

Boosting

Peter Bühlmann

buhlmann@stat.math.ethz.ch

Seminar für Statistik

ETH Zürich

Zürich, CH-8092, Switzerland

Bin Yu

binyu@stat.berkeley.edu

Department of Statistics

University of California

Berkeley, CA 94720-3860, USA

March 30, 2009

Abstract

In this contribution, we review boosting, one of most effective machine learning methods for classification and regression. Most of the article takes the gradient descent point of view, even though we do include the margin point of view as well. In particular, AdaBoost in classification and various versions of L2boosting in regression are covered. Advice on how to choose base (weak) learners and loss functions and pointers to software are also given for practitioners.

Key Words: boosting, gradient descent, AdaBoost, L2Boost, base learner.

1 Introduction

Boosting is a modern statistical methodology for classification or regression that originated from machine learning in the mid 90's and has seen much progress in recent years in terms of its understanding and effectiveness in modeling real world data. Its successful applications widely vary and include text mining, image segmentation, language translation, bioinformatics, neuroscience, and finance, just to name a few. Boosting has many distinctive characteristics. First, it is an ensemble method through linear combination of so-called weak learners. Second, it is closely connected to optimization (often convex) gradient descent (iterative) algorithms. The first trait is shared by other statistical machine learning methods such as bagging and random forests, and the second by various programming methods for optimization.

A third characteristic of boosting is its early stopping rule to achieve regularization. This trait couples regularization with numerical optimization and makes boosting very special. A fourth characteristic is the re-weighting of data examples (units) based on the outcome of the previous iteration. This trait is a consequence of the form of the gradient of the loss function in the optimization and hence not shared by all variants of boosting as we know now. It is, however, an important feature of the original AdaBoost algorithm (Freund and Schapire, 1996).

Kearns and Valiant (1989) postulated the boosting conjecture in the framework of PAC learning that a weak classifier (with success probability just a bit over 50 %) can be “boosted” into a strong one in the sense that the training error of the new one would go to zero and with a polynomial-time run time. In his award-winning MIT thesis in Computer Science, Schapire (1990) answered the question in its affirmative, and Freund and Schapire (1996) devised the first practical boosting algorithm, *AdaBoost*, in binary classification.

Breiman (1998, 1999) made the crucial connection of AdaBoost to optimization by re-deriving AdaBoost from a gradient descent point of view on an exponential loss function of the margin (which is the product yf of the true label $y \in \{-1, 1\}$ and the prediction f). This connection was further developed by Friedman et al. (2000) and Mason et al. (2000). Many variants of Boosting followed in the statistics community by generalizations of AdaBoost to different loss functions and through different gradient descent methods. The most prominent ones are LogitBoost (Friedman et al., 2000) in classification and L2Boosting (Friedman, 2001) in regression. A separate line of research on AdaBoost has concentrated on the margin view point mainly by the machine learning community (Schapire et al., 1998; Rätsch et al., 2001; Rosset et al., 2004; Rätsch and Warmuth, 2005; Rudin et al., 2007) and recent developments there are emphasizing AdaBoost’s ensemble and re-weighting aspects (Warmuth et al., 2008b) and references therein).

Theoretical analyses of boosting have also been developed by both the theoretical statistics and theoretical machine learning communities. The former community studied, for example, the Bayes and model selection consistency property and minimax rates of convergence of various Boosting procedures (when stopped early) (Bühlmann and Yu, 2003; Jiang, 2004; Lugosi and Vayatis, 2004; Zhang and Yu, 2005; Bühlmann, 2006; Bartlett and Traskin, 2007); while the latter community studied algorithmic convergence of boosting (Rätsch et al., 2001) and upper bounds on the generalization error of AdaBoost via distributions of margins and complexity measures of the model class and the sample size (Schapire et al., 1998). One insight from

these theoretical analyses is that the complexity of the boosting procedure is not simply a linear function of the number of iterations. In a simple regression case studied in Bühlmann and Yu (2003), the complexity increases with the number of iteration in an exponentially diminishing manner – the complexity is upper bounded by the noise variance and the amount of complexity added is exponentially small when the iteration number increases. This partially explains the overfitting resistance of AdaBoost. Eventually, AdaBoost and all other boosting methods will overfit, but it could take a long time especially in classification due to the robustness property of the evaluating 0-1 loss function in addition to the slow increase of the complexity of boosting. Most of the theoretical studies so far have been carried out under the i.i.d. assumption although some generalizations to stationary data sources have been analyzed (Lutz and Bühlmann, 2006; Lozano et al., 2006).

