

MULTIBOOST: A Multi-purpose Boosting Package

Djalel Benbouzid

Róbert Busa-Fekete*

*Linear Accelerator Laboratory
University of Paris-Sud, CNRS
Orsay 91898, France*

DJALEL.BENBOUZID@GMAIL.COM

BUSAROBI@GMAIL.COM

Norman Casagrande

Wavii, Inc.

*13-19 Bevendon Street
London, N1 6AA, United Kingdom*

NORMAN@WAVII.COM

François-David Collin

Balázs Kégl†

*Linear Accelerator Laboratory
University of Paris-Sud, CNRS
Orsay 91898, France*

FRADAV@GMAIL.COM

BALAZS.KEGL@GMAIL.COM

Editor: Sören Sonnenburg

Abstract

The MULTIBOOST package provides a fast C++ implementation of multi-class/multi-label/multi-task boosting algorithms. It is based on ADABOOST.MH but it also implements popular cascade classifiers and FILTERBOOST. The package contains common multi-class base learners (stumps, trees, products, Haar filters). Further base learners and strong learners following the boosting paradigm can be easily implemented in a flexible framework.

Keywords: boosting, ADABOOST.MH, FILTERBOOST, cascade classifier

1. Introduction

ADABOOST (Freund and Schapire, 1997) is one of the best off-the-shelf learning methods developed in the last fifteen years. It constructs a classifier in an incremental fashion by adding simple classifiers to a pool, and uses their weighted “vote” to determine the final classification. ADABOOST was later extended to multi-class classification problems (Schapire and Singer, 1999). Although various other attempts have been made handle the multi-class setting, ADABOOST.MH has become the gold standard of multi-class boosting due to its simplicity and versatility.

Despite the simplicity and the practical success of the ADABOOST, there are relatively few off-the-shelf implementations available in the free software market. Whereas binary ADABOOST with decision stumps is easy to code, multi-class ADABOOST.MH and complex base learners are not straightforward to implement efficiently. The MULTIBOOST software package is intended to fill this gap. Its main boosting engine is based on the ADABOOST.MH algorithm of Schapire and

*. R. Busa-Fekete is on leave from the Research Group on Artificial Intelligence of the Hungarian Academy of Sciences and University of Szeged.

†. Also in the Computer Science Laboratory.

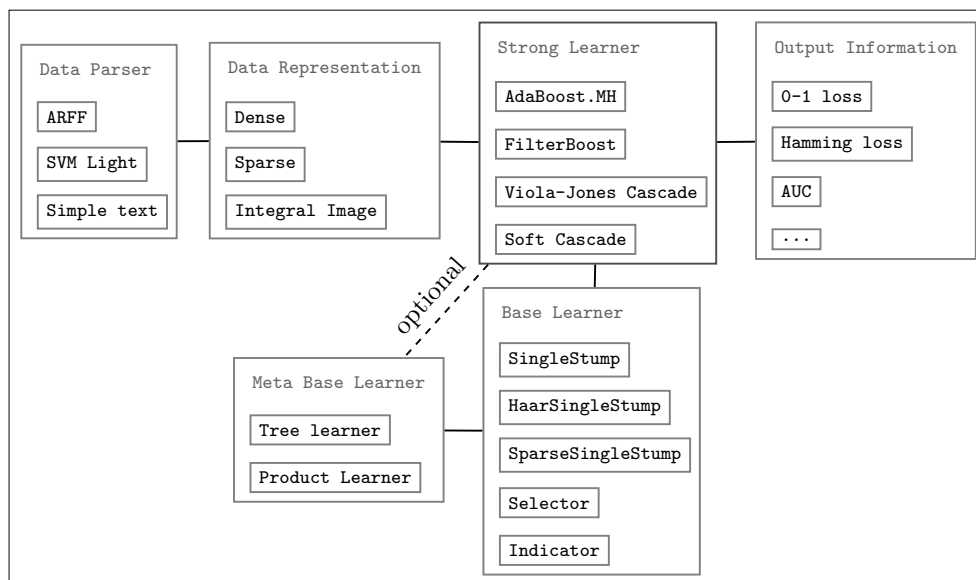


Figure 1: The architecture of MULTIBOOST

Singer (1999), but popular cascade classifiers (VJCASCADE (Viola and Jones, 2004), SOFTCASCADE (Bourdev and Brandt, 2005)) and FILTERBOOST (Bradley and Schapire, 2008) have also been implemented. The package includes common multi-class base learners (real and nominal valued decision stumps, trees, products (Kégl and Busa-Fekete, 2009), and Haar filters), but the flexible architecture makes it simple to add new base learners without interfering with the main boosting engine. MULTIBOOST was designed in the object-oriented paradigm and coded in C++, so it is fast and it provides a flexible base for implementing further modules.

The rest of this paper is organized as follows. Section 2 describes the general architecture and the modules of MULTIBOOST. Section 3 deals with practical issues (website, documentation, licence), and Section 4 describes some of our results obtained on benchmark data sets and in data mining challenges.

2. The Architecture

MULTIBOOST was implemented within the object-oriented paradigm using some design patterns. It consists of several modules which can be changed or extended more or less independently (Figure 1). For instance, an advanced user can implement a data-type/base-learner pair without any need to modify the other modules.

2.1 Strong Learners

The strong learner¹ calls the base learners iteratively, stores the learned base classifiers and their coefficients, and manages the weights of the training instances. The resulting classifier is serialized

1. The name originally comes from the boosting (PAC learning) literature. Here, we use it in a broader sense to mean the “outer” loop of the boosting iteration.

in a human-readable XML format that allows one to resume a run after it was stopped or crashed. MULTIBOOST implements the following strong learners:

- ADABOOST.MH (Schapire and Singer, 1999): a multi-class/multi-label/multi-task version of ADABOOST that learns a “flat” linear combination of vector-valued base classifiers.
- FILTERBOOST (Bradley and Schapire, 2008): an online “filtering” booster.
- VJCASCADE (Viola and Jones, 2004): an algorithm that learns a cascade classifier tree by running ADABOOST at each node.
- SOFTCASCADE (Bourdev and Brandt, 2005): another cascade learner that starts with a set of base classifiers, reorders them, and augments them with rejection thresholds.

