



Adaptive step-length selection in gradient boosting for Gaussian location and scale models

Boyao Zhang¹ · Tobias Hepp² · Sonja Greven³ · Elisabeth Bergherr¹

Received: 10 March 2021 / Accepted: 10 January 2022 / Published online: 30 January 2022
© The Author(s) 2022

Abstract

Tuning of model-based boosting algorithms relies mainly on the number of iterations, while the step-length is fixed at a predefined value. For complex models with several predictors such as Generalized additive models for location, scale and shape (GAMLSS), imbalanced updates of predictors, where some distribution parameters are updated more frequently than others, can be a problem that prevents some sub-models to be appropriately fitted within a limited number of boosting iterations. We propose an approach using adaptive step-length (ASL) determination within a non-cyclical boosting algorithm for Gaussian location and scale models, as an important special case of the wider class of GAMLSS, to prevent such imbalance. Moreover, we discuss properties of the ASL and derive a semi-analytical form of the ASL that avoids manual selection of the search interval and numerical optimization to find the optimal step-length, and consequently improves computational efficiency. We show competitive behavior of the proposed approaches compared to penalized maximum likelihood and boosting with a fixed step-length for Gaussian location and scale models in two simulations and two applications, in particular for cases of large variance and/or more variables than observations. In addition, the underlying concept of the ASL is also applicable to the whole GAMLSS framework and to other models with more than one predictor like zero-inflated count models, and brings up insights into the choice of the reasonable defaults for the step-length in the simpler special case of (Gaussian) additive models.

✉ Boyao Zhang
boyao.zhang@uni-goettingen.de

¹ Faculty of Business and Economics, Georg-August-Universität Göttingen, Platz der Göttinger Sieben 3, 37073 Göttingen, Germany

² Department of Medical Informatics, Biometry and Epidemiology, Friedrich-Alexander Universität Erlangen-Nürnberg, Waldstrasse 6, 91054 Erlangen, Germany

³ School of Business and Economics, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

Keywords Step-Length · Gradient boosting · GAMLSS · Variable selection · Shrinkage

1 Introduction

Generalized additive models for location, scale and shape (GAMLSS) (Rigby and Stasinopoulos 2005) are distribution-based approaches, where all parameters of the assumed distribution for the response can be modelled as additive functions of the explanatory variables (Ripley 2004; Stasinopoulos et al. 2017). Specifically, the GAMLSS framework allows the conditional distribution of the response variable to come from a wide variety of discrete, continuous and mixed discrete-continuous distributions, see Stasinopoulos and Rigby (2007). Unlike conventional generalized additive models (GAMs), GAMLSS not only model the location parameter, e.g. the mean for Gaussian distributions, but also further distribution parameters such as scale (variance) and shape (skewness and kurtosis) through the explanatory variables in linear, nonlinear or smooth functional form.

The coefficients of GAMLSS are usually estimated based on penalized maximum likelihood method (Rigby and Stasinopoulos 2005). However, this approach cannot deal with high dimensional data, or more precisely, the case of more variables than observations (Bühlmann 2006). As the selection of informative covariates is an important part of practical analysis, Mayr et al. (2012) combined the GAMLSS framework with componentwise gradient boosting (Bühlmann and Yu 2003; Hofner et al. 2014; Hothorn et al. 2018) such that variable selection and estimation can be performed simultaneously. The original method cyclically updates the distribution parameters, i.e. all predictors will be updated sequentially in each boosting iteration (Hofner et al. 2016). Because the levels of complexity vary across the prediction functions, separate stopping values are required for each distribution parameter. Consequently, these stopping values have to be optimized jointly as they are not independent of each other. The commonly applied joint optimization methods like grid search are, however, computationally very demanding. For this reason, Thomas et al. (2018) proposed an alternative non-cyclical algorithm that updates only one distribution parameter (yielding the strongest improvement) in each boosting iteration. This way, only one global stopping value is needed and the resulting one-dimensional optimization procedure vastly reduces computing complexity for the boosting algorithm compared to the previous multi-dimensional one. The non-cyclical algorithm can be combined with stability selection (Meinshausen and Bühlmann 2010; Hofner et al. 2015) to further reduce the selection of false positives (Hothorn et al. 2010).

In contrast to the cyclical approach, the non-cyclical algorithm avoids an equal number of updates for all distribution parameters as it is not useful to artificially enforce updates for parameters with a less complex structure than other parameters. However, it becomes even more important to fairly select the predictor to be updated in any given iteration. The current implementation of Thomas et al. (2018), however, uses fixed and equal step-lengths for all updates, regardless of the achieved loss reduction of different distribution parameters. In other words, different parameters affect the loss in different ways, and an update of the same size on all predictors hence results in

different improvement with respect to loss reduction. As a consequence, a more useful update of one parameter could be rejected in favor of the other one just because the relevance in the loss function varies. As we demonstrate later, this leads to imbalanced updates that affect the fair selection and predictors with large number of boosting iterations still tend to be underfitted. This seems inconsistent, since one expects the underfitted predictor to be updated with a small number of iterations. As we show later, a large σ in a Gaussian distribution leads to a small negative gradient of μ and consequently the improvement for μ with fixed small step-lengths in each boosting iteration will also be small. This results in the algorithm needing a lot of updates for μ until its empirical risk decreases to the level of σ . However, the algorithm may stop long before the corresponding coefficients are well estimated.

We address this problem by proposing a version of the non-cyclical boosting algorithm for GAMLSS, especially for Gaussian location and scale models, that adaptively and automatically optimizes the step-lengths for all predictors in each boosting iteration. This ensures no parameters favored over the other by finding the factor that results in the overall best model improvement for each update and then bases the decision on which parameter to update on this comparison. While the adaptive approach does not enforce equal numbers of updates for all distribution parameters, it yields a fair selection of predictors to update and a natural balance in the updates. For the very special Gaussian case, we also derive (semi-)analytical adaptive step-lengths that decrease the need for numerical optimization and discuss their properties. Our findings have implications beyond boosted Gaussian location and scale models for boosting other models with several predictors, e.g. the whole GAMLSS framework in general or for zero-inflated count models, and also give insights into the step-length choice for the simpler special case of (Gaussian) additive models.

The structure of this paper is organized as follows: Section 2 introduces the boosted GAMLSS models including the cyclical and non-cyclical algorithms. Section 3 discusses how to apply the adaptive step-length on the non-cyclical boosted GAMLSS algorithm, and introduces the semi-analytical solutions of the adaptive step-length for the Gaussian location and scale models and discuss their properties. Section 4 evaluates the performance of the adaptive algorithms and the problem of fixed step-length in two simulations. Section 5 presents the application of the adaptive algorithms for two datasets: the malnutrition data, where the outcome variance is very large, and the riboflavin data, which has more variables than observations. Section 6 concludes with a summary and discussion. Further relevant materials and results are included in the appendix.

2 Boosted GAMLSS

In this section, we briefly introduce the GAMLSS models and the two cyclical and noncyclical boosting methods for estimation.

2.1 GAMLSS and componentwise gradient boosting

Conventional generalized additive models (GAM) assume a dependence only of the conditional mean μ of the response on the covariates. GAMLSS, however, also model other distribution parameters such as the scale σ , skewness ν and/or kurtosis τ with a set of statistical models.

