

Twin Boosting

Improved feature selection and prediction

Journal Article**Author(s):**

Buehlmann, Peter; Hothorn, Torsten

Publication date:

2010

Permanent link:

<https://doi.org/10.3929/ethz-b-000017568>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Statistics and Computing 20(2), <https://doi.org/10.1007/s11222-009-9148-5>

Twin Boosting: improved feature selection and prediction

Peter Bühlmann · Torsten Hothorn

Received: 15 January 2009 / Accepted: 10 August 2009 / Published online: 3 September 2009
© Springer Science+Business Media, LLC 2009

Abstract We propose Twin Boosting which has much better feature selection behavior than boosting, particularly with respect to reducing the number of false positives (falsely selected features). In addition, for cases with a few important effective and many noise features, Twin Boosting also substantially improves the predictive accuracy of boosting. Twin Boosting is as general and generic as (gradient-based) boosting. It can be used with general weak learners and in a wide variety of situations, including generalized regression, classification or survival modeling. Furthermore, it is computationally feasible for large problems with potentially many more features than observed samples. Finally, for the special case of orthonormal linear models, we prove equivalence of Twin Boosting to the adaptive Lasso which provides some theoretical aspects on feature selection with Twin Boosting.

Keywords Classification · Gradient descent · High-dimensional data · Regression · Regularization

1 Introduction

Boosting has attracted much attention in the machine learning community (cf. Schapire 2002; Meir and Rätsch 2003, and the references therein) as well as in statistics

(Breiman 1998, 1999; Ridgeway 1999; Friedman et al. 2000; Bühlmann and Yu 2003; Tutz and Binder 2006; Tutz and Reithinger 2007; Bühlmann and Hothorn 2007), mainly because of its excellent performance and computational attractiveness for large datasets. The main breakthrough came with Freund and Schapire's most successful AdaBoost algorithm for binary classification (Freund and Schapire 1996, 1997). Nowadays, boosting is still among the most powerful methods for prediction, feature selection and regularization in high-dimensional problems.

Twin boosting is a very generic boosting-based method which most often yields better feature or variable selection than boosting while keeping or even increasing the prediction accuracy. Twin boosting proceeds roughly as follows. A first round of "classical" boosting is done (the first twin); then, in a second round, another boosting process is run (the second twin) which is forced to resemble the one from the first round. Thus, the two rounds are similar, like twins, and hence the name Twin Boosting.

There have been some attempts to make boosting or also related Lasso-methods (Tibshirani 1996) more powerful, in particular in terms of feature or variable selection but also for prediction. The potential for improvement is mainly given for cases with many ineffective and a few effective covariates. To deal with many ineffective features, a strong regularization is employed in boosting or in related Lasso-methods, creating a large estimation bias. Proposals to effectively avoid such large biases include Sparse Boosting (Bühlmann and Yu 2006), conjugate direction boosting (Lutz and Bühlmann 2006), Lasso with relaxation (Meinshausen 2007) or the adaptive Lasso (Zou 2006). Our approach has some similarity to the latter as it encompasses the adaptive Lasso in the very special case of an orthonormal linear model. An aspect of success of some of these methods is their greater flexibility than what is possible with a

P. Bühlmann (✉)
Seminar für Statistik, ETH Zürich, 8092 Zürich, Switzerland
e-mail: buehlmann@stat.math.ethz.ch

T. Hothorn
Institut für Statistik, Ludwig-Maximilians-Universität, 80539
München, Germany
e-mail: Torsten.Hothorn@stat.uni-muenchen.de

single regularization parameter in boosting (the number of iterations) or in related Lasso-methods (the penalty parameter). Our new Twin Boosting involves two tuning parameters which can be chosen sequentially rather than simultaneously optimizing over a two-dimensional space of regularization parameters.

All these recently proposed methods mentioned above, aiming to improve over boosting or Lasso, are far less general and generic than our Twin Boosting approach. The latter can be easily used with any real-valued weak learner which is a (possibly crude) estimator of the conditional mean function. In particular, Twin Boosting can be easily used with trees and hence, it can be applied to data of mixed types with continuous, ordinal and categorical features. Secondly, Twin Boosting can be used for a rich class of loss functions, including squared and absolute error for regression, logistic or exponential loss for classification or Poisson-loss for count data. Therefore, Twin Boosting is essentially as general and generic as boosting.

2 Boosting algorithms

Boosting is used in supervised learning from data $(X_i, Y_i), \dots, (X_n, Y_n)$, where $X_i \in \mathcal{X}$ is a p -dimensional predictor variable and $Y_i \in \mathcal{Y}$ is a univariate response variable. The space of predictor variables \mathcal{X} is often a subset of \mathbb{R}^p and the space of response variables often a subspace of \mathbb{R} , e.g. $\mathcal{Y} = \{-1, +1\}$ for binary classification or $\mathcal{Y} = \mathbb{R}$ for Gaussian regression.

A boosting algorithm needs the specification of a loss function and a weak learner. Regarding the former, consider a loss function

$$\rho : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}$$

(or a subset of \mathbb{R} ; and we exclude here the case where the range of $\rho(\cdot, \cdot)$ is multivariate). The loss function is assumed to be differentiable (almost everywhere) and typically convex with respect to the second argument. Examples include squared error loss $\rho(y, f) = |y - f|^2/2$ with $y \in \mathbb{R}$ for regression or the logistic loss $\rho(y, f) = \log_2(1 + \exp(-2yf))$ with $y \in \{-1, +1\}$ for binary classification. The weak learner is in our setting a real-valued function estimator:

$$(X_1, U_1), \dots, (X_n, U_n) \xrightarrow{\text{weak learner}} \hat{g}(\cdot),$$

where X_i is the p -dimensional predictor variable and $U_i \in \mathbb{R}$ a pseudo-response variable. The notion of a pseudo-response variable will become clear in the description of the generic boosting algorithm below. An example of a weak learner is a regression tree yielding a regression function estimate $\hat{g}(\cdot)$.

It is instructive to look at the population minimizer of the loss function:

$$f^*(\cdot) = \arg \min_{f(\cdot)} \mathbb{E}[\rho(Y, f(X))]. \quad (1)$$

It is itself a function and it is the target of the boosting algorithm. For example, the squared error loss $\rho(y, f) = |y - f|^2/2$ yields the well-known population minimizer $f^*(x) = \mathbb{E}[Y|X = x]$ and boosting algorithms using the squared error loss are estimators of the regression function $f^*(\cdot)$.

2.1 The generic boosting algorithm

Boosting is based on the empirical risk $n^{-1} \sum_{i=1}^n \rho(Y_i, f(X_i))$ and pursuing iterative steepest descent in function space as described below for estimating the unknown function $f^*(\cdot)$. This very general and useful view of boosting has been pioneered by Breiman (1998, 1999) and further developed by Friedman et al. (2000), Rätsch et al. (2001) and Friedman (2001).

Variable or feature selection is pursued in this paper without any statistical significance testing. The selected features is the set of variables which enter explicitly (via the boosting algorithm) in the final function estimate $\hat{f}(\cdot)$.

Generic boosting algorithm

1. Initialize $\hat{f}^{[0]}$: typical values are $\hat{f}^{[0]} \equiv \bar{Y} = n^{-1} \sum_{i=1}^n Y_i$ or $\hat{f}^{[0]} \equiv 0$. Set $m = 0$.
2. Increase m by 1. Compute negative gradient $-\frac{\partial}{\partial f} \rho(Y, f)$ and evaluate at $\hat{f}_{m-1}(X_i)$:

$$U_i = -\frac{\partial}{\partial f} \rho(Y, f)|_{f=\hat{f}_{m-1}(X_i)}, \quad i = 1, \dots, n.$$

3. Fit negative gradient vector U_1, \dots, U_n by the weak learner

$$(X_1, U_1), \dots, (X_n, U_n) \xrightarrow{\text{weak learner}} \hat{g}^{[m]}(\cdot).$$

Thus, $\hat{g}^{[m]}(\cdot)$ can be viewed as an approximation of the negative gradient vector.

4. Update $\hat{f}^{[m]} = \hat{f}^{[m-1]}(\cdot) + \nu \cdot \hat{g}^{[m]}(\cdot)$, where $0 < \nu \leq 1$ is a step-length, i.e. proceed along an estimate of the negative gradient vector.
5. Iterate steps 2–4 until $m = m_{\text{stop}}$ for some stopping iteration m_{stop} .

The stopping iteration, which is the main tuning parameter, can be estimated via cross-validation. The choice of the step-size ν in step 4. is of minor importance, as long as it is “small” such as $\nu = 0.1$. A smaller value of ν typically requires a larger number of boosting iterations, and thus more computing time, while the predictive accuracy has

been found to be potentially better for ν being “sufficiently small”, e.g. $\nu = 0.1$ (Friedman 2001). Friedman (2001) suggests to use an additional line search between steps 3 and 4 (in case of different loss functions $\rho(\cdot, \cdot)$ than squared error): it yields a slightly different algorithm but the additional line search seems unnecessary to pursue for achieving a good estimator $\hat{f}^{[m_{\text{stop}}]}(\cdot)$.

2.2 L_2 Boosting

L_2 Boosting is the generic boosting algorithm when using the squared error loss $\rho(y, f) = |y - f|^2/2$. It has been proposed by Friedman (2001) and further developed and analyzed in Bühlmann and Yu (2003).

L_2 Boosting algorithm

1. Initialize $\hat{f}^{[0]}$: the default value is $\hat{f}^{[0]} \equiv \bar{Y} = n^{-1} \sum_{i=1}^n Y_i$. Set $m = 0$.
2. Increase m by 1. Compute the residuals $U_i = Y_i - \hat{f}^{[m-1]}(X_i)$ for $i = 1, \dots, n$.
3. Fit the residual vector U_1, \dots, U_n by the weak learner

$$(X_1, U_1), \dots, (X_n, U_n) \xrightarrow{\text{weak learner}} \hat{g}^{[m]}(\cdot).$$

4. Update $\hat{f}^{[m]}(\cdot) = \hat{f}^{[m-1]}(\cdot) + \nu \cdot \hat{g}^{[m]}(\cdot)$, where $0 < \nu \leq 1$ is a step-length factor.
5. Iterate steps 2 to 4 until $m = m_{\text{stop}}$ for some stopping iteration m_{stop} .

The stopping iteration m_{stop} is the main tuning parameter which can be selected using cross-validation.

