

Luo, Ye; Spindler, Martin

Conference Paper

L2-Boosting for Economic Applications

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2017: Alternative Geld- und Finanzarchitekturen - Session: Treatment Effects, No. G03-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Luo, Ye; Spindler, Martin (2017) : L2-Boosting for Economic Applications, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2017: Alternative Geld- und Finanzarchitekturen - Session: Treatment Effects, No. G03-V1, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/168194>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

L_2 Boosting for Economic Applications

By YE LUO AND MARTIN SPINDLER*

In the recent years more and more high-dimensional data sets, where the number of parameters p is high compared to the number of observations n or even larger, are available for applied researchers. Boosting algorithms represent one of the major advances in machine learning and statistics in recent years and are suitable for the analysis of such data sets. While Lasso has been applied very successfully for high-dimensional data sets in Economics, boosting has been underutilized in this field, although it has been proven very powerful in fields like Biostatistics and Pattern Recognition. We attribute this to missing theoretical results for boosting. The goal of this paper is to fill this gap and show that boosting is a competitive method for inference of a treatment effect or instrumental variable (IV) estimation in a high-dimensional setting. First, we present the L_2 Boosting with componentwise least squares algorithm and variants which are tailored for regression problems which are the workhorse for most Econometric problems. Then we show how L_2 Boosting can be used for estimation of treatment effects and IV estimation. We highlight the methods and illustrate them with simulations and empirical examples. For further results and technical details we refer to (?) and (?) and to the online supplement of the paper.

I. L_2 Boosting

To define the boosting algorithm for linear models, we consider the following re-

gression setting:

$$(1) \quad y_i = x_i' \beta + \varepsilon_i, \quad i = 1, \dots, n,$$

with vector $x_i = (x_{i,1}, \dots, x_{i,p_n})$ consisting of p_n predictor variables, β a p_n -dimensional coefficient vector, and a random, mean-zero error term ε_i , $E[\varepsilon_i | x_i] = 0$.

We allow the dimension of the predictors p_n to grow with the sample size n , and even be larger than the sample size, i.e., $\dim(\beta) = p_n \gg n$. But we will impose a sparsity condition. This means that there is a large set of potential variables, but the number of variables which have non-zero coefficients, denoted by s , is small compared to the sample size, i.e. $s \ll n$.

X denotes the $n \times p$ design matrix where the single observations x_i form the rows. X_j denotes the j th column of the design matrix, and $x_{i,j}$ the j th component of the vector x_i . We consider a fixed design for the regressors. We assume that the regressors are standardized with mean zero and variance one, i.e., $E_n[x_{i,j}] = 0$ and $E_n[x_{i,j}^2] = 1$ for $j = 1, \dots, p$,

The basic principle of Boosting can be described as follows. We follow the interpretation of (?) and (?) of Boosting as a functional gradient descent optimization (minimization) method. The goal is to minimize a loss function, e.g., an L_2 -loss or the negative log-likelihood function of a model, by an iterative optimization scheme. In each step the (negative) gradient which is used to update the current solution is modeled and estimated by a parametric or non-parametric statistical model, the so-called base learner. The fitted gradient is used for updating the solution of the optimization problem. A strength of boosting, besides the fact that it can be used for different loss functions, is its flexibility with regard to the base learners. We then repeat this procedure until some stopping criterion is met. The act of stopping is crucial for boosting

* Luo: University of Florida, yeluo@ufl.edu. Spindler: University of Hamburg, Hamburg Business School, Hamburg Center for Health Economics (hche) and Max Planck Society, martin.spindler@uni-hamburg.de. We thank participants and the discussant Hai Wang at the session Machine Learning in Econometrics at the AEA annual meeting 2017 in Chicago for valuable comments.

algorithms, as stopping too late or never stopping leads to overfitting and therefore some kind of penalization is required. A suitable solution is to stop early, i.e., before overfitting takes place. “Early stopping” can be interpreted as a form of penalization. Similar to LASSO, early stopping might induce a bias through shrinkage.