The K distribution parameters $\boldsymbol{\theta}^T = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_K)$ of a density function $f(\mathbf{y}|\boldsymbol{\theta})$ are modelled by a set of up to K additive models. The model class assumes that the observations y_i for $i \in \{1, \dots, n\}$ are conditionally independent given a set of explanatory variables. Let $\mathbf{y}^T = (y_1, y_2, \dots, y_n)$ be a vector of the response variable and \mathbf{X} be a $n \times J$ data matrix. In addition, we denote $\mathbf{X}_{i\cdot}$, $\mathbf{X}_{\cdot j}$ and X_{ij} as the i -th observation (vector of length J), j -variable (vector of length n) and the i -th observation of the j -th variable (a single value) respectively. Let $g_k(\cdot)$, $k = 1, \dots, K$ be known monotonic link functions that relate K distribution parameters to explanatory variables through additive models given by

$$g_k(\boldsymbol{\theta}_k) = \eta_{\boldsymbol{\theta}_k}(\mathbf{X}) = \beta_{0, \boldsymbol{\theta}_k} \mathbf{1}_n + \sum_{j=1}^J f_{j, \boldsymbol{\theta}_k}(\mathbf{X}_{\cdot j} | \beta_{j, \boldsymbol{\theta}_k}), \quad \text{for } k = 1, \dots, K, \quad (1)$$

where $\boldsymbol{\theta}_k = (\theta_{k,1}, \dots, \theta_{k,n})^T$ contains the n parameter values for the n observations and functions are applied elementwise if the argument is a vector, $\eta_{\boldsymbol{\theta}_k}$ is a vector of length n , $\mathbf{1}_n$ is a vector of ones and $\beta_{0, \boldsymbol{\theta}_k}$ is the model parameter specific intercept. Function $f_{j, \boldsymbol{\theta}_k}(\mathbf{X}_{\cdot j} | \beta_{j, \boldsymbol{\theta}_k})$ indicates the effects of the j -th explanatory variable $\mathbf{X}_{\cdot j}$ (vector of length n) for the model parameter $\boldsymbol{\theta}_k$, and $\beta_{j, \boldsymbol{\theta}_k}$ is the parameter of the additive predictor $f_{j, \boldsymbol{\theta}_k}(\cdot)$. Various types of effects (e.g., linear, smooth, random) for $f(\cdot)$ are allowed. If the location parameter ($\theta_1 = \mu$) is the only distribution parameter to be regressed ($K = 1$) and the response variable is from the exponential family, (1) reduces to the conventional GAM. In addition, f_j can depend on more than one variable (interaction), in which case $\mathbf{X}_{\cdot j}$ would be e.g. a $n \times 2$ matrix, but for simplicity we ignore this case in the notation.

A penalized likelihood approach can be used to estimate the unknown quantities; for more details, see Rigby and Stasinopoulos (2005). This approach does not allow parameter estimation in the case of more explanatory variables than observations, and variable selection for high-dimensional data is not possible, which, however, can be well solved by using boosting. The theoretical foundations regarding numerical convergence and consistency of boosting with general loss functions have been studied by Zhang and Yu (2005). The work of Bühlmann and Yu (2003) on L2 boosting with linear learners and Hastie et al. (2007) on the proof of the equivalence of the lasso and forward stagewise regression paved the way of componentwise gradient boosting (Hothorn et al. 2018), which emphasizes the importance of weak learners to reduce the tendency to overfit. To deal with the high-dimensional problems, Mayr et al. (2012) proposed a boosted GAMLSS algorithm, which estimates the predictors in GAMLSS with componentwise gradient boosting. As this method updates in general only one variable in each iteration, it can deal with data that has more variables than observations, and the important variables can be selected by controlling the stopping iterations.

To estimate the unknown predictor parameters $\beta_{j,\theta_k}, j \in \{1, \dots, J\}$ in equation (1), the componentwise gradient boosting algorithm minimizes the empirical risk R , which is also the loss ρ summed over all observations,

$$R = \sum_{i=1}^n \rho(y_i, \boldsymbol{\eta}(X_{i.})),$$

where the loss ρ measures the discrepancy between the response y_i and the predictor $\boldsymbol{\eta}(X_{i.})$. The predictor $\boldsymbol{\eta}(X_{i.}) = (\eta_{\theta_1}(X_{i.}), \dots, \eta_{\theta_K}(X_{i.}))$ is a vector of length K . For the i -th observation $X_{i.}$, each predictor $\eta_{\theta_k}(X_{i.})$ is a single value corresponding to the i -th entry in η_{θ_k} in equation (1). The loss function ρ usually used in GAMLSS is the negative log-likelihood of the assumed distribution of \mathbf{y} (Thomas et al. 2018; Friedman et al. 2000).

The main idea of gradient boosting is to fit simple regression base-learners $h_j(\cdot)$ to the pseudo-residuals vector $\mathbf{u}^T = (u_1, \dots, u_n)$, which is defined as the negative partial derivatives of loss ρ , i.e.,

$$\mathbf{u}_k^{[m]} = \left(-\frac{\partial}{\partial \eta_{\theta_k}} \rho(y, \boldsymbol{\eta}) \Big|_{\boldsymbol{\eta} = \hat{\boldsymbol{\eta}}^{[m-1]}(X_{i.}), y = y_i} \right)_{i=1, \dots, n},$$

where m denotes the current boosting iteration. In a componentwise gradient boosting iteration, each base-learner involves usually one explanatory variable (interactions are also allowed) and is fitted separately to $\mathbf{u}_k^{[m]}$,

$$\mathbf{u}_k^{[m]} \xrightarrow{\text{base-learner}} \hat{h}_{j,\theta_k}^{[m]}(X_{.j}) \quad \text{for } j = 1, \dots, J.$$

For linear base-learner, its correspondence to the model terms in (1) shall be

$$\hat{h}_{j,\theta_k}(X_{.j}) = X_{.j} \hat{\boldsymbol{\beta}}_j,$$

where the estimated coefficients can be obtained by using the maximum likelihood or least square method. The best-fitting base-learner is selected based on the residual sum of squares, i.e.,

$$j^* = \arg \min_{j \in \{1, \dots, J\}} \sum_{i=1}^n (u_{k,i} - \hat{h}_j(X_{ij}))^2,$$

thereby allowing for easy interpretability of the estimated model and also the use of hypothesis tests for single base-learners (Hepp et al. 2019). The additive predictor will be updated based on the best-fitting base-learner $\hat{h}_{j^*,\theta_{k^*}}(X_{.j^*})$ in terms of the best-performing sub-model $\eta_{\theta_{k^*}}$,

$$\hat{\boldsymbol{\eta}}_{\theta_{k^*}}^{[m]}(\mathbf{X}) = \hat{\boldsymbol{\eta}}_{\theta_{k^*}}^{[m-1]}(\mathbf{X}) + v \hat{h}_{j^*,\theta_{k^*}}(\mathbf{X}_{.j^*}), \tag{2}$$

where ν denotes the step-length. In order to prevent overfitting, the step-length is usually set to a small value, in most cases 0.1. Equation (2) updates only the best-performing predictor $\hat{\eta}_{\theta_{k^*}}^{[m]}$, all other predictors (i.e. for $k \neq k^*$) remain the same as in the previous boosting iteration. The best-performing sub-model θ_{k^*} can be selected by comparing the empirical risk, i.e. which model parameter achieves the largest model improvement.

The main tuning parameter in this procedure, as in other boosting algorithms, is how many iterations should be performed before it stops, which is denoted as $m_{\theta_{\text{stop}}}$. As too large or small $m_{\theta_{\text{stop}}}$ leads to over-/underfitting model, cross-validation (Kohavi 1995) is one of the most widely used methods to find the optimal $m_{\theta_{\text{stop}}}$.