Other loss functions, e.g. the logistic loss for classification, and corresponding boosting algorithms are described in Sect. 5. For reasons of clarity, we will first describe all the ideas and motivation for Twin Boosting for the case of the squared error loss. We will then show in Sect. 5 that the concepts easily generalize to general loss functions.

3 Twin L_2 Boosting for linear models

For expository simplicity, we first describe Twin Boosting for linear models:

$$Y_i = \sum_{j=1}^p \beta_j X_i^{(j)} + \varepsilon_i, \tag{2}$$

where $\varepsilon_1, \dots, \varepsilon_n$ are independent, identically distributed (i.i.d.), independent from X_1, \dots, X_n , with $\mathbb{E}[\varepsilon_i] = 0$. We sometimes write in short:

$$\mathbf{Y} = \mathbf{X}\beta + \varepsilon$$

with $\mathbf{Y}_{n \times 1} = (Y_1, \dots, Y_n)^T$, $\varepsilon_{n \times 1} = (\varepsilon_1, \dots, \varepsilon_n)^T$ and $\mathbf{X}_{n \times p} = [X_1, \dots, X_n]^T$.

Fitting of the linear model in (2) with boosting or Twin Boosting is based on the squared error loss (L_2 Boosting) and the componentwise linear least squares weak learner, as described in Sect. 3.1 below. Already L_2 Boosting itself has been proven to be very useful for fitting linear models with potentially many more covariates than samples (Friedman 2001; Bühlmann 2006; Bühlmann and Yu 2006) and we will argue here in which circumstances Twin Boosting is even better.

In the sequel, we will extensively use the following notation:

$$\langle u, v \rangle = \sum_{i=1}^n u_i v_i \quad \text{for some vectors } u, v \in \mathbb{R}^n,$$

$$\|u\|^2 = \langle u, u \rangle = \sum_{i=1}^n u_i^2 \quad \text{for some vector } u \in \mathbb{R}^n.$$

Moreover, denote by $\mathbf{X}^{(j)}$ the j th $n \times 1$ column vector of \mathbf{X} .

3.1 Componentwise linear least squares as weak learner

Consider the following weak learner based on data $(X_1, U_1), \dots, (X_n, U_n)$:

$$\hat{g}(x) = \hat{\gamma}_{\hat{S}} x^{(\hat{S})},$$

$$\hat{\gamma}_j = \langle \mathbf{U}, \mathbf{X}^{(j)} \rangle / \|\mathbf{X}^{(j)}\|^2, \tag{3}$$

$$\hat{S} = \arg \min_{1 \leq j \leq p} \sum_{i=1}^n (U_i - \hat{\gamma}^{(j)} X_i^{(j)})^2.$$

It selects and fits the best predictor variable in a simple linear model in the sense of ordinary least squares fitting. For computational implementation as well as for the construction our Twin Boosting, it is useful to represent the selected predictor variable as:

$$\hat{S} = \arg \max_{1 \leq j \leq p} |\langle \mathbf{U}, \mathbf{X}^{(j)} \rangle|^2 / \|\mathbf{X}^{(j)}\|^2. \tag{4}$$

The derivation of (4) is straightforward.

When using L_2 Boosting with this base procedure, we select in every iteration one predictor variable, not necessarily a different one for each iteration, and we up-date the function $\hat{f}^{[m]}(\cdot)$ linearly as described in Sect. 2.2. For componentwise linear least squares, we can simply up-date the parameter vector of a linear model.

L_2 Boosting with componentwise linear least squares

1. Initialize $\hat{\beta}^{[0]}$. Set $m = 0$.
2. Increase m by 1. Compute the residuals $U_i = Y_i - \bar{Y} - (\mathbf{X}\hat{\beta}^{[m-1]})_i$ for $i = 1, \dots, n$.
3. Compute \hat{S}_m as in (4) and $\hat{\gamma}_{\hat{S}_m}$ as in (3).

4. Update

$$\hat{\beta}^{[m]} = \hat{\beta}^{[m-1]} + v \cdot \hat{\gamma}_{\hat{S}_m}.$$

The notation should be read that only the \hat{S}_m th component of the coefficient estimate is up-dated.

- Iterate steps 2 to 4 until $m = m_{\text{stop}}$ for some stopping iteration m_{stop} .

The selected variables are corresponding to the indices j for which $\hat{\beta}^{[m_1]} \neq 0$.

3.2 Twin L_2 Boosting with componentwise linear least squares

Twin Boosting for linear models uses the squared error loss and componentwise linear least squares as weak learner. For expository simplicity, we consider standardized data where $n^{-1} \sum_{i=1}^n Y_i = 0$, $n^{-1} \sum_{i=1}^n X_i^{(j)} = 0$ and $\sum_{i=1}^n (X_i^{(j)})^2 = 1$ for all $j = 1, \dots, p$ (for unstandardized data, the algorithm has to be slightly re-formulated). This can always be achieved by centering with the empirical mean and scaling with the empirical standard deviation.

Twin L_2 Boosting with componentwise linear least squares

- Run a first round of L_2 Boosting with componentwise linear least squares, using m_1 iterations. Denote the estimated parameter by $\hat{\beta}_{\text{init}}^{[m_1]}$.
- For the second round, run L_2 Boosting with componentwise linear least squares but replace formula (4) by

$$\hat{S} = \arg \max_{1 \leq j \leq p} |(U, \mathbf{X}^{(j)})| |\hat{\beta}_{\text{init};j}^{[m_1]}|^2. \quad (5)$$

Note that $\|X^{(j)}\|^2 = 1$ because of standardized data. Use m_2 iterations and denote the Twin L_2 Boosting estimator by $\hat{\beta}_{\text{TWB}}^{[m_2]}$.

The motivation of the multiplier $|\hat{\beta}_{\text{init};j}^{[m_1]}|^2$ in formula (5) is as follows: if a prediction variable is important, it has a larger multiplier $|\hat{\beta}_{\text{init};j}^{[m_1]}|^2$ and hence, it is more likely to be selected in the criterion (5); and vice-versa. In particular, a non-relevant variable with $\hat{\beta}_{\text{init};j}^{[m_1]} = 0$ will not be selected. We will describe in Proposition 1 that for special cases, the selection with the multiplier in (5) is equivalent to the adaptive Lasso, a method which has shown remarkable success for high-dimensional feature selection (Zou 2006; Huang et al. 2008; Bühlmann and Meier 2008). Two tuning parameters m_1 and m_2 are involved. Instead of optimizing (e.g. cross-validating) over both parameters simultaneously, we use the following, computationally much faster sequential approach: first, an estimate \hat{m}_1 is obtained from cross-validation for L_2 Boosting and then, cross-validation for Twin Boosting with fixed \hat{m}_1 in the initial estimator yields an estimate \hat{m}_2 .

3.3 Connections to the Lasso and the adaptive Lasso

There is an intriguing connection between L_2 Boosting with componentwise linear least squares and the Lasso (Tibshirani 1996). The latter is an ℓ^1 -penalty method for regression defined by

$$\hat{\beta}(\lambda) = \arg \min_{\beta} n^{-1} \sum_{i=1}^n \left(Y_i - \beta_0 - \sum_{j=1}^p \beta_j X_i^{(j)} \right)^2 + \lambda \sum_{j=1}^p |\beta_j|. \quad (6)$$

Efron et al. (2004) consider a version of L_2 Boosting, called forward stagewise linear regression (FSLR), and they show that FSLR with infinitesimally small step-sizes (i.e. the value v) produces a set of solutions which is approximately equivalent to the set of Lasso solutions when varying the regularisation parameter λ in Lasso (see also (6) above). The approximate equivalence is derived by representing FSLR and Lasso as two different modifications of their computationally efficient least angle regression (LARS) algorithm. In special cases where the design matrix satisfies a “positive cone condition”, FSLR, Lasso and LARS all coincide (Efron et al. 2004, p. 425).

Despite the fact that L_2 Boosting and Lasso are not equivalent methods in general, the connection between boosting (as a forward, greedy method) and the Lasso (involving convex optimization) is interesting.

Recently, Zou (2006) has proposed the adaptive Lasso, defined as

$$\hat{\beta}(\lambda) = \arg \min_{\beta} \left(n^{-1} \sum_{i=1}^n \left(Y_i - \beta_0 - \sum_{j=1}^p \beta_j X_i^{(j)} \right)^2 + \lambda \sum_{j=1}^p \frac{|\beta_j|}{|\hat{\beta}_{\text{init};j}|} \right), \quad (7)$$

where β_{init} is an initial estimator. Zou (2006) mainly uses ordinary least squares as initial estimator (for cases with reasonable ratio n/p) and he mentions the Ridge estimator as one among several possible alternatives. In addition, he proposed a more general class of estimators, but the specific form in (7) is useful and often sufficient. By defining $0/0 = 0$, the adaptive Lasso is well-defined even if $\hat{\beta}_{\text{init};j} = 0$ for some j : it implies that $\hat{\beta}_j(\lambda) = 0$. The adaptive Lasso has two advantages over the Lasso. It yields consistent variable selection without imposing severe restriction on the design matrix, at least for the case with fixed predictor dimension p (Zou 2006), whereas the Lasso is inconsistent (typically yields too large models) if the design is roughly speaking “strongly correlated” (Meinshausen and

Bühlmann 2006; Zou 2006; Zhao and Yu 2006). Secondly, adaptive Lasso yields better predictions if the true underlying model has many ineffective (noise) predictor variables. Both of these advantages are closely related to the motivation of our Twin Boosting, see Sect. 1.

In case of an orthonormal linear model, i.e. the model (2) with $\sum_{i=1}^n X_i^{(j)} X_i^{(k)} = \delta_{jk}$ (Kronecker $\delta_{jk} = 1$ if $j = k$ and 0 otherwise), explicit connections between boosting and Lasso exist. It has been shown (constructively) that L_2 Boosting with componentwise linear least squares approximates the solution from Lasso (as $\nu > 0$ tends to zero) which equals the soft-threshold estimator (Bühlmann and Yu 2006). For Twin Boosting, the following holds.