The literature has developed many different forms of boosting algorithms. In this paper we consider L_2 Boosting with componentwise linear least squares, as well as two variants. All three are designed for regression analysis. “ L_2 ” refers to the loss function, which is the typical sum-of-squares of the residuals $Q_n(\beta) = \sum_{i=1}^n (y_i - x'_i\beta)^2$ typical in regression analysis. In this case, the gradient equals the residuals. “Componentwise linear least squares” refers to the base learners. We fit the gradient (i.e. residuals) against each regressor (p univariate regressions) and select the predictor/variable which correlates most highly with the gradient/residual, i.e., decreases the loss function most, and then update the estimator in this direction. We next update the residuals and repeat the procedure until some stopping criterion is met. We consider L_2 Boosting and two modifications: the “classical” one which was introduced in (?) and refined in (?) for regression analysis, an orthogonal variant and post- L_2 Boosting. In signal processing and approximation theory, the first two methods are known as the pure greedy algorithm (PGA) and the orthogonal greedy algorithm (OGA) in the deterministic setting, i.e. in a setting without stochastic error terms.

ALGORITHM 1 (L_2 Boosting): 1) *Start / Initialization: $\beta^0 = 0$ (p -dimensional vector), $f^0 = 0$, set maximum number of iterations m_{stop} and set iteration index m to 0.*

2) *At the $(m + 1)^{th}$ step, calculate the residuals $U_i^m = y_i - x'_i\beta^m$.*

3) *For each predictor variable $j = 1, \dots, p$ calculate the correlation with the resid-*

uals:

$$\gamma_j^m := \frac{\sum_{i=1}^n U_i^m x_{i,j}}{\sum_{i=1}^n x_{i,j}^2} = \frac{\langle U^m, x_j \rangle_n}{E_n[x_{i,j}^2]}.$$

Select the variable j^m that is the most correlated with the residuals, i.e., $\max_{1 \leq j \leq p} |\text{corr}(U^m, x_j)|^1$.

4) *Update the estimator: $\beta^{m+1} := \beta^m + \gamma_{j^m}^m e_{j^m}$ where e_{j^m} is the j^m th index vector and $f^{m+1} := f^m + \gamma_{j^m}^m x_{j^m}$*

5) *Increase m by one. If $m < m_{stop}$, continue with (2); otherwise stop.*

Moreover, we consider two variants, namely post- L_2 Boosting (post-BA) and orthogonal L_2 Boosting (oBA).

Post- L_2 Boosting is a post-model selection estimator that applies ordinary least squares (OLS) to the model selected by the first-step, namely L_2 Boosting. To define this estimator formally, we make the following definitions: $T := \text{supp}(\beta)$ and $\hat{T} := \text{supp}(\beta^{m^*})$, the support of the true model and the support of the model estimated by L_2 Boosting as described above with stopping at m^* . A superscript C denotes the complement of the set with regard to $\{1, \dots, p\}$. In the context of LASSO, OLS after model selection was analyzed in (?). Given the above definitions, the post-model selection estimator or OLS post- L_2 Boosting estimator will take the form

$$(2) \quad \tilde{\beta} = \underset{\beta \in \mathbb{R}^p}{\text{argmin}} Q_n(\beta) : \beta_j = 0 \forall j \in \hat{T}^C,$$

where $Q_n(\beta) = \sum_{i=1}^n (y_i - x'_i\beta)^2$. For oBA, only the updating step is changed: an orthogonal projection of the response variable is conducted on all the variables which have been selected up to this point. The advantage of this method is that any variable is selected at most once in this procedure, while in the previous version the same variable might be selected at different steps which makes the analysis far more complicated. More formally, the method can be described as follows by modifying Step (4):

¹Equivalently, which fits the gradient best in a L_2 -sense.

ALGORITHM 2 (Orthogonal L_2 Boosting):

$$(4') \quad \hat{y}^{m+1} \equiv f^{m+1} = P_m y$$

and

$$U_i^{m+1} = Y_i - \hat{Y}_i^{m+1},$$

where P_m denotes the projection of the variable y on the space spanned by the first m selected variables (the corresponding regression coefficient is denoted β_o^m .)