2.2 Cyclical boosted GAMLSS

The boosted GAMLSS can deal with data that has more variables than observations, as the componentwise gradient boosting updates only one variable in each iteration. It leads to variable selection if some less important variables have never been selected as the best-performing variable and thus are not included in the final model for a given stopping iteration $m_{\theta_{\text{stop}}}$.

The original framework of boosted GAMLSS proposed by Mayr et al. (2012) is a cyclical approach, which means every predictor η_{θ_k} , $k \in \{1, \dots, K\}$ is updated in a cyclical manner inside each boosting iteration. The iteration starts by updating the predictor for the location parameter and uses the predictors from the previous iteration for all other parameters. Then, the updated location model will be used for updating the scale model and so on. A schematic overview of the updating process in iteration $m + 1$ for $K = 4$ is

$$\begin{aligned} (\hat{\mu}^{[m]}, \hat{\sigma}^{[m]}, \hat{\nu}^{[m]}, \hat{\tau}^{[m]}) &\xrightarrow{\text{update}} \hat{\eta}_{\mu}^{[m+1]} \rightarrow \hat{\mu}^{[m+1]} \\ (\hat{\mu}^{[m+1]}, \hat{\sigma}^{[m]}, \hat{\nu}^{[m]}, \hat{\tau}^{[m]}) &\xrightarrow{\text{update}} \hat{\eta}_{\sigma}^{[m+1]} \rightarrow \hat{\sigma}^{[m+1]} \\ (\hat{\mu}^{[m+1]}, \hat{\sigma}^{[m+1]}, \hat{\nu}^{[m]}, \hat{\tau}^{[m]}) &\xrightarrow{\text{update}} \hat{\eta}_{\nu}^{[m+1]} \rightarrow \hat{\nu}^{[m+1]} \\ (\hat{\mu}^{[m+1]}, \hat{\sigma}^{[m+1]}, \hat{\nu}^{[m+1]}, \hat{\tau}^{[m]}) &\xrightarrow{\text{update}} \hat{\eta}_{\tau}^{[m+1]} \rightarrow \hat{\tau}^{[m+1]}. \end{aligned}$$

However, not all of the distribution parameters have the same complexity, i.e., the stopping iterations $m_{\theta_{\text{stop}}}$ should be set separately for different parameters, or jointly optimized, for example by grid search. Since grid search scales exponentially with the number of distribution parameters, such optimization can be very slow.

2.3 Non-cyclical boosted GAMLSS

In order to deal with the issues of a cyclical approach, Thomas et al. (2018) proposed a *non-cyclical* version, that updates only one distribution parameter instead of successively updating all parameters in each boosting iteration by comparing the model improvement (negative log-likelihood) of each model parameter, see Algo-

rithm 1 (especially step 11). Consequently, instead of specifying separate stopping iterations $m_{\theta_{\text{stop}}}$ for different parameters and tuning them with the computationally demanding grid search, only one overall stopping iteration, denoted as m_{stop} , needs to be tuned with e.g. the line search (Friedman 2001; Brent 2013). The tuning problem thus reduces from a multi-dimensional to a one-dimensional problem, which vastly reduces the computing time.

Algorithm 1 has a nested structure, with the outer loop executing the boosting iterations and the inner loops addressing the different distribution parameters. The best-fitting base-learner and their contribution to the model improvement for every parameter is selected in the inner loop and compared in the outer loop (step 11). Therefore, only the best performing base-learner is updated in a single iteration by adding $v\hat{h}(X_{j^*})$ to the predictor of the corresponding parameter θ_{k^*} . Over the course of the iterations, the boosting algorithm steadily increases the model in small steps and the final estimates for the different base-learners are simply the sum of all their updates they may have received.

The cyclical approach led to an inherent but somewhat artificial balance between the distribution parameters, as the predictors for all distribution parameters are updated in each iteration. The different final stopping values $m_{\theta_{\text{stop}}}$ for the different distribution parameters - chosen by tuning methods such as cross-validation - allow to stop updates at different times for distribution parameters of different complexity to avoid overfitting. In the non-cyclical algorithm, especially when m_{stop} is not large enough, there is the danger of an imbalance between predictors. If the selection between predictors to update is not fair, this could lead to iterations primarily updating some of the predictors and underfitting others. We will provide a detailed example for the Gaussian distribution with large σ in Sect. 4.2.

A related challenge is to choose an appropriate step-length $v_{\theta_k}^{[m]}$ for both the cyclical and non-cyclical approaches. Tuning the parameters when boosting GAMLSS models relies mainly on the number of boosting iterations (m_{stop}), with the step-length v usually set to a small value such as 0.1. Bühlmann and Hothorn (2007) argued that using a small step-length like 0.1 (potentially resulting in a larger number of iterations m_{stop}) had a similar computing speed as using an adaptive step-length performed by doing a line search, but meant an easier tuning task for one parameter (m_{stop}) instead of two. However, this result referred to models with a single predictor. A fixed step-length can lead to an imbalance in the case of several predictors that may live on quite different scales. For example, 0.1 may be too small for μ but large for σ . We will discuss such cases analytically and with empirical evidence in the later sections. Moreover, varying the step-lengths for the different sub-models directly influences the choice of the best performing sub-model in the non-cyclical boosting algorithm, thus choosing a subjective step-length is not appropriate. In the following, we denote a fixed predefined step-length such as 0.1 as the *fixed step-length* (FSL) approach.

To overcome the problems stated above, we propose using an *adaptive step-length* (ASL) while boosting. In particular, we propose to optimize the step-length for each predictor in each iteration to obtain a fair comparison between the predictors. While the adaptive step-length has been used before, the proposal to use different ASLs for

Algorithm 1 Non-cyclical componentwise gradient boosting in multiple dimensions - Basic algorithm

- 1: Initialize the additive predictors $\hat{\eta}^{[0]} = (\hat{\eta}_{\theta_1}^{[0]}, \dots, \hat{\eta}_{\theta_K}^{[0]})$ with offset values.
- 2: For each distribution parameter $\theta_k, k = 1, \dots, K$, specify a set of base-learners, i.e., for parameter θ_k define $h_{1, \theta_k}(\cdot), \dots, h_{J_k, \theta_k}(\cdot)$ where J_k is the cardinality of the set of base-learners specified for θ_k .
- 3: **for** $m = 1$ to m_{stop} **do**
- 4: **for** $k = 1$ to K **do**
- 5: Compute negative partial derivatives $-\frac{\partial}{\partial \eta_{\theta_k}} \rho(y, \eta)$ and plug in the current estimates $\hat{\eta}^{[m-1]}(\cdot)$:

$$u_k^{[m]} = -\frac{\partial}{\partial \eta_{\theta_k}} \rho(y, \eta),$$

where $\eta = \hat{\eta}^{[m-1]}(X_{i \cdot})$ and $y = y_i$ for $i = 1, \dots, n$.

- 6: Fit (e.g. with the least square method) the negative gradient vector $u_k^{[m]}$ separately to every base-learner:

$$u_k^{[m]} \xrightarrow{\text{base-learner}} \hat{h}_{j, \theta_k}(X_{\cdot j}) \quad \text{for } j = 1, \dots, J_k.$$

- 7: Select the best-fitting base-learner $\hat{h}_{j^*, \theta_k}(X_{\cdot j^*})$ by inner loss, i.e., the residual sum of squares of the base-learner fit w.r.t. $u_k^{[m]} = (u_{k,1}^{[m]}, \dots, u_{k,n}^{[m]})^T$:

$$j^* = \arg \min_{j \in \{1, \dots, J_k\}} \sum_{i=1}^n (u_{k,i}^{[m]} - \hat{h}_{j, \theta_k}(X_{ij}))^2,$$

where we dropped the dependence of j^* on k in the notation for simplicity.