Proposition 1 *Consider an orthonormal linear model and any sequence of regularization parameters $(\lambda_k)_{k \in \mathbb{N}}$ for the adaptive Lasso in (7). Then, for Twin L_2 Boosting with componentwise linear least squares using a step-size ν , there exists a Twin Boosting iteration m_2 , typically depending on λ_k, ν and the data, such that*

$$\hat{\beta}_{TWB; j}^{[m_2]} = \hat{\beta}_{TWB; j}^{[m_2]}(\nu) = \hat{\beta}_{AdaLasso, j}(\gamma_j)$$

$$\text{as in (7) with } \beta_{init} = \hat{\beta}_{init}^{[m_1]},$$

$$\gamma_j = \lambda_k(1 + e_j(\nu)), \quad \max_{1 \leq j \leq n} |e_j(\nu)| \rightarrow 0 \quad (\nu \rightarrow 0).$$

A proof is given in the Appendix. We remark that the adaptive Lasso in the orthonormal linear model is of the form

$$\hat{\beta}_{AdaLasso, j} = \text{sign}(Z_j) \left(\left| Z_j \right| - \frac{\lambda}{2|\hat{\beta}_{init, j}^{[m_1]}|} \right)_+,$$

where $(x)_+ = \max(0, x)$ is the positive part and $Z_j = (\mathbf{X}^T \mathbf{Y})_j$. Together with Proposition 1, this illustrates that in the simple case of an orthonormal linear model, Twin Boosting yields in the limiting case with $\nu \rightarrow 0$ the adaptive Lasso which equals the adaptive soft-threshold estimator. The connection is interesting but the real power of Twin Boosting is its generic applicability to very general weak learners and loss functions.

4 General Twin L_2 Boosting with trees and general weak learners

The most popular weak learners for boosting are decision trees. Our proposal for Twin Boosting with trees is simple, easy to implement and effective. We could represent trees in terms of basis functions (with indicator functions of rectangles) and then employ Twin Boosting methodology as for linear models. However, such an approach becomes computationally very impractical due to the huge number of basis

functions (particularly for larger trees) and in addition, such an approach would (adaptively) encourage sparseness in the space of basis functions rather than sparseness in the space of predictor variables; the latter is often much more interesting for many datasets and scientific problems.

Twin Boosting with any real-valued (regression-type) weak learner is defined as follows.

General Twin L_2 Boosting with general weak learner

1. Run a first round of L_2 Boosting and denote by $\hat{\mathbf{f}}_{init}^{[m_1]} = (\hat{f}_{init}^{[m_1]}(X_1), \dots, \hat{f}_{init}^{[m_1]}(X_n))$ the fitted function at the data points and by $\hat{\mathcal{V}}^{[m_1]} \subseteq \{1, \dots, p\}$ the subset of indices corresponding to selected predictor variables, both based on m_1 boosting iterations. (If the weak learner is not doing any variable selection, then $\hat{\mathcal{V}}^{[m_1]} = \{1, \dots, p\}$ is the full set.)
2. For the second round, initialize $\hat{f}^{[0]}$: the default value is $\hat{f}^{[0]} \equiv \bar{Y} = n^{-1} \sum_{i=1}^n Y_i$. Set $m = 0$.
3. Increase m by 1. Compute the residuals $U_i = Y_i - \hat{f}^{[m-1]}(X_i)$ for $i = 1, \dots, n$.
4. For every subset $\mathcal{W} \subseteq \hat{\mathcal{V}}^{[m_1]}$,¹ fit the residual vector U_1, \dots, U_n to $X_1^{\mathcal{W}}, \dots, X_n^{\mathcal{W}}$ with the weak learner; here $X^{\mathcal{W}}$ denotes $\{X^{(j)}; j \in \mathcal{W}\}$. Denote this fitted function by $\hat{h}_{\mathcal{W}}(\cdot)$ and by $\hat{\mathbf{h}}_{\mathcal{W}} = (\hat{h}_{\mathcal{W}}(X_1), \dots, \hat{h}_{\mathcal{W}}(X_n))$. Then, choose the best \mathcal{W} according to:

$$\hat{\mathcal{W}} = \arg \max_{\mathcal{W}} C_{\mathcal{W}}^2(2\langle \mathbf{U}, \hat{\mathbf{h}}_{\mathcal{W}} \rangle - \|\hat{\mathbf{h}}_{\mathcal{W}}\|^2),$$

$$C_{\mathcal{W}} = \langle \hat{\mathbf{f}}_{init}^{[m_1]} - \overline{\hat{f}_{init}^{[m_1]}}, \hat{\mathbf{h}}_{\mathcal{W}} \rangle / \|\hat{\mathbf{h}}_{\mathcal{W}}\|, \tag{8}$$

$$\overline{\hat{f}_{init}^{[m_1]}} = n^{-1} \sum_{i=1}^n \hat{f}_{init}^{[m_1]}(X_i).$$

Denote by $\hat{g}^{[m]}(\cdot) = \hat{h}_{\hat{\mathcal{W}}}(\cdot)$.

5. Update $\hat{f}^{[m]} = \hat{f}^{[m-1]}(\cdot) + \nu \cdot \hat{g}^{[m]}(\cdot)$, where $0 < \nu \leq 1$ is a step-length factor.
6. Iterate steps 3 to 5 until $m = m_2$ for some stopping iteration m_2 .

The stopping iteration m_2 is the main tuning parameter (for a given boosting estimator of the first round) which can be selected using cross-validation.

From a computational point of view, it seems awkward to consider all subsets $\mathcal{W} \subseteq \hat{\mathcal{V}}^{[m_1]}$ in step 4. However, if the weak learner does variable selection, selecting at most d predictor variables (e.g. a tree with at most $d + 1$ terminal nodes), we only have to consider in step 4 all subsets \mathcal{W} having cardinality $|\mathcal{W}| = d$. For example, when using stumps, $d = 1$ and step 4 becomes:

¹See also below the modification with random feature subsets.

Step 4 for stumps. For every $j \in \hat{\mathcal{V}}^{[m_1]}$, fit the residual vector with a stump and denote it by $\hat{h}_j(\cdot)$. The remaining part of step 4 is then as above.

The selected variables from Twin Boosting arise from the set of variables which occur in $\hat{\mathcal{W}}$ from step 4 for at least one iteration. For example with trees as weak learner, the selected variables arise from variable selection of the tree-type weak learner during the Twin Boosting iterations.

Using trees as weak learner, $d = 1$ (stumps) or $d = 2$ is often a good choice (see also the discussion below in Sect. 4.2). More generally, for a tree with $d + 1$ terminal nodes (d splits), the computation in step 4 is of the order $O(dnp|\hat{\mathcal{V}}^{[m_1]}|^d)$. It should be noted here that usually $|\hat{\mathcal{V}}^{[m_1]}| \leq \min(n, p)$, often even $\ll \min(n, p)$.

In case of large trees as weak learner, we propose the following:

General Twin L_2 Boosting with random feature subsets

Step 4 with random feature subsets. Generate B independent random feature subsets $\mathcal{W} \subseteq \hat{\mathcal{V}}^{[m_1]}$ with cardinality $|\mathcal{W}| = s_{\mathcal{W}}$ where the elements of \mathcal{W} are sampled i.i.d. $\sim \text{Uniform}(\hat{\mathcal{V}}^{[m_1]})$. Then proceed as in step 4 above.

The value $s_{\mathcal{W}}$ is typically chosen to be of the magnitude of the depth of the tree learner and we use $B = 500$ as default value for the number of random subsets. We also emphasize that, unlike in Random Forests (Breiman 2001), there is no averaging operation involved over the weak learners (e.g. trees) resulting from the fits with the random feature subsets: here, the best (random) weak learner is selected according to the criterion in (8).

4.1 Relation to Twin L_2 Boosting for linear models

We give now a motivation for the construction in step 4. Consider Twin L_2 Boosting for linear models, as described in Sect. 3.2. There, it happens automatically that the second round of Twin Boosting considers only the set of predictor variables $\hat{\mathcal{V}}^{[m_1]}$ which has been chosen in the first round. Next, we consider the formula (8). The residual sum of squares is

$$\begin{aligned} \|\mathbf{U} - \hat{\mathbf{h}}_{\mathcal{W}}\|^2 &= \|\mathbf{U}\|^2 - 2\langle \mathbf{U}, \hat{\mathbf{h}}_{\mathcal{W}} \rangle + \|\hat{\mathbf{h}}_{\mathcal{W}}\|^2 \\ &= \text{const.} - (2\langle \mathbf{U}, \hat{\mathbf{h}}_{\mathcal{W}} \rangle - \|\hat{\mathbf{h}}_{\mathcal{W}}\|^2). \end{aligned}$$

L_2 Boosting would proceed by choosing the best \mathcal{W} maximizing

$$G(\mathcal{W}) = 2\langle \mathbf{U}, \hat{\mathbf{h}}_{\mathcal{W}} \rangle - \|\hat{\mathbf{h}}_{\mathcal{W}}\|^2.$$

For Twin Boosting, we want to multiply weights $C_{\mathcal{W}}$ into the criterion $G(\mathcal{W})$. For the form of these weights, it is

instructive to consider Twin L_2 Boosting for linear models with normed predictor variables having $\|\mathbf{X}^{(j)}\|^2 = 1$: there, \mathcal{W} is an element of $\{1, \dots, p\}$ and $\hat{h}_j(X_i) = \langle \mathbf{U}, \mathbf{X}^{(j)} \rangle X_i^{(j)}$. Using this, we easily obtain for $\mathcal{W} = j$,

$$G(\mathcal{W}) = G(j) = |\langle \mathbf{U}, \mathbf{X}^{(j)} \rangle|^2. \tag{9}$$

Twin Boosting for linear models does nothing else than multiplying $G(j)$ by the weights $\beta_{\text{init},j}^2$. Therefore, we want to multiply $G(\mathcal{W})$ in (9) with the square of a suitable regression coefficient. Our $C_{\mathcal{W}}$ in formula (8) is the standardized regression coefficient when regressing the fitted function from the first round of Twin Boosting $\hat{\mathbf{f}}^{[m_1]}$ against the candidate estimate $\hat{\mathbf{h}}_{\mathcal{W}}$; the standardization is a multiplication by $\|\hat{\mathbf{h}}_{\mathcal{W}}\|$. The standardization is useful as it implicitly measures the regression coefficient on a scale where $\hat{\mathbf{h}}_{\mathcal{W}}$ would have been standardized to $\|\hat{\mathbf{h}}_{\mathcal{W}}\| \equiv 1$ for all \mathcal{W} . We end our motivation of step 4 with the following result.