In this contribution, we review the general boosting algorithm mainly from the gradient descent point of view. In particular, we cover the AdaBoost, LogitBoost and L2Boosting algorithms. After a short section on the margin view point, we end with a discussion on the choices of loss functions, base learners and stopping rules.

2 The gradient descent point of view

We describe here the general boosting algorithm from the gradient descent point of view (Breiman, 1998, 1999; Friedman et al., 2000; Mason et al., 2000; Friedman, 2001). In particular, we review AdaBoost and LogitBoost for classification (corresponding to different loss functions L) and various versions of L2Boosting for regression (corresponding to the L2-loss function). Finally, we mention briefly the generalization of boosting from empirical loss to penalized empirical loss such as Lasso (Tibshirani, 1996) in regression.

Suppose that we observe

$$(X_1, Y_1), \dots, (X_n, Y_n),$$

where $X_i \in \mathbb{R}^p$ denotes a p -dimensional predictor variable and Y_i a univariate response, for example taking values in \mathbb{R} as in regression or in $\{-1, +1\}$ as in binary classification. In the sequel, we denote by $X^{(j)}$ the j th component of a vector $X \in \mathbb{R}^p$. We usually assume that the pairs (X_i, Y_i) are i.i.d. or from a stationary process. The goal is to estimate the regression function $F(x) = \mathbb{E}[Y|X = x]$ or to find a classifier $\text{sign}(F(x))$ where $F(x) : \mathbb{R}^p \rightarrow \mathbb{R}$.

The estimation performance is evaluated via a real-valued loss function in the sense that we want to minimize the expected loss or risk:

$$\mathbb{E}L(Y, F(X)),$$

based on data $(X_i, Y_i)(i = 1, \dots, n)$. The loss function L is assumed to be smooth and convex in the second argument so that the gradient method can be applied. Boosting algorithms are obtainable by minimizing the empirical loss function

$$n^{-1} \sum_{i=1}^n L(Y_i, F(X_i)),$$

over an additive expansion of base learners for F via functional gradient descent. The base learners take the form of $h(x, \hat{\theta})$ ($x \in \mathbb{R}^p$) where $\hat{\theta}$ is an estimated parameter of finite or infinite dimension. For example, the base learners could be stumps with $\hat{\theta}$ describing the axis to split, the split points and the fitted values for the two terminal nodes. Fitting the base learner based on data is part of the base learner.

The boosting methodology in general builds on a user-determined base procedure or weak learner and uses it repeatedly on modified data which are typically outputs from the previous iterations. The final boosted procedure takes the form of linear combinations of the base procedures. Precisely, given a base learner $h(x, \theta)$, boosting is derivable as functional gradient descent on the loss function L .

Boosting (gradient descent view)

1. Start with $F_0(x) = 0$.
2. Given $F_{m-1}(x)$, let

$$(\beta_m, h(x, \hat{\theta}_m)) = \operatorname{argmin}_{\beta \in \mathbb{R}, \theta} \sum_{i=1}^n L(Y_i, F_{m-1}(X_i) + \beta h(x, \theta)).$$

3. Set

$$F_m(x) = F_{m-1}(x) + \beta_m h(x, \hat{\theta}_m).$$

4. Stop when $m = M$.

The AdaBoost classifier is $\operatorname{sign}(F_M(x))$.

2.1 Classification

In binary classification, $y \in \{-1, +1\}$ and the most commonly used loss is the 0-1 loss. That is, for a classifier $\text{sign}(F(x)) \in \{-1, +1\}$ if the label of x is $y \in \{-1, +1\}$, the 0-1 loss can be written as a function of the margin $yF(x)$:

$$L_{01}(y, F(x)) = I\{yF(x) < 0\}.$$

It is easy to see that the exponential loss function

$$L_{exp}(y, F(x)) = \exp(-yF(x))$$

is an upper bound on L_{01} and its population minimizer is half of the log odds ratio

$$F(x) = \frac{1}{2} \log \frac{\mathbb{P}(Y = 1|X = x)}{\mathbb{P}(Y = -1|X = x)}. \quad (1)$$

For e.g. AdaBoost, the weak or base learner is a procedure that maps from a given training data set to a classifier with a training error better than a random guess, or less than 50%. Often used are tree-based classifiers. AdaBoost improves upon the current fit in an additive way to minimize the empirical exponential loss function over the base learner (acting on a modified data set) and a multiplier.