- 8: Set the step-length to a fixed value v_0 , usually $v_0 = 0.1$:

$$v_{\theta_k}^{[m]} = v_0$$

- 9: Compute the possible improvement of this update regarding the outer loss

$$\Delta \rho_k = \sum_{i=1}^n \rho \left(y_i, \hat{\eta}_{\theta_k}^{[m-1]}(X_{i \cdot}) + v_{\theta_k}^{[m]} \cdot \hat{h}_{j^*, \theta_k}(X_{ij^*}) \right).$$

- 10: **end for**
- 11: Update, depending on the value of the loss reduction, only the overall best-fitting base-learner $k^* = \arg \min_{k \in \{1, \dots, K\}} \Delta \rho_k$:

$$\hat{\eta}_{\theta_{k^*}}^{[m]}(X) = \hat{\eta}_{\theta_{k^*}}^{[m-1]}(X) + v_{\theta_{k^*}}^{[m]} \cdot \hat{h}_{j^*, \theta_{k^*}}(X_{\cdot j^*}).$$

- 12: Set $\hat{\eta}_{\theta_k}^{[m]} := \hat{\eta}_{\theta_k}^{[m-1]}$ for all $k \neq k^*$.
- 13: **end for**

different predictors is new and we will see that this leads to balanced updates of the different predictors.

3 Adaptive step-length

In this section, we first introduce the general idea of the implementation of adaptive step-lengths for different predictors to GAMLSS. For the important special case of a Gaussian location and scale models with two model parameters (μ and σ), we will derive and discuss their adaptive step-lengths and properties, which also serves as an important illustration of the relevant issues more generally.

3.1 Boosted GAMLSS with adaptive step-length

Unlike the step-length in Eq. (2) and Algorithm 1, step 11, the adaptive step-length may also vary in different boosting iterations according to the loss reduction.

The adaptive step-length can be derived by solving the optimization problem

$$v_{j^*, \theta_k}^{*[m]} = \arg \min_v \sum_{i=1}^n \rho \left(y_i, \hat{\eta}_{\theta_k}^{[m-1]}(X_{i.}) + v \cdot \hat{h}_{j^*, \theta_k}(X_{ij^*}) \right), \quad (3)$$

note that $v_{j^*, \theta_k}^{*[m]}$ is the *optimal step-length* of the model parameter θ_k dependent on j^* in iteration m . The optimal step-length is a value that leads to the largest decrease possible of the empirical risk and usually leads to overfitting of the corresponding variable if no shrinkage is used (Hepp et al. 2016). Therefore the actual adaptive step-length (ASL) we apply in the boosting algorithm is the product of two parts, the shrinkage parameter λ and the optimal step-length $v_{j^*, \theta_k}^{*[m]}$, i.e.,

$$v_{j^*, \theta_k}^{[m]} = \lambda \cdot v_{j^*, \theta_k}^{*[m]}.$$

In this article, we take $\lambda = 0.1$, thus 10% of the optimal step-length. By comparison, the fixed step-length $v = 0.1$ would correspond to a combination of a shrinkage parameter $\lambda = 0.1$ with the “optimal” step-length v^* set to one.

The non-cyclical algorithm with ASL can be improved by replacing the fixed step-length in step 8 of algorithm 1 with the adaptive one. We formulate this change in Algorithm 2.

As the parameters in GAMLSS may have quite different scales, updates with fixed step-length can lead to an imbalance between model parameters, especially when m_{stop} is not large enough. When using FSL, a single update for predictor η_{θ_1} may achieve the same amount of global loss reduction than several updates of another predictor η_{θ_2} even if the actually possible contribution of the competing base-learners is similar, because for different scales the loss reductions of η_{θ_2} in these iterations are always smaller than that of η_{θ_1} . However, such unfair selections can be avoided by using ASL, because the model improvement depends on the largest decrease possible of

Algorithm 2 Non-cyclical componentwise gradient boosting with adaptive step-length - Extension of basic algorithm 1

... Steps 1-7 equal to algorithm 1 ... , in addition, choose shrinkage parameter λ .
 8: Find the optimal step-length $v_{\theta_k}^{[m]}$ by optimizing the outer loss:

$$v_{j^*, \theta_k}^{[m]} = \arg \min_v \sum_{i=1}^n \rho \left(y_i, \hat{\eta}_{\theta_k}^{[m-1]}(X_{i.}) + v \cdot \hat{h}_{j^*, \theta_k}(X_{ij^*}) \right),$$

and set adaptive step-length $v_{j^*, \theta_k}^{[m]}$ as the optimal value with shrinkage λ :

$$v_{j^*, \theta_k}^{[m]} = \lambda \cdot v_{j^*, \theta_k}^{*[m]}.$$

... Steps 9-13 equal to those in algorithm 1 ...

each predictor, i.e., the potential reduction in the empirical risks of all predictors are on the same level and their comparison therefore is fair. Fair selection does not enforce an equal number of updates as in the cyclical approach. The ASL approach can lead to imbalanced updates of predictors, but such imbalance actually reveals the intrinsically different complexities of each sub-model.

The main contribution of this paper is the proposal to use ASLs for each predictor in GAMLSS. This idea can also be applied to other complex models (e.g. zero-inflated count models) with several predictors for the different parameters, because these models meet the same problem, i.e. the scale of these parameters might differ considerably. If a boosting algorithm is preferred for estimation of such a model, we provide a new solution to address these kinds of problems, i.e. separate adaptive step-lengths for each distribution parameter.

3.2 Gaussian location and scale models

In general, the adaptive step-length v can be found by optimizing procedures such as a line search. However, such methods do not help to reveal the properties of adaptive step-lengths and its relationship with model parameters. Moreover, a line search method searches for the optimal value from a predefined search interval, which can be difficult to find out since too narrow intervals might not include the optimal value and too large intervals increase the searching time. The direct computation from an analytical expression is faster than a search. By investigating the important special case of a Gaussian distribution with two parameters, we will learn a lot about the adaptive step-length for the general case. Nevertheless, we must underline that for many cases an explicit closed form for the adaptive step-length may not exist and line search still plays an irreplaceable role. We perform the following study of the analytical solutions for the Gaussian special case out of the wish of finding its inner relationship with the model parameters, in order to better understand the limitation of fixed step-length and how adaptive values improve the learning process.

Consider the data points $(y_i, \mathbf{x}_i), i \in \{1, \dots, n\}$, where \mathbf{x} is a $n \times J$ matrix. Assume the true data generating mechanism is the normal model

$$\begin{aligned} y_i &\sim N(\mu_i, \sigma_i) \\ \mu_i &= \eta_{\boldsymbol{\mu}}(\mathbf{x}_i) \\ \sigma_i &= \exp(\eta_{\boldsymbol{\sigma}}(\mathbf{x}_i)). \end{aligned}$$

As we talk about the observed data, we replace η_{θ_k} , where $k = 1, 2$ for Gaussian distribution, with $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, and replace \mathbf{X} with \mathbf{x} . The identity and exponential functions for $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are thus the corresponding inverse link. Taking the negative log-likelihood as the loss function, its negative partial derivatives $\mathbf{u}_{\boldsymbol{\mu}}$ and $\mathbf{u}_{\boldsymbol{\sigma}}$ in iteration m for both parameters can then be modelled with the base-learners $\hat{h}_{j,\boldsymbol{\mu}}^{[m]}$ and $\hat{h}_{j,\boldsymbol{\sigma}}^{[m]}$. The optimization process can then be divided into two parts: one is the ASL for the location parameter $\boldsymbol{\mu}$, and the other is for the scale parameter $\boldsymbol{\sigma}$. As the ASL shrinks the optimal value, we consider only the optimal step-lengths for both parameters.