Proposition 2 Consider a linear model as in (2) with $\sum_{i=1}^n X_i^{(j)} = 0$, $\|\mathbf{X}^{(j)}\| = 1$ for all $j = 1, \dots, p$. Then, for the general Twin L_2 Boosting algorithm with the component-wise linear least squares weak learner:

1. If the underlying regression model is orthonormal with $\sum_{i=1}^n X_i^{(j)} X_i^{(k)} = \delta_{jk}$, this algorithm coincides exactly with Twin L_2 Boosting for linear models, as described in Sect. 3.2.
2. In general, this algorithm equals a modified Twin L_2 Boosting algorithm for linear models, as described in Sect. 3.2, but where we replace $|\hat{\beta}_{\text{init},j}^{[m]}|^2$ in (5) by $|\sum_{k=1}^p \hat{\beta}_{\text{init},k}^{[m]} \langle \mathbf{X}^{(k)}, \mathbf{X}^{(j)} \rangle|^2$.

A proof is given in the Appendix. We have discussed in Sect. 3.3 some connections to the adaptive Lasso, generally defined as

$$\begin{aligned} \hat{\beta}(\lambda) = \arg \min_{\beta} & \left(n^{-1} \sum_{i=1}^n \left(Y_i - \beta_0 - \sum_{j=1}^p \beta_j X_i^{(j)} \right)^2 \right. \\ & \left. + \lambda \sum_{j=1}^p w_j |\beta_j| \right). \end{aligned}$$

Here, we consider the weights

$$w_j^{-1} = \left| \sum_{k=1}^p \hat{\beta}_{\text{init},k}^{[m]} \langle \mathbf{X}^{(k)}, \mathbf{X}^{(j)} \rangle \right| \tag{10}$$

(instead of $w_j^{-1} = |\beta_{\text{init},j}|$ as in (7)). Denote by $W = \text{diag}(w_1, \dots, w_p)$ and $\mathbf{X}_W = \mathbf{X}W^{-1}$ the new $n \times p$ design matrix. Then, it is straightforward to see that L_2 Boosting based on the design \mathbf{X}_W and the general Twin L_2 Boosting algorithm based on the original design \mathbf{X} , as in statement

2 of Proposition 2, coincide (up to the fact that regression coefficients need to be scaled). Likewise and easy to see, the Lasso based on design \mathbf{X}_W and the adaptive Lasso, with weights as in (10) and based on the original design \mathbf{X} , are equivalent (up to scaling of the regression coefficients). Therefore, whenever Boosting and Lasso, both based on \mathbf{X}_W are equivalent, the general Twin L_2 Boosting and the adaptive Lasso, both based on the original design \mathbf{X} , are equivalent. This holds, for example, if \mathbf{X}_W is orthonormal: more generally, there is an equivalence of forward stage-wise linear regression (FSLR) and Lasso if \mathbf{X}_W satisfies the positive cone condition (Efron et al. 2004). We will report in Sect. 6.1.3 some empirical results comparing Twin L_2 Boosting for linear models as in Sect. 3.2 (formula (5) and general Twin L_2 Boosting, as arising in assertion 2 of Proposition 2.

4.2 Stumps and larger trees as weak learners

The choice of the weak learner is usually driven by optimizing the prediction performance. In addition, some structural properties can be another useful criterion as well.

The generic boosting and Twin Boosting estimator is a linear combination of weak learners

$$\hat{f}^{[m]}(\cdot) = v \sum_{k=1}^m \hat{g}^{[k]}(\cdot).$$

Therefore, structural properties of the boosting function estimator are given by linear combination of structural characteristics of the weak learner.

Trees are among the most popular base procedures in machine learning. They have the advantage to be invariant under monotone transformations of predictor variables, i.e., we do not need to search for good data transformations.

When using stumps, i.e., a tree with two terminal nodes, the boosting and Twin Boosting estimate will be an additive model in the original predictor variables, because every stump-estimate is a function of a single predictor variable only. Similarly, boosting trees with (at most) $d + 1$ terminal nodes results in a nonparametric model having at most interactions of order $d - 1$: e.g. for $d = 2$, we would pick up interaction terms between pairs of predictor variables. Thus, if we want to constrain the degree of interactions, we can easily do this by constraining the (maximal) number of nodes in the tree learner. For many real datasets, it seems that low-order interaction (or even additive) models are sufficiently rich for good prediction and interpretation. For example, the naive Bayes classifier or linear discriminant analysis, based on an additive or linear decision function respectively, works surprisingly well in many applications (Jamain and Hand 2005; Hand 2006). Also boosting with stumps, yielding an additive model, has proven to be successful in many areas,

e.g. winning the performance prediction challenge of the IEEE World Congress on Computational Intelligence 2006 (Lutz 2006). Thus, we often get good performance with trees having 2 or 3 terminal nodes ($d = 1$ or 2, respectively). With such small values of d , Twin Boosting is computationally fast, as discussed after the description of the Twin Boosting algorithm in Sect. 4.

5 Other loss functions and generic Twin Boosting

For other loss functions than squared error (i.e. other boosting algorithms than L_2 Boosting) we can use the general functional gradient descent approach as described in Sect. 2.1.

Interesting examples of loss functions include the following. For binary classification with $y \in \{-1, +1\}$, the logistic loss is

$$\rho_{\text{logit}}(y, f) = \log_2(1 + \exp(-2yf)), \quad (11)$$

and the exponential loss is

$$\rho_{\text{exp}}(y, f) = \exp(-yf). \quad (12)$$

Boosting with the logistic loss or exponential loss function is essentially LogitBoost (Friedman et al. 2000) (also called BinomialBoosting) or AdaBoost (Freund and Schapire 1996), respectively. For both loss function, the population minimizer is

$$f^*(x) = \frac{1}{2} \log\left(\frac{p(x)}{1-p(x)}\right), \quad p(x) = \mathbb{P}[Y = 1|X = x].$$

For cases where $Y \in \{0, 1, 2, \dots\}$, the Poisson log-likelihood is often appropriate:

$$\rho(y, f) = -yf + \exp(f), \quad f = \log(\lambda).$$

In survival analysis, we can derive the loss function from the partial likelihood in the Cox model (Cox 1975).

5.1 General Twin Boosting with general weak learners

If the loss function $\rho(\cdot, \cdot)$ is differentiable (almost everywhere) with respect to the second argument, the generic boosting algorithm from Sect. 2.1 can be used.

General Twin Boosting is defined exactly as the general Twin L_2 Boosting algorithm from Sect. 4, except that in step 3, instead of using residuals U_i , we will use

$$U_i = -\frac{\partial}{\partial f} \rho(Y, f)|_{f=\hat{f}_{m-1}(X_i)}, \quad i = 1, \dots, n,$$

as in the generic boosting algorithm from Sect. 2.1. In the special case of Twin Boosting with the componentwise linear least squares learner, we would modify the residual vector U in step 2 of the algorithm in Sect. 3.2. From an implementation point of view, (general) Twin Boosting with differentiable loss functions is as simple as (general) Twin L_2 Boosting.

6 Empirical results

We report here some results on Twin Boosting for regression and classification and we compare them with boosting. We will demonstrate that Twin Boosting has a clear advantage over boosting if the truth has many ineffective predictor variables. Given the success of boosting algorithms in many application areas, Twin Boosting exhibits a substantial potential for further improvements over boosting. All of our results are displayed in Figures, giving a better summary how the methods behave as a function of boosting iterations.

6.1 Regression

The response variables Y_i are real-valued and the goal is estimation of the function $\mathbb{E}[Y|X = x]$ or prediction of new observations Y .

6.1.1 Simulated data

Consider the linear models:

model (2) with $p = 500, \beta_1 = 5,$
 $\beta_j = 0 \quad (j = 2, \dots, p),$
 $X_i \sim \mathcal{N}_p(0, I) \quad \text{and} \quad \varepsilon_i \sim \mathcal{N}(0, 1);$ (13)

model (2) with $p = 500, \beta_1 = \dots = \beta_5 = 1.175,$
 $\beta_j = 0 \quad (j = 6, \dots, p),$
 $X_i \sim \mathcal{N}_p(0, \Sigma), \quad \Sigma_{ij} = 0.8^{|i-j|}, \quad \text{and}$
 $\varepsilon_i \sim \mathcal{N}(0, 1).$ (14)

Both models (13) and (14) have the same signal to noise ratio $\mathbb{E}[|f(X)|^2]/\mathbb{E}[|\varepsilon|^2]$, where $f(x) = \sum_{j=1}^p \beta_j x^{(j)}$. Sample size is chosen as $n = 50$ and the number of independent simulation runs is 100.

We first use L_2 Boosting and Twin L_2 Boosting for linear models, using the componentwise linear least squares weak learner. The step-length factor is chosen as $\nu = 0.1$ and the number of boosting iterations in the first round of Twin Boosting is chosen as $m_1 = 50$ which is a reasonable value according to the performance of L_2 Boosting. Figure 1 displays the mean squared error (MSE) $\mathbb{E}[(\hat{\beta} - \beta)^T X_{new}]^2 = \mathbb{E}[(\hat{\beta}_j - \beta_j)^T \Sigma (\hat{\beta}_j - \beta_j)]$, with $\Sigma = \text{Cov}(X)$, (i.e. generalization error) and the number of selected and incorrectly




Fig. 1 L_2 Boosting (solid line) and Twin L_2 Boosting (dashed line) with componentwise linear least squares for model (13). Mean squared error (MSE) (left), average number of selected predictor variables (middle) and number of incorrectly selected predictor variables (right)

as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively). Simulation accuracy is indicated by dotted lines as 95% confidence intervals