Define X_o^m as the matrix which consists only of the columns which correspond to the variables selected in the first m steps, i.e. all X_{jk} , $k = 0, 1, \dots, m$. Then we have:

$$(3) \quad \beta_o^m = (X_o^{m'} X_o^m)^{-1} X_o^{m'} y$$

$$(4) \quad \hat{y}^{m+1} = f_o^{m+1} = X_o^m \beta_o^m$$

II. Inference with L_2 Boosting

In many cases the researcher is interested in valid inference of a low-dimensional parameter α in the presence of a high-dimensional nuisance parameter η where modern methods of machine learning are used to estimate the nuisance parameter η . To achieve this, two conditions are sufficient. First, it is important that estimating equations used to draw inferences about α satisfy a key orthogonality or immunization condition. When estimation and inference for α are based on the system of equations $M(\alpha, \eta) = 0$, this condition is fulfilled if

$$\partial_\eta M(\alpha, \eta) = 0.$$

This is an important element in providing an inferential procedure for α that remains valid when η is estimated using regularized machine learning methods, like Lasso or boosting. This orthogonality condition can generally be established. Second, it is important to use high-quality, structured estimators of η . Additional structure is usually required for η and the imposed estimator shall replicate this structure. It has been shown that under a sparsity condition Lasso is such a high-quality estimator. For a detailed description of this approach to valid

post selection inference we refer to another article in this session, (?), and (?).

(?) show that both post and orthogonal L_2 Boosting fulfill the conditions for a high quality estimator and have the same rate of convergence as Lasso.² Hence, those two variants of boosting can be used in a high-dimensional setting to estimate the nuisance parameter η and finally provide valid inference for the target parameter of interest, α . Formal statements of those results are provided in (?). In the following sections we show how L_2 Boosting can be used in practical economic problems, namely IV estimation with many potential IVs, and inference on treatment effects after selection among high-dimensional controls with L_2 Boosting. (?) and (?) provide the underlying theory for high-dimension, including the immunized moment conditions.

III. Estimation of Treatment Effects

In this section we report briefly the results from the simulation studies and the applications. A detailed description can be found in the supplement to the paper.

A. IV estimation with many instruments

Here we use boosting for estimation of the first stage in an IV setting with very many potential instrumental variables. The simulations show that boosting performs well in common settings, giving a lower bias than post-Lasso and rejection rates close to the nominal 5% level. In the application the influence of property rights protection, measured by federal appellate court decisions regarding in eminent domain, on GDP is analyzed. The boosting estimates replicate the Lasso estimates but with smaller standard errors. The economic conclusions remain unchanged.

²They also derive an upper bound for the rate of convergence of L_2 Boosting which is slower than the Lasso rate. Under additional assumptions an improvement might be possible.

B. Inference on treatment effects after selection among high-dimensional controls

For estimation and inference of a treatment effect in a setting with very many control variables we apply the so-called double selection method. It consists of two separate selection steps to determine the final controls for the regression of the outcome variable on the treatment variable and the selected controls. In the simulations Lasso shows a lower bias, but the rejection rates are too small compared to the nominal 5% level. The boosting estimates are close to the nominal level. The application, analyzing the convergence hypothesis in Macroeconomics, replicates the Lasso estimates.

IV. Conclusion

In this paper we define and explain briefly L_2 Boosting and two variants, namely post- L_2 Boosting and orthogonal boosting. We show, how these methods can be used within the orthogonal moment conditions framework for valid post-selection inference on treatment effects, either in an IV estimation setting with many instruments or in a setting with very many controls. Although only post- and orthogonal boosting have been shown to have rate of convergence in prediction norm, allowing valid post-selection inference, all three versions show very good properties in the simulations. In the applications we present boosting replicates the results from Lasso estimates. In sum, boosting seems to be a useful tool for the (micro-)econometrician's toolbox. The strength of boosting is particular in more complex models and offers many interesting questions for future research.

Online Appendix for “ L_2 Boosting for Economic Applications”

L_2 Boosting for Economic Applications]

Abstract. In this supplement additional material, in particular simulation results and applications, for the paper “ L_2 Boosting for Economic Applications” is presented.