3.2.1 Optimal step-length for $\boldsymbol{\mu}$

The analytical optimal step-length for $\boldsymbol{\mu}$ in iteration m is obtained through minimizing the empirical risk

$$\begin{aligned} v_{j^*,\boldsymbol{\mu}}^{*[m]} &= \arg \min_v \sum_{i=1}^n \rho\left(y_i, \{\hat{\eta}_{\boldsymbol{\mu}}^{[m]}(\mathbf{x}_i), \hat{\eta}_{\boldsymbol{\sigma}}^{[m-1]}(\mathbf{x}_i)\}\right) \\ &= \arg \min_v \sum_{i=1}^n \frac{\left(y_i - \hat{\eta}_{\boldsymbol{\mu}}^{[m-1]}(\mathbf{x}_i) - v \hat{h}_{j^*,\boldsymbol{\mu}}^{[m]}(x_{ij^*})\right)^2}{2\hat{\sigma}_i^{2[m-1]}}, \end{aligned} \tag{4}$$

where the expression $\hat{\sigma}_i^{2[m-1]}$ represents the square of the standard deviation in the previous iteration, i.e. $\hat{\sigma}_i^{2[m-1]} = (\hat{\sigma}_i^{[m-1]})^2$. The optimal value of $v_{j^*,\boldsymbol{\mu}}^{*[m]}$ is obtained by letting the derivative of the equation equal zero, so we get the analytical ASL for $\boldsymbol{\mu}$ (for more derivation details, see also appendix A.1):

$$v_{j^*,\boldsymbol{\mu}}^{*[m]} = \frac{\sum_{i=1}^n \left(\hat{h}_{j^*,\boldsymbol{\mu}}^{[m]}(x_{ij^*})\right)^2}{\sum_{i=1}^n \frac{\left(\hat{h}_{j^*,\boldsymbol{\mu}}^{[m]}(x_{ij^*})\right)^2}{\hat{\sigma}_i^{2[m-1]}}}. \tag{5}$$

It is obvious, that $v_{j^*,\boldsymbol{\mu}}^{*[m]}$ is not an independent parameter in GAMLSS but depends on the base-learner $\hat{h}_{\boldsymbol{\mu}}^{[m]}(x_{ij^*})$ with respect to the best performing variable $\mathbf{x}_{.j^*}$ and the estimated variance in the previous iteration $\hat{\sigma}_i^{2[m-1]}$.

In the special case of a Gaussian additive model, the scale parameter σ is assumed to be constant, i.e. $\hat{\sigma}_i^{[m-1]} = \hat{\sigma}^{[m-1]}$ for all $i \in \{1, \dots, n\}$. We then obtain

$$v_{j^*, \mu}^{*[m]} = \frac{\sum_{i=1}^n \left(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*})\right)^2}{\frac{1}{\hat{\sigma}^{2[m-1]}} \sum_{i=1}^n \left(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*})\right)^2} = \hat{\sigma}^{2[m-1]}. \tag{6}$$

This gives us an interesting property of the optimal step-length or ASL, i.e., the analytical ASL for μ in the Gaussian distribution is actually the variance (as computed in the previous boosting iteration). This property enables this paper to be not only applicable for the special GAMLSS case, but also for the boosting of additive models with normal responses. Therefore, in the case of Gaussian additive models, we can use $v_{j^*, \mu}^{[m]} = \lambda \hat{\sigma}^{2[m-1]}$ as the step-length, which has a stronger theoretical foundation, instead of the common choice 0.1.

Back to the general GAMLSS case, we can further investigate the behavior of the step-length by considering the limiting case of $m \rightarrow \infty$. For large m , all base-learner fits $\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*})$ converge to zero or are similarly small. If we consequently approximate all $\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*})$ by some small constant h , this gives an approximation of the analytical optimal step-length of

$$v_{j^*, \mu}^{*[m]} \approx \frac{\sum_{i=1}^n h^2}{\sum_{i=1}^n \frac{h^2}{\hat{\sigma}_i^{2[m-1]}}} = \frac{nh^2}{h^2 \sum_{i=1}^n \frac{1}{\hat{\sigma}_i^{2[m-1]}}} = \frac{n}{\sum_{i=1}^n \frac{1}{\hat{\sigma}_i^{2[m-1]}}}, \tag{7}$$

which is the harmonic mean of the estimated variances $\hat{\sigma}_i^{2[m-1]}$ in the previous iteration. While this expression requires m to be large, which may not be reached if m_{stop} is of moderate size to prevent overfitting, the expression still gives an indication of the strong dependence of the optimal step-length on the variances $\hat{\sigma}_i^{2[m-1]}$, which generalizes the optimal value of the additive model in (6).

3.2.2 Optimal step-length for σ

The optimal step-length for the scale parameter σ can be obtained analogously by minimizing the empirical risk, now with respect to $v_{j^*, \sigma}^{*[m]}$. We obtain

$$\begin{aligned} v_{j^*, \sigma}^{*[m]} &= \arg \min_v \sum_{i=1}^n \rho \left(y_i, \{ \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i), \hat{\eta}_\sigma^{[m]}(\mathbf{x}_i) \} \right) \\ &= \arg \min_v \sum_{i=1}^n \left(\hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i) + v \hat{h}_\sigma^{[m]}(x_{ij^*}) \right) \\ &\quad + \sum_{i=1}^n \frac{\left(y_i - \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) \right)^2}{2 \exp \left(2 \hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i) + 2v \hat{h}_\sigma^{[m]}(x_{ij^*}) \right)}. \end{aligned} \tag{8}$$

After checking the positivity of the second-order derivative of the expression in equation (8), the optimal value can be obtained by setting the first-order derivative equal to zero:

$$\sum_{i=1}^n \hat{h}_{\sigma}^{[m]}(x_{ij^*}) - \sum_{i=1}^n \frac{(\hat{h}_{\sigma}^{[m]}(x_{ij^*}) + \epsilon_{i,\sigma} + 1) \hat{h}_{\sigma}^{[m]}(x_{ij^*})}{\exp(2v_{j^*,\sigma}^{*[m]} \hat{h}_{\sigma}^{[m]}(x_{ij^*}))} \stackrel{!}{=} 0, \tag{9}$$

where $\epsilon_{i,\sigma}$ denotes the residuals when regressing the negative partial derivatives $u_{\sigma,i}^{[m]}$ on the base-learner $\hat{h}_{\sigma}^{[m]}(x_{ij^*})$, i.e., $u_{\sigma,i} = \hat{h}_{\sigma}^{[m]}(x_i) + \epsilon_{i,\sigma}$. Unfortunately, equation (9) cannot be further simplified, which means that there is no analytical ASL for the scale parameter σ in the Gaussian distribution. Hence, the optimal ASL must be found by performing a conventional line search. For more details, see also Appendix A.2.