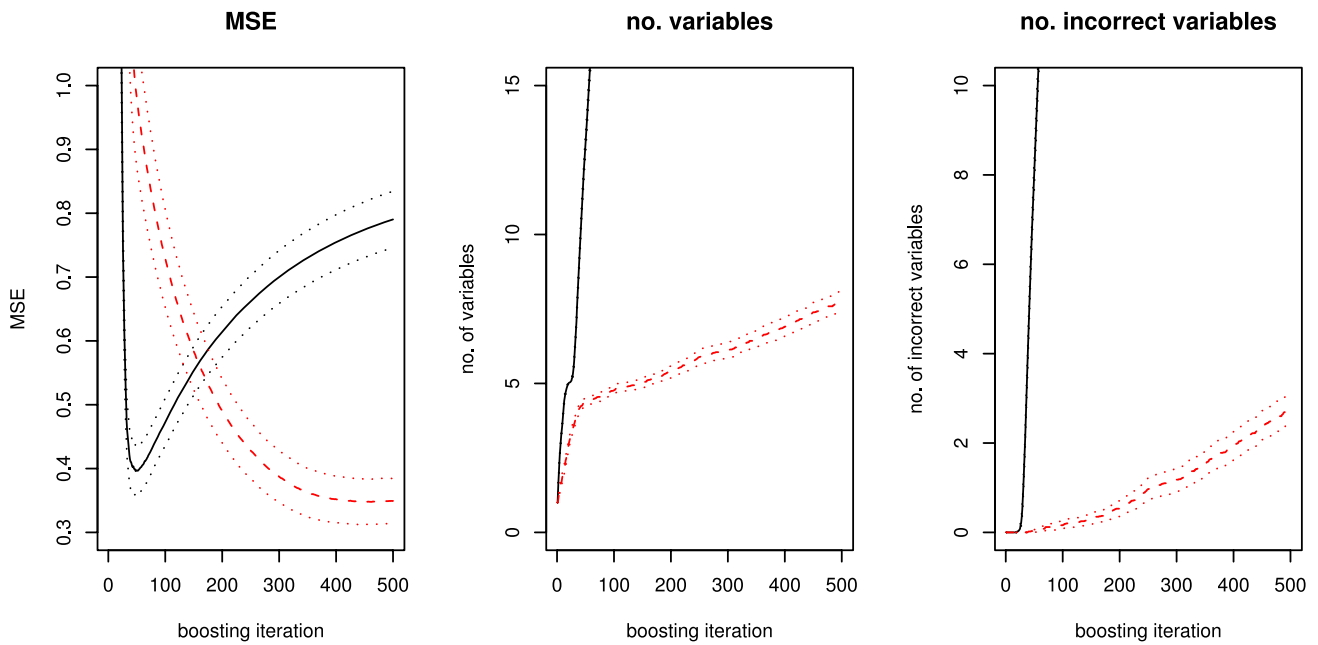


Fig. 2 L_2 Boosting (solid line) and Twin L_2 Boosting (dashed line) with componentwise linear least squares for model (14). Other specifications as in Fig. 1

Table 1 Performances of boosting with componentwise linear least squares (comp. LS) or stumps as weak learners, for models (13) and (14), at stopping iteration which minimizes mean squared error. Mean

squared error (MSE), number of selected variables/features and number of incorrectly selected variables/features (i.e. false positives). Standard errors are given into parentheses

Model, method	MSE	No. variables	No. incorrect variables
(13), L_2 Boost comp. LS	0.22 (0.015)	5.97 (0.124)	4.97 (0.124)
(13), Twin L_2 Boost comp. LS	0.05 (0.005)	1.01 (0.010)	0.01 (0.010)
(14), L_2 Boost comp. LS	0.40 (0.020)	12.41 (0.165)	7.41 (0.165)
(14), Twin L_2 Boost comp. LS	0.35 (0.018)	7.39 (0.167)	2.40 (0.166)
(13), L_2 Boost stumps	2.91 (0.087)	10.67 (0.233)	9.67 (0.237)
(13), Twin L_2 Boost stumps	2.25 (0.070)	4.52 (0.173)	3.52 (0.173)
(14), L_2 Boost stumps	3.45 (0.080)	31.98 (0.352)	26.98 (0.352)
(14), Twin L_2 Boost stumps	2.84 (0.069)	11.12 (0.231)	6.13 (0.213)

selected predictor variables (false positives), as a function of boosting iterations. Figures 1 and 2 illustrate very clearly that Twin Boosting is substantially better than boosting in terms of variable selection. For the very sparse case in (13), Fig. 1 also indicates relevant improvements in terms of prediction. In Table 1, we report some exact numbers.

Next, we consider L_2 Boosting and Twin L_2 Boosting with stumps. The results are displayed in Figs. 3 and 4; for Twin Boosting, we used $m_1 = 50$ iterations in the first round.

Although boosting (and Twin Boosting) with stumps yields an additive model fit, the weak learner is “mis-specified” (as is often the case in practice). This explains why the mean squared error is much larger than with componentwise linear least squares. Also here, L_2 Boosting se-

lects way too many predictor variables while Twin Boosting is very effective and substantially reduces the number of selected variables. Table 1 reports some numerical values. The results shown here are quite representative for many other simulation settings.

We will show in Fig. 5 the sensitivity of Twin L_2 Boosting to the choice of the tuning parameter m_1 . The final performance of Twin L_2 Boosting is not heavily depending on the choice of m_1 , as long as m_1 is within a reasonable range, and hence, some rough estimate of m_1 will be sufficient for good performance of Twin L_2 Boosting (when considering it as a function of the second tuning parameter m_2). In particular, the found insensitivity yields some evidence that the sequen-

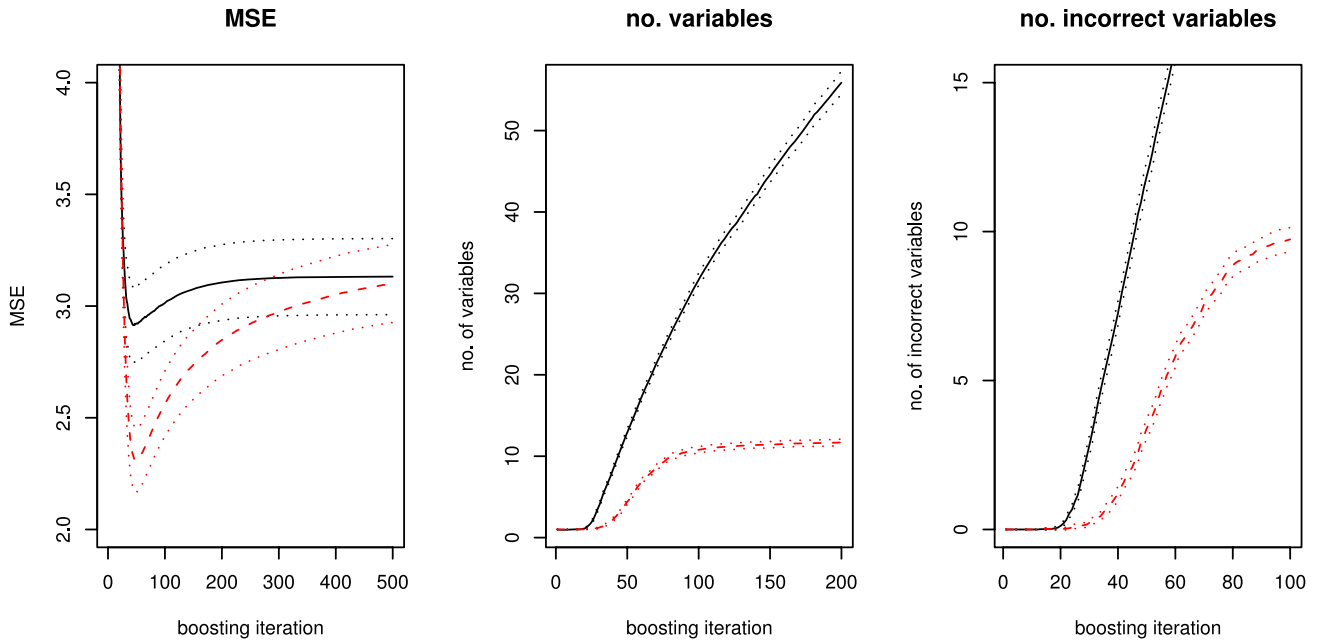


Fig. 3 L_2 Boosting (solid line) and Twin L_2 Boosting (dashed line) with stumps for model (13). Other specifications as in Fig. 1




Fig. 4 L_2 Boosting (solid line) and Twin L_2 Boosting (dashed line) with stumps for model (14). Other specifications as in Fig. 1

tial cross-validation scheme described in Sect. 3.2 should work reasonably well.

6.1.2 Real data

We consider two real data sets: *Ozone* concentration and *Motif* regression from molecular biology. The first data set is

available from <ftp://ftp.stat.math.ethz.ch/Research-Reports/Other-Manuscripts/buhlmann/ozone.dat>

The *Ozone* data is about daily ozone concentration in the Los Angeles basin as a function of $p = 8$ meteorological predictor variables. Sample size is $n = 330$. From a prediction point of view, the componentwise linear least squares weak learner is inferior than stumps. Thus, Fig. 6 reports