AdaBoost

1. Start with $F_0(x) = 0$;
2. Given $F_{m-1}(x)$, let

$$w_i^{(m)} = \exp(-Y_i F_{m-1}(X_i)),$$

$$h(x, \hat{\theta}_m) = \underset{\theta}{\text{argmin}} \sum_{i=1}^n w_i^{(m)} I(Y_i \neq h(X_i, \theta)),$$

and denote $h(\cdot, \hat{\theta}_m)$'s associated error by

$$\text{err}_m = \frac{\sum_{i=1}^n w_i^{(m)} I(Y_i \neq h(X_i, \hat{\theta}_m))}{\sum_{i=1}^n w_i^{(m)}}.$$

Furthermore let

$$\beta_m = \frac{1}{2} \log \frac{1 - \text{err}_m}{\text{err}_m}.$$

3. Set

$$F_m(x) = F_{m-1}(x) + \beta_m h(x, \hat{\theta}_m).$$

4. Stop when $m = M$.

5. The AdaBoost classifier is $y = \text{sign}(F_M(x))$.

Clearly, two key inputs to the AdaBoost algorithm are the choices of base learner $h(\cdot, \theta)$ and the stopping rule for M . Relatively simple tree classifiers such as stumps or 8-node trees have been used effectively as the base learners in many empirical studies. The stopping iteration M acts as a regularization parameter. A test data set, when the data set is large, or cross-validation can be used to find such an M . When M tends to infinity, the AdaBoost estimator has been shown to converge to the minimizer of the empirical exponential loss over the set of linear combinations of base learners (Zhang and Yu (2005); Bickel et al. (2006); Bartlett and Traskin (2007)).

We defer further discussions on the choices of the base procedure and some stopping rules for M to Section 4.

LogitBoost

If the loss function is the negative log likelihood function from a logistic model with a logit link function, we get the loss function

$$\log(1 + \exp^{-yF}),$$

or equivalently

$$L_{logit}(y, F) = \log_2(1 + \exp^{-yF}),$$

which is an upper bound on the 0-1 loss function L_{01} as a function of the margin yF . Moreover, the expected loss of L_{logit} is minimized by the same function as in (1), as for the exponential loss function.

In multi-class situations, a log-likelihood function in multinomial models can be used to apply the gradient descent algorithm to arrive at a multi-class boosting algorithm. However, often in practice, the one-against-all approach is used to turn a multi-class problem into many binary classification problem such that AdaBoost or LogitBoost can be applied.

2.2 Regression: Boosting with the squared error loss

In regression, a natural loss function is the squared error loss. With this loss function, we get *L2Boosting* (Friedman, 2001). When applying the gradient descent Boosting algorithm with the squared loss, we end up with a repeated fitting of residuals with the base learner. Analogously as before, *L2Boosting* is a “constrained” minimization of the empirical squared error risk $n^{-1} \sum_{i=1}^n (Y_i - F(X_i))^2$ (with respect to $F(\cdot)$) which yields an estimator $\hat{F}(\cdot)$. The regularization of the empirical risk minimization comes in again implicitly by the choice of a base procedure and by algorithmical constraints such as early stopping or some penalty barriers.

L2Boosting (with base procedure $h(\cdot, \theta)$)

1. Start with $F_0 = 0$.
2. Given $F_{m-1}(x)$, Compute residuals $U_i = Y_i - F_{m-1}(X_i)$ ($i = 1, \dots, n$). Fit the base procedure to the current residuals:

$$h(x, \hat{\theta}_m) = \operatorname{argmin}_{\theta} \sum_{i=1}^n (U_i - h(X_i, \theta))^2,$$

- 3.

$$F_m(x) = F_{m-1}(x) + \beta_m h(x, \hat{\theta}_m),$$

where the line search turns out to give $\beta_m \equiv 1$.