Even without an analytical solution, we can still use (9) to further study the behavior of the ASL. Analogous to the derivation of (7), $\hat{h}_{\sigma}^{[m]}(x_{ij^*})$ converges to zero for $m \rightarrow \infty$. If we approximate with a (small) constant $\hat{h}_{\sigma}^{[m]}(x_{ij^*}) \approx h, \forall i \in \{1, \dots, n\}$. Then (9) simplifies to

$$\begin{aligned} \sum_{i=1}^n h - \sum_{i=1}^n \frac{(h + \epsilon_{i,\sigma} + 1)h}{\exp(2v_{j^*,\sigma}^{*[m]}h)} &= 0 \\ \Leftrightarrow v_{j^*,\sigma}^{*[m]} &= \frac{1}{2h} \log\left(h + 1 + \frac{1}{n} \sum_{i=1}^n \epsilon_{i,\sigma}\right) \\ \Leftrightarrow v_{j^*,\sigma}^{*[m]} &= \frac{1}{2h} \log(h + 1), \end{aligned} \tag{10}$$

where $\frac{1}{n} \sum_{i=1}^n \epsilon_{i,\sigma} = 0$ in the regression model. Equation (10) can be further simplified by approximating the logarithm function with a Taylor series at $h = 0$, thus

$$\begin{aligned} v_{j^*,\sigma}^{*[m]} &= \frac{1}{2h} \left(h - \frac{h^2}{2} + O(h^3) \right) \\ &= \frac{1}{2} - \frac{h}{4} + O(h^2). \end{aligned}$$

As $h \rightarrow 0$ for $m \rightarrow \infty$, the limit of this approximate optimal step-length for σ is

$$\lim_{m \rightarrow \infty} v_{j^*,\sigma}^{*[m]} = \lim_{h \rightarrow 0} \frac{1}{2} - \frac{h}{4} = \frac{1}{2}. \tag{11}$$

Thus, the ASL for σ approaches approximately 0.05 if we take the shrinkage parameter $\lambda = 0.1$ and iterations run for a longer time (and the boosting algorithm is not stopped too early to prevent overfitting for this trend to show).

3.3 (Semi-)Analytical adaptive step-length

Knowing the properties of the analytical ASL in boosting GAMLSS for the Gaussian distribution, we can replace the line search with the analytical solution for the location parameter μ . If we keep the line search for the scale parameter σ , we call this the *Semi-Analytical Adaptive Step-Length (SAASL)*. Moreover, we are interested in the performance of combining the analytical ASL for μ with the approximate value $0.05 = \lambda \cdot \frac{1}{2}$ (with $\lambda = 0.1$) for the ASL for σ , which is motivated by the limiting considerations discussed above and has a better theoretical foundation than selecting an arbitrary small value in the common FSL. We call this step-length setup *SAASL05*. In either of these cases, it is straightforward and computationally efficient to obtain the (approximate) optimal value(s) and both alternatives are faster than performing two line searches.

The semi-analytical solution avoids the need for selecting a search interval for the line search, at least for the ASL for μ in the case of SAASL and for both parameters for SAASL05. This is an advantage, since too large search intervals will cause additional computing time, but too small intervals may miss the optimal ASL value and again lead to an imbalance of updates between the parameters. Also note that the value 0.5 gives an indication for a reasonable range for the search interval for $v_{j^*,\sigma}^{*[m]}$ if a line search is conducted after all.

The boosting GAMLSS algorithm with ASL for the Gaussian distribution is shown in Algorithm 3.

For a chosen shrinkage parameter of $\lambda = 0.1$, the ν_σ in SAASL05 would be 0.05, which is a smaller or "less aggressive" value than 0.1 in FSL, leading to a somewhat larger number of boosting iterations but a smaller risk of overfitting, and to a better balance with the ASL for μ .

4 Simulation study

In the following, two simulations are shown to demonstrate the performance of the adaptive algorithms. The first one compares the estimation accuracy between the different non-cyclical boosted GAMLSS algorithms with FSL or ASL in a Gaussian regression model for location and scale. The second one underlines the problem of FSL and the performance of the adaptive approaches if the variance in this setting is large.

4.1 Gaussian location and scale model

The simulation study in Thomas et al. (2018) showed that their FSL non-cyclical approach outperforms the classical cyclical approach. We use the same setup to show that the ASL approach performs at least as good as the FSL non-cyclical approach (and hence also outperforms the classical cyclical approach). At the end of this subsection we will show that the reason for the good performance of FSL is due to the chosen simulated data structure. The setup is the following: the response y_i is drawn from $N(\mu_i, \sigma_i)$ for $n = 500$ observations, with 6 informative covariates x_{ij} , $j \in \{1, \dots, 6\}$

Algorithm 3 Non-cyclical componentwise gradient boosting for the Gaussian location and scale models with different step-lengths - Extension of basic algorithm 1

... Steps 1-7 equal to algorithm 1 ... , in addition, choose shrinkage parameter λ .

8: Set or find the step-length $v_{j^*, \theta_k}^{[m]}$ for $\theta_k \in \{\mu, \sigma\}$ by one of the followings:

- Adaptive step-length (ASL):

$$v_{j^*, \theta_k}^{*[m]} = \arg \min_v \sum_{i=1}^n \rho \left(y_i, \hat{\eta}_{\theta_k}^{[m-1]}(x_{i \cdot}) + v \cdot \hat{h}_{j^*, \theta_k}(x_{ij^*}) \right);$$

- Semi-analytical adaptive step-length (SAASL):
if $\theta_k = \mu$,

$$v_{j^*, \mu}^{*[m]} = \frac{\sum_{i=1}^n \left(\hat{h}_{j^*, \mu}(x_{ij^*}) \right)^2}{\sum_{i=1}^n \frac{\left(\hat{h}_{j^*, \mu}(x_{ij^*}) \right)^2}{\hat{\sigma}_i^{2[m-1]}}}$$

if $\theta_k = \sigma$, same as for ASL.

- Semi-analytical adaptive step-length (SAASL05):

if $\theta_k = \mu$, same as for SAASL,

if $\theta_k = \sigma$, $v_{j^*, \theta_k}^{*[m]} = 0.5$.

and set adaptive step-length $v_{j^*, \theta_k}^{[m]}$ as the optimal value with shrinkage λ :

$$v_{j^*, \theta_k}^{[m]} = \lambda \cdot v_{j^*, \theta_k}^{*[m]}.$$

... Steps 9-13 equal to those in algorithm 1 ...

drawn independently from $\text{Uni}(-1, 1)$. The predictors of both distribution parameters are:

$$\begin{aligned} \eta_{\mu}(x_{i \cdot}) &= \mu_i = x_{i1} + 2x_{i2} + 0.5x_{i3} - x_{i4} \\ \eta_{\sigma}(x_{i \cdot}) &= \log(\sigma_i) = 0.5x_{i3} + 0.25x_{i4} - 0.25x_{i5} - 0.5x_{i6}, \end{aligned}$$

where x_3 and x_4 are shared between both μ and σ . Moreover, $p_{n\text{-inf}} = 0, 50, 250$ or 500 non-informative variables sampled from $\text{Uni}(-1, 1)$ are also added to the model. We conduct $B = 100$ simulation runs.

The estimated coefficients of η_{μ} and η_{σ} , whose values are taken at stopping iterations tuned by 10-fold CV with the maximum number of boosting iterations set to 1000, are shown in Appendix Figs. 8 and 9.

Overall, the estimated coefficients are similar between all four methods, with the shrinkage bias of boosting only becoming apparent with an increasing number of noise variables.