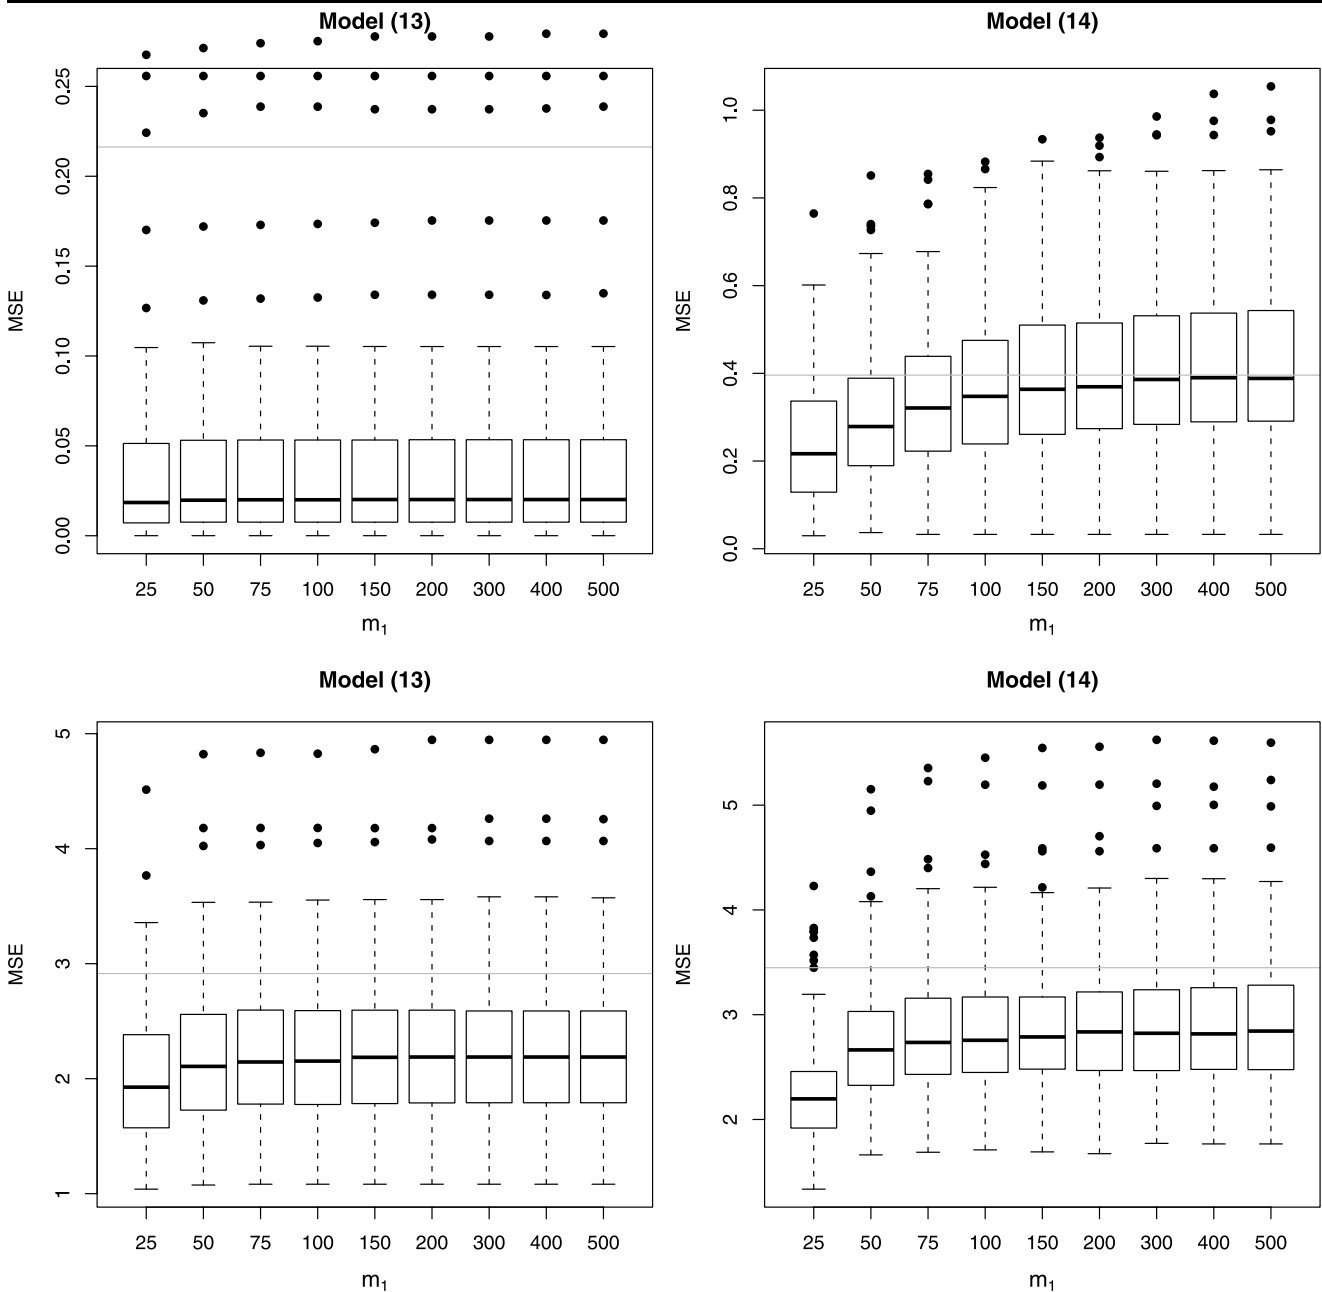


Fig. 5 Squared prediction error of Twin L_2 Boosting (boxplots), as a function of m_1 , and mean squared error of L_2 Boosting (grey horizontal line). Left: model (13); Right: model (14). Base procedure: componentwise linear least squares (top) or stumps (bottom)

only for boosting and Twin Boosting (with $m_1 = 100$ iterations in the first round) with stumps.

In addition, we look at a synthetically enlarged problem. We add 500 additional, ineffective noise predictor variables $X_{add} \sim \mathcal{N}_{500}(0, I)$. The problem has then dimension $p = 508$ with at most 8 effective predictors. This will enable us to see whether and how many from the obviously ineffective variables will be selected; we do not know whether all of the 8 original predictor variables are effective or not. We refer to an obviously incorrectly selected predictor vari-

able if it is one of the 500 synthetically added predictors. Figure 7 reports the results (with $m_1 = 100$ iterations in the first round of Twin Boosting). Twin L_2 Boosting has slightly better prediction performance than L_2 Boosting and is much better with respect to obviously incorrectly selected variables.

The *Motif* regression data models gene expression as a function of MDSCAN motif scores (Conlon et al. 2003, p. 3343: Spellman et al. data, 15th time point). This data is representative for many gene expression—motif scores data-

sets, all of them being very noisy. Our data has $p = 4312$ motif scores (predictor variables) and sample size (number of genes) is $n = 4443$. Figure 8 displays the results for the componentwise linear least squares weak learner.

Although L_2 Boosting is performing as well as Twin L_2 Boosting from a prediction point of view, the sparsity of Twin L_2 Boosting in terms of selected variables, and hence with a lower number of false positives, is crucial in this application. When using about 600–800 boosting iterations, L_2 Boosting selects 144–178 predictors while Twin

Boosting uses 41–53 variables only. Biological validation of about 50 potential motifs (cis-regulatory elements) is much more realistic than for three times as many candidates. L_2 Boosting and Twin L_2 Boosting with trees did not improve prediction performance while it selected more predictor variables than what is reported above for componentwise linear least squares.

6.1.3 Twin L_2 Boosting and general Twin L_2 Boosting for linear models

We briefly compare here the Twin L_2 Boosting algorithm for linear models, as described in Sect. 3.2, with general Twin L_2 Boosting for linear models as described in Sect. 4.1 and assertion (ii) of Proposition 2.

The simulation model is from Zou and Hastie (2005, model (d), Sect. 5) with strongly correlated covariates: a linear model as in (2) with $n = 50$, $p = 40$,

$$\beta_j = 3 \quad (j = 1, \dots, 15), \quad \beta_j = 0 \quad (j = 16, \dots, 40),$$

$$\text{Var}(\varepsilon_i) = 15^2,$$

and the covariates are generated (i.i.d. with respect to all i 's) as

$$X_i^{(j)} = Z_i^{(1)} + \eta_i^{(j)}, \quad j = 1, \dots, 5,$$

$$X_i^{(j)} = Z_i^{(2)} + \eta_i^{(j)}, \quad j = 6, \dots, 10,$$




Fig. 6 L_2 Boosting (solid line) and Twin Boosting (dashed line) with stumps for *Ozone* data. 10-fold cross-validation of: Squared error (left) and number of selected predictor variables (right), as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively)

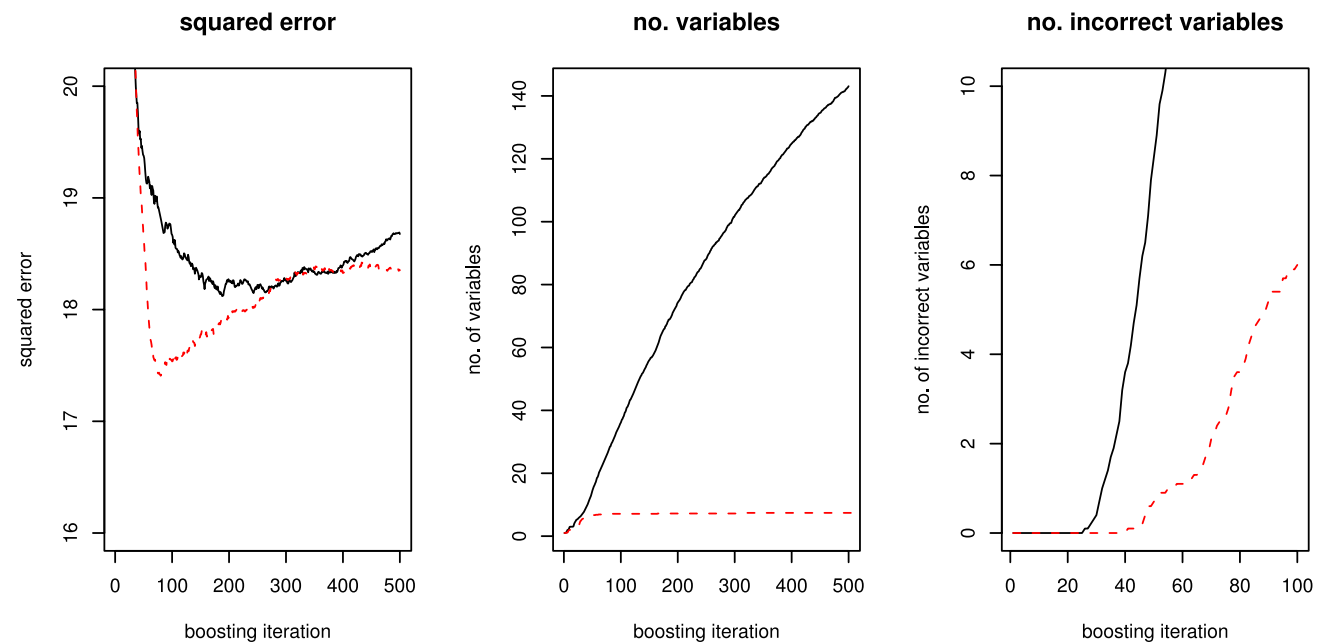


Fig. 7 L_2 Boosting (solid line) and Twin Boosting (dashed line) with stumps for *Ozone* data with synthetically enlarged predictor space ($p = 508$). 10-fold cross-validation of: Squared error (left), number of

selected predictor variables (middle) and number of obviously incorrectly selected variables (right), as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively)




Fig. 8 L_2 Boosting (solid line) and Twin L_2 Boosting (dashed line) with componentwise linear least squares for *Motif* regression data. Other specifications as in Fig. 6

$$X_i^{(j)} = Z_i^{(3)} + \eta_i^{(j)}, \quad j = 11, \dots, 15,$$

$$X_i^{(16)}, \dots, X_i^{(40)} \text{ i.i.d. } \sim \mathcal{N}(0, 1),$$

where $Z_i^{(1)}, Z_i^{(2)}, Z_i^{(3)}$ i.i.d. $\sim \mathcal{N}(0, 1)$ and $\eta_i^{(1)}, \dots, \eta_i^{(15)}$ i.i.d. $\sim \mathcal{N}(0, \sigma^2 = 0.01)$.