Key words: L_2 Boosting, High-dimensional, instrumental variables, treatment effects, post-selection inference.

In this supplement additional material for the article “ L_2 Boosting for Economic Applications” is presented. First, a brief literature review of using boosting in Economics and Finance is given. The main part shows – by simulations and applications – how boosting can be used for estimation of treatment effects in a setting with very many control variables and with very many potential instrumental variables.

V. A Brief Review of the Literature

In this section we give a very brief review of applications of boosting in Economics and Finance. As the strength of machine learning is in prediction and model selection, boosting has been mainly used in these domains. Although boosting has been shown to be a useful approach in many statistical applications, it has been more or less ignored in empirical economics and finance. Some of the few exceptions include the following applications.

Boosting has been used for modeling and predicting volatility, amongst others, by (?), (?) and (?). (?) model stock-index volatility in a GARCH framework and employ boosting for componentwise knot selection

* Luo: University of Florida, yeluo@ufl.edu. Spindler: University of Hamburg, Hamburg Business School, Hamburg Center for Health Economics (hche) and Max Planck Society, martin.spindler@uni-hamburg.de. We thank participants and the discussant Hai Wang at the session Machine Learning in Econometrics at the AEA annual meeting 2017 in Chicago for valuable comments.

† Luo: University of Florida. Spindler: University of Hamburg, Hamburg Business School, Moorweidenstr. 18, 20148 Hamburg, Germany, martin.spindler@uni-hamburg.de.

in bivariate-spline estimation. (?) also employ a GARCH framework for modeling volatility but allow for a large set of macroeconomic variables which drive volatility. Their data set consists of monthly data with 253 months in total and 40 macroeconomic variables leading to more than 80 predictors (allowing lags). They employ boosting for model estimation and variable selection.

(?) use boosting to select the predictors in factor-augmented autoregressions. (?) classifies and predicts recessions with boosting.

VI. IV estimation with many instruments

In this section we demonstrate how boosting can be used for IV estimation in a setting with very many instruments.

A. Simulation

The simulations are based on a simple instrumental variables model data-generating process (DGP):

$$(5) \quad y_i = \beta d_i + e_i,$$

$$(6) \quad d_i = z_i \Pi + v_i,$$

$$(7) \quad (e_i, v_i) \sim N \left(0, \begin{pmatrix} \sigma_e^2 & \sigma_{ev} \\ \sigma_{ev} & \sigma_v^2 \end{pmatrix} \right) i.i.d.,$$

where $\beta = 1$ is the parameter of interest. The regressors $Z_i = (z_{i1}, \dots, z_{i100})'$ are normally distributed $N(0, \Sigma_Z)$ with $E[z_{ih}^2] = \sigma_z^2$ and $Corr(z_{ih}, z_{ij}) = 0.5^{|j-h|}$. σ_z^2 and σ_e^2 are set to unit, $Corr(e, v) = 0.6$. $\sigma_v^2 = 1 - \Pi' \Sigma_Z \Pi$ so that the unconditional variance of the endogenous variable equals 1. The first stage coefficients are set according to $\Pi = C \tilde{\Pi}$. For $\tilde{\Pi}$ we use a sparse design, i.e., $\tilde{\Pi} = (1, \dots, 1, 0, \dots, 0)$ with s coordinates equal to one and all other $p - s$ equal to zero. C is set in such a way that

we generate target values for the concentration parameter $\mu = \frac{n\Pi'\Sigma_z\Pi}{\sigma_v^2}$ which determines the behavior of IV estimators. We set the sample size equal n to 100, $s = 5$, $p = 100$ and the concentration parameter equal to 180. We estimate the first stage and calculate the first stage predictions with L_2 Boosting and its variants. The simulation results in Table ?? reveal that boosting has a smaller bias than post-Lasso in this setting. While post-Lasso produces rejection rates below the nominal 5% level, boosting is slightly above.