Figure 1 shows the comparison of the mean squared error (MSE) among non-cyclical boosted algorithms for μ and σ , where the MSEs are defined on the predictor level as $\text{MSE}_{\mu} = \frac{1}{n} \sum_{i=1}^n (\mu_i - \eta_{\mu}(x_{i \cdot}))^2$ and $\text{MSE}_{\sigma} = \frac{1}{n} \sum_{i=1}^n (\log(\sigma_i) - \eta_{\sigma}(x_{i \cdot}))^2$, respectively. In general, all methods have a similar MSE, with the MSE of FSL

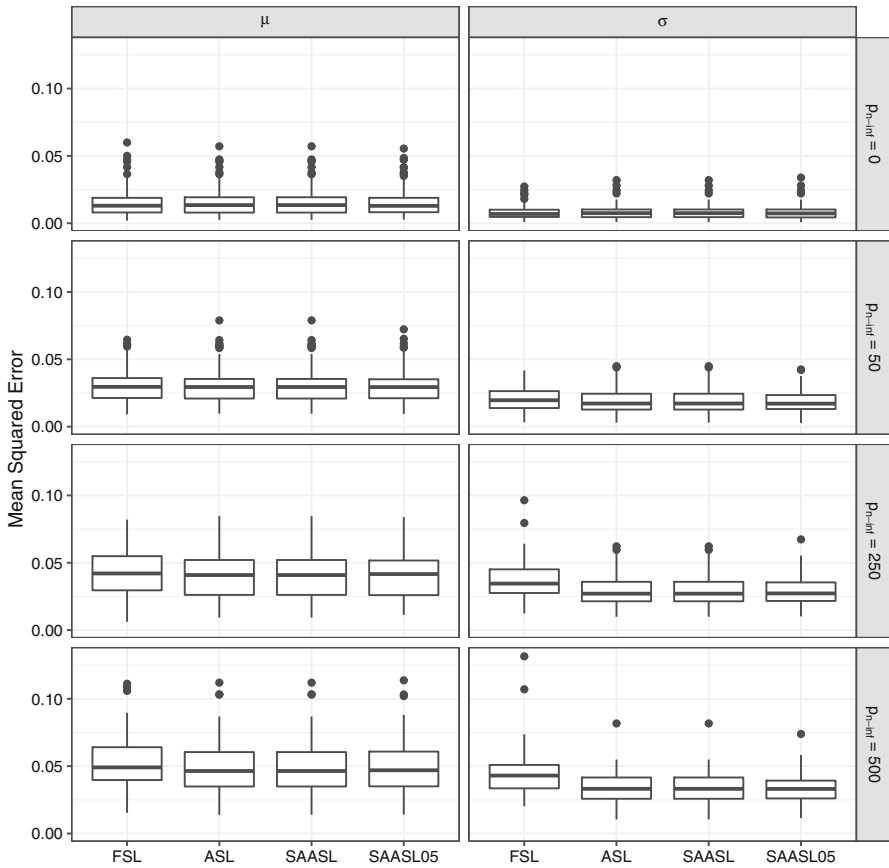


Fig. 1 Comparison between mean squared error for FSL and the three ASL methods. The left column comprises the MSE for η_μ , the right column for η_σ . The different numbers of non-informative variables are represented row-wise

increasing more strongly with the number of non-informative variables $p_{n\text{-inf}}$ and ASL methods hence slightly outperform FSL in the variance predictor for a high number of non-informative variables. ASL and SAASL show identical results, as they should if the line search is correctly conducted, with results returned by SAASL05 very similar.

Computing the negative log-likelihood in sample of the model fits reveals a slight advantage for FSL (see Appendix Fig. 10). However, this can be linked to the fact that FSL selects more false positive variables on average than the adaptive approaches and thus shows a relatively stronger tendency to overfit the training data (Fig. 2).

Figure 2 illustrates the false positives of each methods for each parameter. For σ , even if $p_{n\text{-info}}$ is small, the false positive rates of the adaptive approaches are notably smaller than those of FSL. As discussed above, $v_{j^*,\sigma}^{[m]} \approx 0.05$ for large m in the adaptive approach is smaller than $v_\sigma = 0.1$ for FSL. An update with a smaller, conservative step-length can apparently help to avoid overfitting and the adaptive step-length here

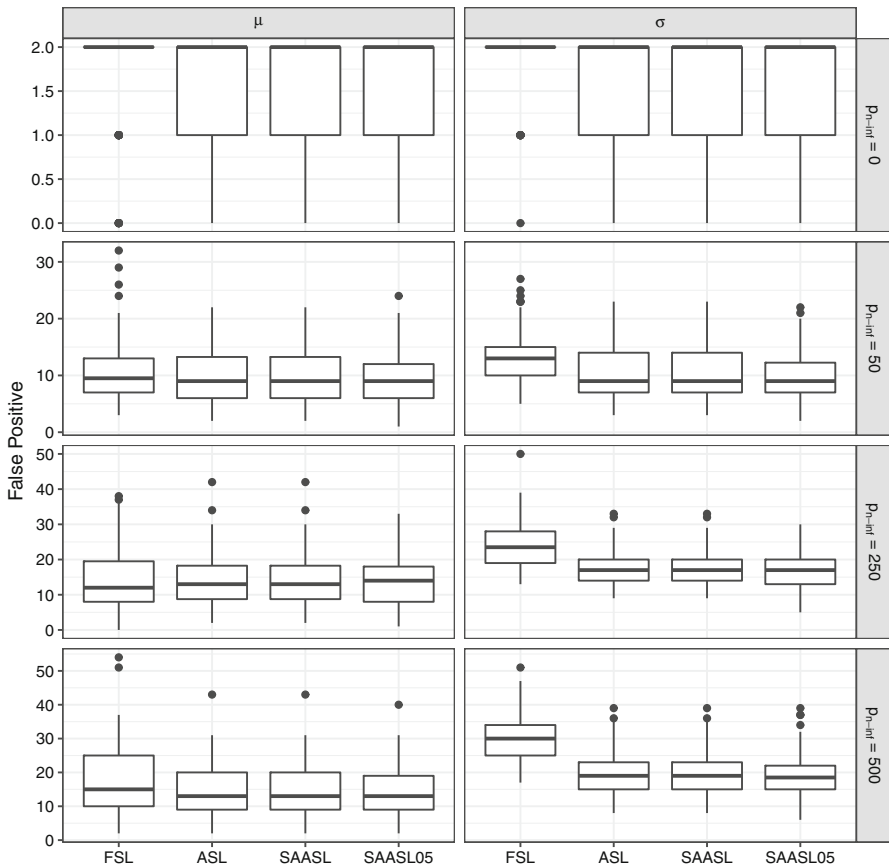


Fig. 2 Comparison between false positives for FSL and the three ASL methods. The left column comprises the false positives for μ , and the right column for σ . The different numbers of non-informative variables settings are represented row-wise

seems to strike the balance between learning speed and the number of false positives. While it would also be possible to lower the step-length for FSL to reduce the number of non-informative variables included in the final model, this would increase the number of boosting iterations and the computing time, and it would not address the imbalance between updates for μ and σ . The optimal choice of the step-length is also difficult without further tuning or an automatic selection as in ASL.

With respect to the neglecting of actually informative variables, i.e. false negatives, all four methods are able to find and select all variables for μ in all of the simulation runs. Regarding σ , the risk of false negatives slightly increases with the number of noise variables in the setting. However, even in the case of 500 noise-variables, only a single false negative is observed in between 3% and 6% of the runs, independently of the algorithm in question.