The results in terms of the mean squared error are reported in Fig. 9. We do not detect any substantial differences. This may be due to the simulation model, and we do not know whether there are examples where one version is clearly better than the other.

6.2 Classification

We consider some binary classification problems and use exclusively the logistic loss in (11) for boosting, i.e. Binomial or LogitBoosting. The classifier is given by $\text{sign}(\hat{f}(x))$ where $\hat{f}(\cdot)$ is the estimated function from boosting or Twin Boosting, respectively. This rule is equivalent to classify to the label with larger (conditional) class-probability.

6.2.1 Simulated data

We modify model (13) as follows:

$$p = 500, \quad \beta_1 = 2, \quad \beta_j = 0 \quad (j = 2, \dots, p),$$

$$X_i \sim \mathcal{N}_p(0, I), \quad \log(\pi_i / (1 - \pi_i)) = \sum_{j=1}^p \beta_j X_i^{(j)}, \quad (15)$$

$$Y_i \sim \text{Bernoulli}(\pi_i).$$

Sample size is again chosen as $n = 50$ and the number of independent simulation runs is 100. We reduced the size of the coefficient β_1 in comparison to model (13) to decrease the signal to noise ratio in the problem.

Figure 10 reports the results for Binomial/LogitBoosting and its Twin Boosting version (with $m_1 = 10$ iterations in

the first round of Twin Boosting) with componentwise linear least squares (which yields a logistic linear model). The results are qualitatively comparable to the case of regression in Fig. 1, demonstrating a clear advantage of Twin Boosting.

6.2.2 Real data

We consider the *Sonar* dataset ($n = 208, p = 60$) from the Statlog project, available from <ftp://ftp.stat.math.ethz.ch/Research-Reports/Other-Manuscripts/buhlmann/sonar.dat>, the *Ionosphere* ($n = 351, p = 34$) and the *monk* dataset (*Monk1*) ($n = 432, p = 6$) from the UCI machine learning repository (<http://www.ics.uci.edu/mllearn/MLSummary.html>), and the three datasets *Arcene*, *Madelon* and *Gisette* from the NIPS 2003 feature selection challenge (Guyon et al. 2006). We note that *Monk1* is a synthetic dataset: however, it is not generated by ourselves. For the *Sonar* and *Ionosphere* dataset, we also consider synthetically enlarged predictor spaces where we add 500 ineffective predictor variables; for the *Monk1* data, we exclusively consider the case with an enlarged feature space. We consider boosting and Twin Boosting with stumps for the *Sonar*, *Ionosphere*, *Arcene*, *Madelon* and *Gisette* data while for *Monk1*, we use larger trees as weak learners and the corresponding Twin Boosting with random feature subsets as described in Sect. 4. Logistic linear models, fitted by using the componentwise linear least squares weak learner, were not competitive for all six data-sets.

For the *Sonar* data, we use $m_1 = 100$ iterations in the first round of Twin Boosting. The results are displayed in Fig. 11. The classification accuracy is about the same for boosting and Twin Boosting while the latter selects about 30% fewer variables (when using reasonable stopping iterations which differ for the two methods).

We enlarge the number of features by adding 500 additional, ineffective noise predictor variables $X_{add} \sim \mathcal{N}_{500}(0, I)$. Then, the classification problem involves dimension $p = 560$ with at most 60 effective predictors. We refer to an obviously incorrectly selected predictor variable if it belongs to one of the 500 synthetically added features. Results are given in Fig. 12.

The interpretation is similar (even more in favor of Twin Boosting) as for the original *Sonar* data. In addition, Twin Boosting is much better in terms of obviously incorrectly selected variables.

For the *Ionosphere* data, $m_1 = 500$ iterations in the first round of Twin Boosting is a reasonable value. The results are displayed in Fig. 13.

Twin Boosting has marginally better prediction power while being more sparse in the selected variables. When adding 500 additional, ineffective noise predictor variables $X_{add} \sim \mathcal{N}_{500}(0, I)$ the problem has dimension $p = 534$ with at most 34 effective predictors. Results are displayed in




Fig. 9 Mean squared error for model from Sect. 6.1.3 (Zou and Hastie 2005). *Top:* Twin L_2 Boosting for linear models; *Bottom:* General Twin L_2 Boosting for linear models. *Left:* Contour plots as a function of the

boosting iterations m_1 and m_2 ; *Right:* L_2 Boosting (black line) and (general) Twin L_2 Boosting for various values of m_1 as a function of m_2 (grey lines)

Fig. 14, based on $m_1 = 200$ iterations in the first round of Twin Boosting; as above for the *Sonar* data, the obviously incorrectly selected variables can be determined.

Interestingly, the classification performance does not degrade for both boosting and Twin Boosting. The reason is probably due to the increased resistance of overfitting (e.g.




Fig. 10 LogitBoosting (solid line) and corresponding Twin Boosting (dashed line) with componentwise linear least squares in model (15). Misclassification error (left), average number of selected predictor variables (middle) and average number of incorrectly selected predic-

tor variables (right) as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively). Simulation accuracy is indicated by dotted lines as 95% confidence intervals




Fig. 11 LogitBoosting (solid line) and corresponding Twin Boosting (dashed line) with stumps for Sonar data. 10-fold cross-validation of: Misclassification error rate (left) and number of selected predictor variables (right, as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively))

when selecting wrong features) when using the misclassification error (Friedman et al. 2000, pp. 400–404) and having a situation with low noise (low misclassification error). Regarding the quality of feature selection, however, Twin Boosting is much better than boosting in terms of selecting obviously incorrect predictor variables.

In addition, LogitBoosting and TwinBoosting are benchmarked on three high-dimensional problems, the Arcene, Gisette, and Madelon datasets, see also <http://www.nipsfsc.ecs.soton.ac.uk/> for a detailed description. Here, our interest

is to compare both algorithms in real high-dimensional situations with respect to their feature selection properties for a varying number of initial boosting steps m_1 . The data comes with separate learning and validation samples, and we report performance measures (balanced misclassification error and negative binomial log-likelihood) for the validation samples.

Arcene offers a binary response and $p = 10'000$ features, the learning sample consists of $n_{\text{train}} = 100$ observations, additional $n_{\text{valid}} = 100$ observations are available for validation. The log-likelihood (Fig. 15, left panel) suggests to stop LogitBoosting after ca. 25 iterations to prevent overfitting. The balanced misclassification error attains its minimum after ca. 20–30 iterations as well. Four runs of TwinBoosting (with $m_1 \in \{25, 50, 75, 100\}$) have been performed as well. Based on the validation log-likelihood, between 15 and 20 iterations should be enough, the balanced misclassification error for all four values of m_1 is practically equivalent. Note that the number of selected features is smaller compared to LogitBoosting.

Roughly the same conclusions can be drawn for the Madelon problem ($n_{\text{train}} = 2000$, $n_{\text{valid}} = 600$, $p = 500$), see Fig. 16). The optimal number of boosting iterations is smaller for TwinBoosting whereas the prediction performance is slightly better for Boosting. However, TwinBoosting leads to a sparser model. It should be noted that the performance of TwinBoosting seems to be rather robust against different choices of m_1 . This is even more pronounced for the Gisette problem ($n_{\text{train}} = 6000$, $n_{\text{valid}} = 1000$, $p =$




Fig. 12 LogitBoosting (solid line) and corresponding Twin Boosting (dashed line) with stumps for *Sonar* data with synthetically enlarged predictor space. 10-fold cross-validation of: Misclassification error rate (left), number of selected predictor variables (middle) and number of

obviously incorrectly selected variables (right), as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively)




Fig. 13 LogitBoosting (solid line) and corresponding Twin Boosting (dashed line) with stumps for *Ionosphere* data. Other specifications as in Fig. 11

5000), see Fig. 17, where the four different models (based on $m_1 = (100, 200, 300, 400)$) are practically not distinguishable. Boosting (with 400 iterations) requires more than twice as many variables entering the model than TwinBoosting to achieve a similar performance.

6.2.3 Large trees and random feature subsets in Twin Boosting

For the *Monk1* data, boosting and Twin Boosting with stumps has a cross-validated misclassification error of about 0.25. Boosting with larger trees yields substantial improve-

ments with a misclassification error of 0.03. Thus, for this problem, it is essential to allow for interactions among the predictor variables (which is well known due to the construction of the *Monk1* data).

There are only little differences between boosting and Twin Boosting with larger trees. However, when adding 500 ineffective noise variables $X_{add} \in \{-1, +1\}^{500}$ with independent components and $\mathbb{P}[X_{add}^{(j)} = 1] = 0.5$ for all j (the original 6 predictors are categorical, often binary), the situation is very different. As weak learners, we use trees whose depths are at most 4 (which allows for interactions among four predictor variables, at least). Furthermore, for Twin Boosting, we sample $B = 500$ random feature subsets of size $|\mathcal{W}| = s_{\mathcal{W}} = 4$ (per boosting iteration), see Sect. 4. We use $m_1 = 10$ iterations in the first round of Twin Boosting. In addition, we add an additional iteration to Twin Boosting, termed Triple Boosting: i.e. the Twin Boosting fit (with $m_2 = 50$) is used as initialization and we then proceeded exactly as with Twin Boosting. Figure 18 illustrates the results: Twin Boosting improves upon boosting with respect to prediction and feature selection. Furthermore, Triple Boosting yields additional improvements over Twin Boosting.

This example with the *Monk1* data (with enlarged feature space) is demonstrating that Twin Boosting with random feature subsets and with larger trees leads to similar qualitative conclusions as for Twin Boosting with stumps. Moreover, we see that further (smaller) gains can be achieved by




Fig. 14 LogitBoosting (solid line) and corresponding Twin Boosting (dashed line) with stumps for *Ionosphere* data with synthetically enlarged predictor space. Other specifications as in Fig. 12




Fig. 15 LogitBoosting (solid line) and corresponding Twin Boosting, using various values of m_1 (various dashed lines), with stumps for *Arcene* data. Validation set error of log-likelihood (left) and misclassification rate (middle), and number of selected variables (right), as a function of boosting iterations (or iterations from the second round in Twin Boosting, respectively)

pursuing Triple Boosting invoking one stage more than Twin Boosting.