2.2 Base Learners

MULTIBOOST implements the following base learners.

- The STUMP learner is a one-decision two-leaf tree learned on real-valued features. It is indexed by the feature it cuts and the threshold where it cuts the feature.
- SELECTOR is a one-decision two-leaf tree learned on nominal features. It is indexed by the feature it cuts and the value of the feature it selects.
- INDICATOR is similar to SELECTOR but it can select several values for a given feature (that is, it can *indicate* a subset of the values).
- HAARSTUMP is a STUMP learned over a feature space generated using rectangular filters on images.
- TREE is a *meta* base learner that takes any base learner as input and learns a vector-valued multi-class decision tree using the input base learner as the basic cut.
- PRODUCT is another meta base learner that takes any base learner as input and learns a vector-valued multi-class decision product (Kégl and Busa-Fekete, 2009) using the input base learner as terms of the product.

2.3 The Data Representation

The multi-class data structure is a set of observation-label pairs, where each observation is a vector of feature values, and each label is a vector of binary class indicators. In binary classification, we also allow one single label that indicates the class dichotomy. In single-label multi-class classification, only one of the K labels is 1 and the others are -1 , but the framework also allows multi-label classification with several positive classes per instance. In addition, multi-task classification can be encoded by letting each label column represent a different task. We implement a sparse data representation for both the observation matrix and the label matrix. In general, base learners were implemented to work with their own data representation. For example, SPARSESTUMP works on sparse observation matrices and HAARSTUMP works on an integral image data representation.

2.4 The Data Parser and the Output Information

The training and test sets can be input in the attribute-relation file format (ARFF),² in the SVM-LIGHT format,³ or using a comma separated text file. We augmented the first two formats with initial label weighting, which is an important feature in the boosting framework (especially in the multi-class/multi-label setup).

In each iteration, MULTIBOOST can output several metrics (specified by the user), such as the 0-1 error, the Hamming loss, or the area under the ROC curve. New metrics can also be implemented without modifying other parts of the code.

2.5 LAZYBOOST and BANDITBOOST

When the number of features is large, *featurewise* learners (STUMP, SELECTOR, and INDICATOR) can be accelerated by searching only a subset of the features in each iteration. MULTIBOOST implements two options, namely, LAZYBOOST (Escudero et al., 2000), where features are sampled randomly, and BANDITBOOST (Busa-Fekete and Kégl, 2010), where the sampling is biased towards “good” features learned using a multi-armed bandit algorithm.

3. Documentation and License

The code of MULTIBOOST has been fully documented in Doxygen.⁴ It is available under the GPL licence at multiboost.org. The website also provides documentation that contains detailed instructions and examples for using the package along with tutorials explaining how to implement new features. The documentation also contains the pseudo-code of the multi-class base learners implemented in MULTIBOOST.

4. Challenges and Benchmarks

We make available reproducible test results (validated test errors, learning curves) of MULTIBOOST on the web site as we produce them. Among other results, the boosted decision product is one of the best reported no-domain-knowledge algorithms on MNIST.⁵ An early version of the program (Bergstra et al., 2006) was the best genre classifier out of 15 submissions and the second-best out of 10 submissions at recognizing artists in the MIREX 2005 international contest on music information extraction. More recently, we participated in the Yahoo! Learning to Rank Challenge⁶ using a pointwise ranking approach based on hundreds of MULTIBOOST classifiers. We finished 6th in Track 1 and 11th in Track 2 out of several hundred participating teams (Busa-Fekete et al., 2011).

Acknowledgments

This work was supported by the ANR-2010-COSI-002 grant of the French National Research Agency.

2. See www.cs.waikato.ac.nz/~ml/weka/arff.html.

3. See svmlight.joachims.org.

4. See www.doxygen.org.

5. See yann.lecun.com/exdb/mnist.

6. See learningtorankchallenge.yahoo.com.

References

- J. Bergstra, N. Casagrande, D. Erhan, D. Eck, and B. Kégl. Aggregate features and AdaBoost for music classification. *Machine Learning Journal*, 65(2/3):473–484, 2006.
- L. Bourdev and J. Brandt. Robust object detection via soft cascade. In *Conference on Computer Vision and Pattern Recognition*, volume 2, pages 236–243. IEEE Computer Society, 2005.
- J.K. Bradley and R.E. Schapire. FilterBoost: Regression and classification on large datasets. In *Advances in Neural Information Processing Systems*, volume 20. The MIT Press, 2008.
- R. Busa-Fekete and B. Kégl. Fast boosting using adversarial bandits. In *International Conference on Machine Learning*, volume 27, pages 143–150, 2010.
- R. Busa-Fekete, B. Kégl, Éltető T., and Gy. Szarvas. Ranking by calibrated AdaBoost. In (*JMLR W&CP*), volume 14, pages 37–48, 2011.
- G. Escudero, L. Màrquez, and G. Rigau. Boosting applied to word sense disambiguation. In *Proceedings of the 11th European Conference on Machine Learning*, pages 129–141, 2000.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997.
- B. Kégl and R. Busa-Fekete. Boosting products of base classifiers. In *International Conference on Machine Learning*, volume 26, pages 497–504, Montreal, Canada, 2009.
- R.E. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3):297–336, 1999.
- P. Viola and M. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57:137–154, 2004.