4. Stop when $m = M$ and $F_M(x)$ is the final estimator of the regression function.

As before, the number of iterations M is the main tuning parameter for *L2Boosting* in addition to the base learner B . The number of boosting iterations may be estimated by cross-validation. As an alternative, one may use some model selection criteria to bypass cross-validation and bring computation savings.

When the base procedure selects just one predictor (feature) among all of the p variables, gradient descent becomes coordinatewise gradient descent. For example, for the componentwise smoothing spline base procedure which fits a smoothing spline to the predictor variable reducing residual sum of squares most, the boosting estimator yields an additive model fit (Bühlmann and Yu, 2003). Another widely used procedure is componentwise linear least

squares yielding a linear model fit.

L2Boosting: coordinatewise descent for linear models

1. Start with $F_0 = 0$.
2. Given $F_{m-1}(x)$, Compute residuals $U_i = Y_i - F_{m-1}(X_i)$ ($i = 1, \dots, n$).
Let $X_i^{(j)}$ be the j th component of $X_i \in \mathbb{R}^p$,

$$\hat{j}_m = \operatorname{argmin}_{j=1, \dots, p} \sum_{i=1}^n (U_i - \hat{\beta}_m X_i^{(j)})^2,$$

$$\hat{\beta}_m = \operatorname{argmin}_{\beta} \sum_{i=1}^n (U_i - \beta X_i^{(\hat{j}_m)})^2,$$

- 3.

$$F_m(x) = F_{m-1}(x) + \hat{\beta}_m X^{(\hat{j}_m)},$$

4. Stop when $m = M$ and $F_M(x)$ is the final estimator of the linear regression function.

Friedman (2001) introduced shrinkage to L2Boosting by shrinking the base learner using $\nu \cdot h(\cdot, \theta)$ with $0 < \nu \leq 1$. Empirical evidence suggests that the choice for the step-size ν isn't crucial as long as ν is small; we usually use $\nu = 0.1$. A related version of L2Boosting has been termed e-L2Boosting in the literature (Hastie et al. (2001); Efron et al. (2004)):

e-L2Boosting Normalize all predictors to the same scale and use $\hat{\beta}_m \equiv \nu > 0$, a fixed “step-size”.

With $m = 2$ and $\nu = 1$, L2Boosting has already been proposed by Tukey (1977) under the name “twicing”. e-L2Boosting is also called Forward Stagewise Fitting in Efron et al. (2004). It is shown there to be connected to the L1-penalized least squares method Lasso (Tibshirani, 1996) and often shares the sparsity property of Lasso. In fact, Lasso and e-L2Boosting are connected through the Blasso algorithm (Zhao and Yu, 2007) that consists of e-L2boosting steps (forward steps) and appropriately chosen backward steps where predictors could be removed from the fitted function.

Instead of L1-penalized least squares one can take L0-penalized least squares, although the latter is computationally awkward and infeasible to solve. This is related to SparseBoosting (Bühlmann and Yu, 2006) which

employs some information criteria such as AIC, BIC or gMDL, while gMDL was recommended by Bühlmann and Yu (2006) for overall performance of prediction and sparsity. Recently, Zhang (2008) combines forward and backward steps to minimize L0-penalized LS and Friedman and Popescu (2005) devise a gradient direction regularization algorithm which does not necessarily come from an explicit penalty function.

3 The margin point of view in classification

Schapire et al. (1998) suggest a new explanation for the effectiveness of AdaBoost via boosting the margin distribution. Note that the margin is a key concept in the methodology and theory of support vector machines. This work started a separate trend of generalizations of AdaBoost in the machine learning community through maximizing various versions of the margin, leading to e.g. LPBoost (Demiriz et al., 2002), SoftBoost (Warmuth et al., 2008a) and Entropy regularized LPBoost (Warmuth et al., 2008b). Among these margin-based algorithms, some are seen to be only “corrective” in the sense that they base the re-weighting on only the outcomes of the previous iteration, and some are “totally corrective” because the re-weighting takes into account the outcomes of all previous iterations. Some studies (with C4.5 or radial functions as base learners) indicate that these margin based methods have similar or slightly better classification accuracy than AdaBoost (and possibly than LogitBoost as well). In a computational speed comparison study in Warmuth et al. (2008b), LPBoost is found to be the fastest, Entropy regularized LPBoost slightly slower while SoftBoost is much slower. These algorithms rely on primal-dual formulations in convex optimization from which some convergence rate analyses are known for SoftBoost and Entropy regularized LPBoost.