7 Conclusions

We proposed Twin Boosting which is as general and generic as (gradient-based) boosting. It can be used with general

weak learners, for example with trees enabling the applicability for mixed data types with continuous, ordinal and categorical features, and it is suitable in a wide variety of situations, including regression, classification, Poisson regression or survival analysis (using the loss function from the partial likelihood in the Cox model). It is easy to implement and computationally feasible for large problems with




Fig. 16 LogitBoosting (solid line) and corresponding Twin Boosting, using various values of m_1 (various dashed lines), with stumps for Madelon data. Other specifications as in Fig. 15




Fig. 17 LogitBoosting (solid line) and corresponding Twin Boosting, using various values of m_1 (various dashed lines), with stumps for Gisette data. Other specifications as in Fig. 15

potentially very many features (or predictors or covariates) and/or large sample size. Furthermore, it is useful for high-dimensional situations where the number of features is much larger than sample size.

We have empirically shown that Twin Boosting has much better feature or variable selection behavior than boosting. In particular, Twin Boosting leads to sparser solutions which implies a reduction in the number of false positives (fewer falsely selected features): a low number of false positives is sometimes highly desirable, e.g. in computational biology

where only a few features or variables (e.g. genes) will be biologically validated in follow-up experiments. For cases with a small number of important effective covariates and many noise features, Twin Boosting also improves the predictive accuracy of boosting; for other situations, we never found it substantially worse for prediction than boosting. For the special case of orthonormal linear models, we prove equivalence to the adaptive Lasso (Zou 2006) which provides some theoretical aspects for explaining our general empirical findings with Twin Boosting.




Fig. 18 LogitBoosting (solid line), corresponding Twin Boosting (dashed line) and Triple Boosting (dashed-dotted line) with larger trees for *Monk1* data with synthetically enlarged predictor space. Twin and Triple Boosting with random feature subsets. Other specifications as in Fig. 12

Appendix

Proof of Proposition 1 The proof of Theorem 2 in Bühlmann and Yu (2006) can be adapted. The main modification is needed for formula (22) and its previous five lines. We denote in short by $\beta_{init} = \hat{\beta}_{init}^{[m_1]}$ and $\hat{\beta}^{[m]} = \hat{\beta}_{TWB}^{[m]}$.

The residual sum of squares of Twin L_2 Boosting at iteration m , denoted by RSS_m , decreases monotonically in m . The difference in residual sum of squares is:

$$RSS_m - RSS_{m+1} = |\langle \mathbf{U}, \mathbf{X}^{(S_{m+1})} \rangle|^2,$$

where \mathbf{U} denotes the residual vector $\mathbf{Y} - \mathbf{X}\hat{\beta}^{[m]}$ and S_{m+1} the selected variable in iteration $m + 1$. In every step of Twin L_2 Boosting, a maximal reduction of the weighted difference in residual sum of squares is used:

$$G_{m+1} = (RSS_m - RSS_{m+1})|\beta_{init, S_{m+1}}|^2 = |\langle \mathbf{U}, \mathbf{X}^{(S_{m+1})} \rangle|^2 |\beta_{init, S_{m+1}}|^2,$$

and the sequence $G_{m+1}, m = 1, 2, \dots$ is monotonically decreasing (because of the definition of Twin L_2 Boosting, the independence of fitting the i th component of β from the j th component ($i \neq j$) and the form of the decay of differences of residual sum of squares). Therefore, every stopping iteration corresponds to a tolerance δ^2 as in formula (22) in Bühlmann and Yu (2006), using here G_{m+1} instead of $RSS_m - RSS_{m+1}$. The remaining part of the proof is exactly as in Bühlmann and Yu (2006): the additional factor $|\beta_{init, i}|^2$ leads to the assertion of Proposition 1. \square

Proof Proposition 2 For both assertions, it holds that $\|\mathbf{X}^{(j)}\| = 1$ for all $j = 1, \dots, p$. By formula (9), we only have to deal with the form of the coefficient $C_{\mathcal{W}} = C_j$

in (8). Denote by $\hat{\gamma}_j = \langle \mathbf{U}, \mathbf{X}^{(j)} \rangle$ the estimated regression coefficient of \mathbf{U} versus $\mathbf{X}^{(j)}$. Then,

$$C_j = \langle \hat{\mathbf{f}}_{init}^{[m_1]}, \hat{\gamma}_j \mathbf{X}^{(j)} \rangle / \|\hat{\gamma}_j \mathbf{X}^{(j)}\| \tag{16}$$

$$= \left\langle \sum_{k=1}^p \hat{\beta}_{init, k}^{[m_1]} \mathbf{X}^{(k)}, \hat{\gamma}_j \mathbf{X}^{(j)} \right\rangle / \|\hat{\gamma}_j \mathbf{X}^{(j)}\|. \tag{17}$$

For assertion (i), we re-write (16) using orthonormality

$$C_j = \hat{\beta}_{init, j}^{[m_1]} \hat{\gamma}_j / |\hat{\gamma}_j|$$

and hence $C_j^2 = (\hat{\beta}_{init, j}^{[m_1]})^2$ equaling the factor in formula (5).

For assertion (ii), using that $\|\mathbf{X}^{(j)}\| = 1$, (16) yields

$$C_j^2 = \left(\sum_{k=1}^p \hat{\beta}_{init, k}^{[m_1]} \langle \mathbf{X}^{(k)}, \hat{\gamma}_j \mathbf{X}^{(j)} \rangle \right)^2$$

which equals the multiplier in formula (5) from assertion (ii). \square

References

Breiman, L.: Arcing classifiers (with discussion). *Ann. Stat.* **26**, 801–849 (1998)

Breiman, L.: Prediction games & arcing algorithms. *Neural Comput.* **11**, 1493–1517 (1999)

Breiman, L.: Random forests. *Mach. Learn.* **45**, 5–32 (2001)

Bühlmann, P.: Boosting for high-dimensional linear models. *Ann. Stat.* **34**, 559–583 (2006)

Bühlmann, P., Hothorn, T.: Boosting algorithms: regularization, prediction and model fitting (with discussion). *Stat. Sci.* **22**, 477–505 (2007)

Bühlmann, P., Meier, L.: Discussion of “One-step sparse estimates in nonconcave penalized likelihood models” (H. Zou and R. Li, auths.). *Ann. Stat.* **36**, 1534–1541 (2008)

Bühlmann, P., Yu, B.: Boosting with the L_2 loss: regression and classification. *J. Am. Stat. Assoc.* **98**, 324–339 (2003)

Bühlmann, P., Yu, B.: Sparse boosting. *J. Mach. Learn. Res.* **7**, 1001–1024 (2006)

Conlon, E., Liu, X., Lieb, J., Liu, J.: Integrating regulatory motif discovery and genome-wide expression analysis. *Proc. Natl. Acad. Sci. USA* **100**, 3339–3344 (2003)

Cox, D.: Partial likelihood. *Biometrika* **62**, 269–276 (1975)

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression (with discussion). *Ann. Stat.* **32**, 407–451 (2004)

Freund, Y., Schapire, R.: Experiments with a new boosting algorithm. In: *Proceedings of the Thirteenth International Conference on Machine Learning*. Morgan Kaufmann, San Francisco (1996)

Freund, Y., Schapire, R.: A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.* **55**, 119–139 (1997)

Friedman, J.: Greedy function approximation: a gradient boosting machine. *Ann. Stat.* **29**, 1189–1232 (2001)

Friedman, J., Hastie, T., Tibshirani, R.: Additive logistic regression: a statistical view of boosting (with discussion). *Ann. Stat.* **28**, 337–407 (2000)

Guyon, I., Gunn, S., Nikravesh, M., Zadeh, L.: *Feature Extraction, Foundations and Applications, Studies in Fuzziness and Soft Computing*. Springer, Heidelberg (2006)

- Hand, D.: Classifier technology and the illusion of progress (with discussion). *Stat. Sci.* **21**, 1–34 (2006)
- Huang, J., Ma, S., Zhang, C.-H.: Adaptive Lasso for sparse high-dimensional regression models. *Stat. Sin.* **18**, 1603–1618 (2008)
- Jamain, A., Hand, D.: The naive Bayes mystery. *Pattern Recogn. Lett.* **26**, 1752–1760 (2005)
- Lutz, R.: Logitboost with trees applied to the WCCI 2006 performance prediction challenge datasets. In: *Proceedings of the IJCNN 2006*
- Lutz, R.W., Bühlmann, P.: Conjugate direction boosting. *J. Comput. Graph. Stat.* **15**, 287–311 (2006)
- Meinshausen, N.: Relaxed Lasso. *Comput. Stat. Data Anal.* **52**, 374–393 (2007)
- Meinshausen, N., Bühlmann, P.: High-dimensional graphs and variable selection with the Lasso. *Ann. Stat.* **34**, 1436–1462 (2006)
- Meir, R., Rätsch, G.: An introduction to boosting and leveraging. In: Mendelson, S., Smola, A. (eds.) *Advanced Lectures on Machine Learning*. Lecture Notes in Computer Science. Springer, Berlin (2003)
- Rätsch, G., Onoda, T., Müller, K.: Soft margins for AdaBoost. *Mach. Learn.* **42**, 287–320 (2001)
- Ridgeway, G.: The state of boosting. *Comput. Sci. Stat.* **31**, 172–181 (1999)
- Schapire, R.: The boosting approach to machine learning: an overview. In: Denison, D., Hansen, M., Holmes, C., Mallick, B., Yu, B. (eds.) *MSRI Workshop on Nonlinear Estimation and Classification*. Springer, Berlin (2002)
- Tibshirani, R.: Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc., Ser. B* **58**, 267–288 (1996)
- Tutz, G., Binder, H.: Generalized additive modeling with implicit variable selection by likelihood-based boosting. *Biometrics* **62**, 961–971 (2006)
- Tutz, G., Reithinger, F.: A boosting approach to flexible semiparametric mixed models. *Stat. Med.* **26**, 2872–2900 (2007)
- Zhao, P., Yu, B.: On model selection consistency of Lasso. *J. Mach. Learn. Res.* **7**, 2541–2563 (2006)
- Zou, H.: The adaptive Lasso and its oracle properties. *J. Am. Stat. Assoc.* **101**, 1418–1429 (2006)
- Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *J. R. Stat. Soc., Ser. B* **67**, 301–320 (2005)