENE < DEXTER < GISETTE < MADÉLON. Accordingly, it makes good sense that the naïve Bayes method (the simplest) was used for DOROTHEA, while the non-linear SVM (the most complex) was used for GISETTE and MADÉLON. DEXTER in the middle, uses the linear SVM of middle range complexity. It easily understood why relief performs well for MADÉLON and ARCENE: both datasets have classes containing multiple clusters. So we need a filter that performs a “local” feature selection. For the same reason, these two datasets benefit from using a non-linear classifier (particularly MADÉLON). For DOROTHEA, the problem of

class imbalance has been addressed by using a special feature filter (TF) and a special post-processor to adjust the bias. For a more detailed analysis dataset by dataset, see (Guyon, 2006).

6. Conclusions and future work

A challenge can be more than a one-time event. It can become an on-going life benchmark and a teaching tool. Leaving the website of the NIPS2003 feature selection challenge open for post-challenge submissions has given to graduate students and researchers the opportunity to compare their algorithms to well established baseline results. Since the end of the challenge, the number of entrants has almost doubled.

In this paper, we have demonstrated that even undergraduate students can get their hands dirty and “learn machine learning from examples”, with a success that exceeded our expectations. All of them easily outperformed the baseline methods we provided them and most of them matched the performances of the best challengers (within the statistical error bar) or even exceeded them. We hope that this experience will be followed by similar other attempts. In the mean time, we make available all of our teaching material, data and code.

The results obtained mark also a victory of simple methods. All the models used to match or outperform the best challenge entries use a combination of simple normalization, feature selection filters (signal-to-noise ratio, Relief, or fraction of true positive), and a naïve Bayes or a support vector machine classifier. There was no need to use ensemble methods or transduction. However, univariate feature selection and linear classifiers did not always suffice.

With this study, we could reach more conclusive results regarding the effectiveness of feature selection than by analyzing the results of the challenge. The best challenge entries significantly outperformed other entries without using feature selection. In contrast, the student entries used a reduced feature set, while matching or outperforming the performance of the best challenge entries. For three datasets the reduced model using a small fraction of the original feature set significantly outperformed the full model and the false discovery rate approached zero.

This paper by no means marks an end point to the problem of feature selection or even to solving the tasks of the NIPS2003 feature selection challenge. Our explorations indicate that there is still much room for improvement. In particular, since we have released the identity of the features, it is now possible to introduce domain knowledge in the feature construction process. To a limited extent we have seen that this strategy show promises on the G_{ISETTE} datasets: a simple smoothing of the pixel image allowed us to boost performances.

Acknowledgements

We are very thankful to the institutions that have contributed data (see Guyon, 2003), to the developers of the Spider who make their code available, and to the co-developers of the CLOP package Amir Reza Saffari Azar and Gideon Dror. The kind support of Prof. Joachim Buhmann and the encouragements and advices of Peter Orbanz are gratefully acknowledged. The CLOP project is partially supported by the National Science Foundation under Grant No. ECS-0424142. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Amir Reza Saffari Azar Alamdari and Isabelle Guyon, 2006. Quick start guide for CLOP. Technical Report, May 2006. <http://ymer.org/research/files/clop/QuickStartV1.0.pdf>.
- Anon. Feature extraction course, ETH WS 2005/2006. <http://clopinet.com/isabelle/Projects/ETH>.
- Bishop, Christopher M., 1996. *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford, UK.
- Boser, Bernhard, Guyon, Isabelle, Vapnik, Vladimir, 1992. A training algorithm for optimal margin classifiers. In: *Fifth Annual Workshop on Computational Learning Theory*. ACM, Pittsburgh, pp. 144–152.
- Breiman, Leo, 2001. Random forests. *Machine Learn.* 45 (1), 5–32 <http://citeseer.ist.psu.edu/breiman01random.html>.
- Chih-Chung Chang, Chih-Jen Lin, 2001. LIBSVM: A Library for Support Vector Machines. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Feature selection challenge. NIPS 2003. <http://www.nipsfsc.ecs.soton.ac.uk>.
- Golub, T.R. et al., 1999. Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring. *Science* 286, 531–537.
- Guyon, Isabelle, Gunn, Steve, Nikraves, Masoud, Zadeh, Lofti (Eds.), *Feature Extraction, Foundations and Applications*. Studies in Fuzziness and Soft Computing. Physica-Verlag, Springer.
- Isabelle Guyon, 2003. Design of experiments of the NIPS 2003 variable selection benchmark. Technical Report. <http://www.nipsfsc.ecs.soton.ac.uk/papers/Datasets.pdf>.
- Isabelle Guyon, 2005. Kernel ridge regression tutorial. Technical Report. <http://clopinet.com/isabelle/Projects/ETH/KernelRidge.pdf>.
- Isabelle Guyon, 2005. Naïve bayes algorithm in CLOP. Technical Report. <http://clopinet.com/isabelle/Projects/ETH/NaiveBayesAlgorithm.pdf>.
- Isabelle Guyon, et al. 2006. Feature selection with the CLOP package. Technical Report. <http://clopinet.com/isabelle/Projects/ETH/TM-fextract-class.pdf>.
- Isabelle Guyon, Steve R. Gunn, Asa Ben-Hur, Gideon Dror, 2004. Result analysis of the NIPS 2003 feature selection challenge. In: *NIPS*. http://books.nips.cc/papers/files/nips17/NIPS2004_0194.pdf.
- Kira, K., Rendell, L., 1992. A practical approach to feature selection. In: Sleeman, D., Edwards, P. (Eds.), *Proc. Internat. Conf. on Machine Learning*, Aberdeen, July 1992. Morgan Kaufmann, pp. 368–377.
- Performance prediction challenge. WCCI 2006. <http://www.modelselect.inf.ethz.ch>.
- Weston, Jason, Pérez-Cruz, Fernando, Bousquet, Olivier, Chapelle, Olivier, Elisseeff, André, Schölkopf, Bernhard, 2003. Feature selection and transduction for prediction of molecular bioactivity for drug design. *Bioinformatics* 19 (6), 764–771.
- Weston, J., Elisseeff, A., BakIr, G., Sinz, F., 2005. *The Spider Machine Learning Toolbox*. <http://www.kyb.tuebingen.mpg.de/bs/people/spider>.