4 Practical issues in applying Boosting

The boosting methodology as we know now has three ingredients: loss function, base learner and stopping rule. To apply boosting in practice, choices have to be made for these three components. Similar to applying any other methodology in practice, these choices are often subjective, depending on the familiarity of the practitioner, the availability of software or the time for computational implementation. Nevertheless, we give here some rule of thumb advice on these choices based on our and others experience of using boosting in various problems.

Loss function

In classification, AdaBoost and LogitBoost are the most popular choices of loss functions with LogitBoost’s performance slightly better than that of AdaBoost. Among margin-based generalizations of AdaBoost, LPBoost seems to be a good choice (Warmuth et al. (2008b)).

In regression, L2Boosting and often also its shrunken version e-L2Boosting is used. Arguably, more robust loss functions and base procedures should be used (Lutz et al., 2008), but in practice, they are not in common use perhaps because of available software.

Base learner

At a high level, the conventional wisdom on the choice of base learner is that one should use a “weak” learner or a simple procedure in terms of complexity. This allows the boosting methodology to adaptively build up the complexity of the boosting procedure tailored to the particular problem through the iterative fitting recipe. For example, in L2Boosting, if one starts with a strong learner such as projection pursuit, things can get worse even at the second step because boosting can not correct the earlier iterations of overfitting.

The most used base learner with boosting (AdaBoost or its generalizations) in classification is CART or C4.5 (i.e. tree-based classifiers). Stumps have been used in many problems and slightly stronger learners are trees with a moderate number of nodes, say 8-node trees.

In regression, the coordinatewise descent version of L2Boosting with small steps and e-L2boosting have become popular for high-dimensional linear and additive modeling.

Early stopping rule

Performance on a test set or cross-validation is usually the choice of early stopping if the goal is prediction. If the sample size is large (relative to the complexity of the fitting), a single test-set is often sufficient. Alternatively, if the sample size is small (relative to the complexity of the fitting), the single test-set idea is not accurate and cross-validation should be used. However, cross-validated prediction errors incur large variances and hence may not be reliable. A few alternatives exist based on model selection criteria such as AIC (or AIC_c), BIC or gMDL. The idea is to use a possibly biased estimated prediction error with a smaller variance than the cross-validation error. Moreover, the computation cost may be many fold reduced in comparison to cross-validation and thus, for very large data sets, these model-based

prediction error estimates are also desirable due to computational savings.

Software in R

The R-package `mboost` provides many specific versions of the gradient descent Boosting algorithm, including choices for stopping rules. In addition, `mboost` allows to fit boosting with a user-specific loss function L . A review of Boosting including examples using `mboost` is given in Bühlmann and Hothorn (2007).

References

- BARTLETT, P. and TRASKIN, M. (2007). AdaBoost is consistent. *Journal of Machine Learning Research* **8** 2347–2368.
- BICKEL, P., RITOV, Y. and ZAKAI, A. (2006). Some theory for generalized boosting algorithms. *Journal of Machine Learning Research* **7** 705–736.
- BREIMAN, L. (1998). Arcing classifiers (with discussion). *The Annals of Statistics* **26** 801–849.
- BREIMAN, L. (1999). Prediction games & arcing algorithms. *Neural Computation* **11** 1493–1517.
- BÜHLMANN, P. (2006). Boosting for high-dimensional linear models. *The Annals of Statistics* **34** 559–583.
- BÜHLMANN, P. and HOTHORN, T. (2007). Boosting algorithms: regularization, prediction and model fitting (with discussion). *Statistical Science* **22** 477–505.
- BÜHLMANN, P. and YU, B. (2003). Boosting with the L_2 loss: Regression and classification. *Journal of the American Statistical Association* **98** 324–339.
- BÜHLMANN, P. and YU, B. (2006). Sparse boosting. *Journal of Machine Learning Research* **7** 1001–1024.
- DEMIRIZ, A., BENNETT, K. and SHAWE-TAYLOR, J. (2002). Linear programming boosting via column generation. *Journal of Machine Learning Research* **46** 225–254.
- EFRON, B., HASTIE, T., JOHNSTONE, I. and TIBSHIRANI, R. (2004). Least angle regression (with discussion). *The Annals of Statistics* **32** 407–451.