B. Application

We consider IV estimation of the effects of federal appellate court decisions regards in eminent domain on macroeconomic outcomes, here in particular the log of the GDP.¹ The structural model is given by

$$(8) \quad y_{ct} = \alpha_c + \alpha_t + \gamma_{ct} + \beta TakingsLaw_{ct} + W'_{ct}\delta + \epsilon_{ct}$$

where y_{ct} is the economic outcome, here log of GDP, for circuit c at time t , $TakingsLaw_{ct}$ number of pro-plaintiff appellate takings decisions in circuit c and time t , W_{ct} judicial pool characteristics, a dummy for whether there were no cases in that circuit-year, and the number of takings appellate decisions; α_c , α_t and γ_{ct} denote circuit-specific, time-specific and circuit-specific time trends. The parameter of interest, β , represents the effect of an additional decision upholding individual property interpreted as more protective individual property rights. The sample size is 312. The analysis of the causal effect of takings law is complicated by potential endogeneity between taking law decisions and economic variables. We employ an instrumental variables strategy that relies on the random assignment of judges to federal appellate panels and uses characteristics of federal circuit court judges (e.g. gender, race, religion, political affiliation, etc.) as instruments. This gives 138 instruments. We estimate the effect β by doing the selection of IVs and

¹We refer to (?) for more information on this application.

estimation the first-stage predicted values $\hat{TakingsLaw}_{ct}$ by employing the boosting algorithms introduced before. The results are given in Table ???. The boosting estimates agree with the Lasso estimate but give smaller standard errors. The economic conclusions remain unchanged.

VII. Inference on treatment effects after selection among high-dimensional controls

A. Simulation

Here we consider the following data-generating process:

$$(9) \quad y_i = d_i\alpha_0 + x_i'\theta_g + \xi_i$$

$$(10) \quad d_i = x_i'\theta_m + \nu_i,$$

where $(\xi_i, \nu_i)' \sim N(0, I_2)$ with I_2 the 2×2 identity matrix, $p = 200$, $x_i \sim N(0, \Sigma)$ with $\Sigma_{kj} = .5^{|j-k|}$. The parameter of interest, α_0 , is set equal to 0.5 and the sample size is $n = 100$. We consider a design with a decaying sequence of θ_m and θ_g , namely $1/j^2$ for $j = 1, \dots, p$. The results in Table ??? show that post-Lasso has a smaller bias than the boosting algorithms, but too small rejection rates (RP) compared to the nominal 5% level. Boosting has rejection rates close to the nominal level.

B. Application

In Macroeconomics an important question is how the rates (Y) at which economies of different countries grow are related to the initial wealth levels in each country (D) controlling for country's institutional, educational, and other similar characteristics (W). The relationship is captured by β_1 , the "speed of convergence/divergence", it measures the speed at which poor countries catch up $\beta_1 < 0$ or fall behind $\beta_1 > 0$ rich countries, after controlling for W . Hence the model is given as

$$(11) \quad Y = \beta_1 D + \beta_2' W + \epsilon.$$

For the analysis we use the Barro-Lee data set with 90 countries (observations) and

TABLE 1—SIMULATION RESULTS.

	post-Lasso	BA	post-BA	oBA
bias	0.194	0.142	0.142	0.141
RP	0.032	0.060	0.064	0.056

TABLE 2—EFFECT OF FEDERAL APPELLATE TAKINGS LAW DECISIONS ON ECONOMIC OUTCOMES.

	post-Lasso	BA	post-BA	oBA
$\hat{\beta}$	0.005	0.005	0.004	0.008
se	0.012	0.007	0.006	0.006

about 60 controls. We estimate the parameter of interest by the double selection method employing both Lasso and L_2 Boosting for the two selection steps. The double selection method implicitly constructs an orthogonal moment condition. The results are given in Table ???. Here again, the boosting estimates agree with the Lasso estimate and confirm the convergence hypothesis.

TABLE 3—SIMULATION RESULTS.

	post-Lasso	BA	post-BA	oBA
bias	0.082	0.121	0.136	0.121
RP	0.002	0.042	0.054	0.042

TABLE 4—EFFECT OF INITIAL GDP LEVEL ON GROWTH.

	post-Lasso	BA	post-BA	oBA
$\hat{\beta}$	-0.040	-0.042	-0.042	-0.041
se	0.015	0.012	0.012	0.013