- FREUND, Y. and SCHAPIRE, R. (1996). Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA.
- FRIEDMAN, J. (2001). Greedy function approximation: a gradient boosting machine. *The Annals of Statistics* **29** 1189–1232.
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2000). Additive logistic regression: a statistical view of boosting (with discussion). *The Annals of Statistics* **28** 337–407.
- FRIEDMAN, J. H. and POPESCU, B. E. (2005). Gradient directed regularization. Tech. rep., Stanford University.
URL <http://www-stat.stanford.edu/~jhf/#reports>
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning; Data Mining, Inference and Prediction*. Springer, New York.
- JIANG, W. (2004). Process consistency for AdaBoost (with discussion). *The Annals of Statistics* **32** 13–29 (disc. pp. 85–134).
- KEARNS, M. and VALIANT, L. G. (1989). Cryptographic limitations on learning Boolean formulae and finite automata. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*. ACM Press, New York.
- LOZANO, A., KULKARNI, S. and SCHAPIRE, R. (2006). Convergence and consistency of regularized boosting algorithms with stationary β -mixing observations. In *Advances in Neural Information Processing Systems* (Y. Weiss, B. Schölkopf and J. Platt, eds.), vol. 18. MIT Press.
- LUGOSI, G. and VAYATIS, N. (2004). On the Bayes-risk consistency of regularized boosting methods (with discussion). *The Annals of Statistics* **32** 30–55 (disc. pp. 85–134).
- LUTZ, R. and BÜHLMANN, P. (2006). Boosting for high-multivariate responses in high-dimensional linear regression. *Statistica Sinica* **16** 471–494.
- LUTZ, R., KALISCH, M. and BÜHLMANN, P. (2008). Robustified L2 boosting. *Computational Statistics & Data Analysis* **52** 3331–3341.

- MASON, L., BAXTER, J., BARTLETT, P. and FREAN, M. (2000). Functional gradient techniques for combining hypotheses. In *Advances in Large Margin Classifiers* (A. Smola, P. Bartlett, B. Schölkopf and D. Schuurmans, eds.). MIT Press, Cambridge.
- RÄTSCH, G., ONODA, T. and MÜLLER, K. (2001). Soft margins for AdaBoost. *Machine Learning* **42** 287–320.
- RÄTSCH, G. and WARMUTH, M. (2005). Efficient margin maximizing with boosting. *J. Machine Learning Research* **6** 2131–2152.
- ROSSET, S., ZHU, J. and HASTIE, T. (2004). Boosting as a regularized path to a maximum margin classifier. *Journal of Machine Learning Research* **5** 941–973.
- RUDIN, X., SCHAPIRE, R. E. and DAUBECHIES, I. (2007). Analysis of boosting algorithms using the smooth margin function. *The Annals of Statistics* **35** 2723–2768.
- SCHAPIRE, R. (1990). The strength of weak learnability. *Machine Learning* **5** 197–227.
- SCHAPIRE, R., FREUND, Y., BARTLETT, P. and LEE, W. (1998). Boosting the margin: a new explanation for the effectiveness of voting methods. *The Annals of Statistics* **26** 1651–1686.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B* **58** 267–288.
- TUKEY, J. (1977). *Exploratory Data Analysis*. Addison-Wesley, Reading, MA.
- WARMUTH, M., GLOCER, K. and RÄTSCH, G. (2008a). Boosting algorithms for maximizing the soft margin. In *Advances in Neural Information Processing Systems 20* (J. Platt, D. Koller, Y. Singer and S. Roweis, eds.). MIT Press.
- WARMUTH, M., GLOCER, K. and VISHWANATHAN, S. (2008b). Entropy regularized LP Boost. In *Advanced Lectures on Machine Learning* (Y. Freund, L. Györfi, G. Turán and T. Zeugmann, eds.). Lecture Notes in Computer Science, Springer.
- ZHANG, T. (2008). Adaptive forward-backward greedy algorithm for learning sparse representations. Tech. rep., Rutgers Statistics Department. A short version appeared in NIPS 08.

ZHANG, T. and YU, B. (2005). Boosting with early stopping: convergence and consistency. *The Annals of Statistics* **33** 1538–1579.

ZHAO, P. and YU, B. (2007). Stagewise Lasso. *Journal of Machine Learning Research* **8** 2701–2726.