To some extent, the low false negative rate could be expected considering the somewhat greedy nature of boosting algorithms. For this reason, performance in terms




Fig. 3 Comparison of the optimal step-lengths $v_{j^*,\mu}^{*[m]}$ and $v_{j^*,\sigma}^{*[m]}$ in SAASL from one of the 100 simulation runs. The step-lengths for μ are in black dots, the step-lengths for σ in grey cross. Different horizontal layers of dots/crosses correspond to different covariates

of false-positive selections is arguably the more important aspect and speaks to the adaptive updates.

In Fig. 3 we show an example of the comparison between the optimal step-lengths in this case. As can be seen, the step-lengths for σ (depicted in grey) converge to 0.5 as shown in Sect. 3.2.2. The second fact that becomes obvious when looking at the figure is that the optimal step-lengths for both predictors do not differ a lot. Even though differences can be observed in early iterations in particular, the step-lengths still have the same order of magnitude. This is not only the case for this example but overall in this simulation setup. Having this in mind, the similar results for both approaches (FSL and ASL) are not very surprising anymore: there is hardly any difference in the approaches, since the updates do not need different step-lengths to be balanced. In the next subsection we will examine a case in which the data calls for different step-lengths, and see how both methods perform under those changed circumstances.

4.2 Large variance with resulting imbalance between location and scale

As discussed above, the Gaussian location and scale model in Sect. 4.1 did not lead to a large difference between FSL and ASL, as the optimal step-lengths for μ and σ were roughly similar and the imbalance between the updates for the two predictors in FSL was thus not large. In this section, we investigate a setting with a large variance, which leads to a stronger imbalance between the two parts of the model.

In the following, we use SAASL as a representative of the adaptive approaches in our presentation, as it yields identical results to ASL, but avoids the numerical search for the optimal v_μ by using the analytical result (5). Since estimated effects generally deviated more strongly from the theoretical values than before due to the large variance




Fig. 4 Distribution of coefficient estimates from $B = 100$ simulation runs. The true coefficients are marked by the dashed horizontal lines

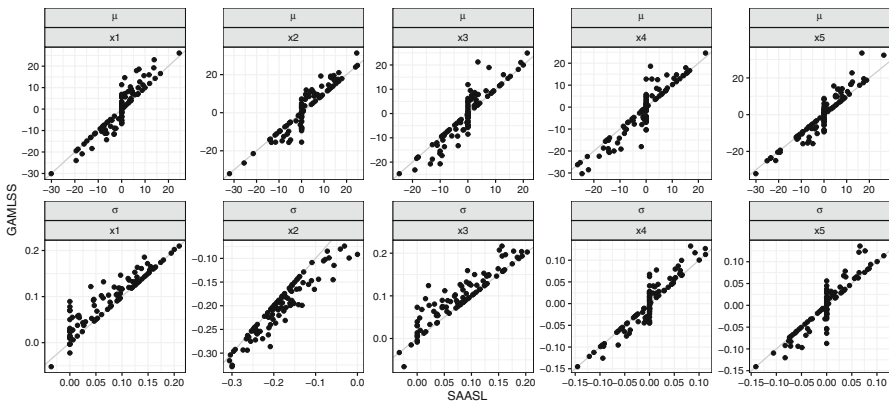


Fig. 5 Pairwise comparison of the estimated coefficients between GAMLSS and SAASL for both model parameter μ (top row) and σ (bottom row)

(details will be discussed later), we additionally compared the results to those obtained using GAMLSS with penalized maximum likelihood estimation as implemented in the R-package `gamlss` (Rigby and Stasinopoulos 2005).

Consider the data generating mechanism $y_i \sim N(\mu_i, \sigma_i), i \in \{1, \dots, 500\}$ with $B = 100$ simulation runs. The predictors are determined by

$$\begin{aligned} \eta_\mu(\mathbf{x}_{i.}) &= \mu_i = 1 + x_{i1} + 2x_{i2} - x_{i3} \\ \eta_\sigma(\mathbf{x}_{i.}) &= \log(\sigma_i) = 5 + 0.1x_{i1} - 0.2x_{i2} + 0.1x_{i3}, \end{aligned}$$

where $x_{.j} \sim \text{Uni}(-1, 1), j \in \{1, 2, 3, 4, 5\}, x_{.4}$ and $x_{.5}$ are noise variables. The choice of η_σ leads to an extremely large standard deviation in the order of 150 due to the large intercept 5. The stopping iteration is obtained by 10-fold CV, and the maximum number of iterations is 3000 and 2,000,000 for SAASL and FSL respectively. The

Table 1 Summary of the in-sample MSE for each estimation methods, i.e. $\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
FSL	19848	21796	22547	22688	23579	27026
GAMLSS	19707	21687	22414	22586	23515	26883
SAASL	19679	21663	22372	22554	23443	26883

main goal of this simulation setting is to highlight the imbalance problem of FSL when the scale parameter is large. As many noise variables will make it difficult to demonstrate the differences between FSL and adaptive approaches, we include only two noise variables in this example for illustration.

As can be seen in Fig. 4, both fixed and adaptive step-lengths yield reasonable estimates regarding η_σ , but FSL results in many false negative estimates equal to zero for η_μ in the majority of the simulation runs. This is of course connected to the relative importance of the variance component in this setting, which should in itself already lead to a preference for updating η_σ rather than η_μ in early boosting iterations due to the fact that the negative gradient for μ (i.e. $u_{\mu,i} = \sum_{i=1}^n (y_i - \mu_i) / \sigma_i^2$ with large σ_i) is actually scaled by the variance (recall the large intercept 5, log-link and the resulting exponential transformation) and hence very small. As a consequence, the impact on the global loss of base-learners fit to the gradient is also small compared to those suggested for updates regarding σ in step 11 of Algorithm 1. Then, using the same step-length for both parameters makes it clearly harder to identify informative effects on μ as they are trivialized in comparisons.

The adaptive step-lengths implemented in SAASL compensates for this disadvantage. Compared to the simulation results in the previous subsection the estimates regarding η_μ are less precise with large variability around the true values. This is not a problem of SAASL but again the consequence of the large variance, obscuring the effects on the mean, and also encountered using the penalized maximum likelihood approach implemented in the `gam1ss`-package (called GAMLSS in Fig. 4). The variability in the estimates is actually somewhat smaller than for GAMLSS due to the regularization inherent in the boosting approach. This is also illustrated in Fig. 5 in the pairwise comparison of the estimated coefficients for both methods, where SAASL leads to similar but slightly closer to zero estimates compared to the penalized maximum likelihood based method GAMLSS.

Interestingly, Fig. 4 also reveals that the inability to identify the informative variables results in the lowest MSE for all three individual coefficients for μ when using FSL (for more numerical details, see Appendix C). As can be seen from Table 1, however, the combined additive predictor performs worse in terms of overall MSE than both GAMLSS and SAASL, with the latter performing best.

To further highlight the differences in the selection behavior between FSL and SAASL, Fig. 6 illustrates the proportion of boosting iterations used to update μ over the course of the model fits, i.e. $p_{m_\mu} = m_\mu / (m_\mu + m_\sigma)$, where $m_\mu + m_\sigma = m_{\text{stop}}$. The bimodal distribution for FSL observed in the histogram in panel (a) demonstrates another problem of the fixed step-lengths in this setting. Considering the many esti-

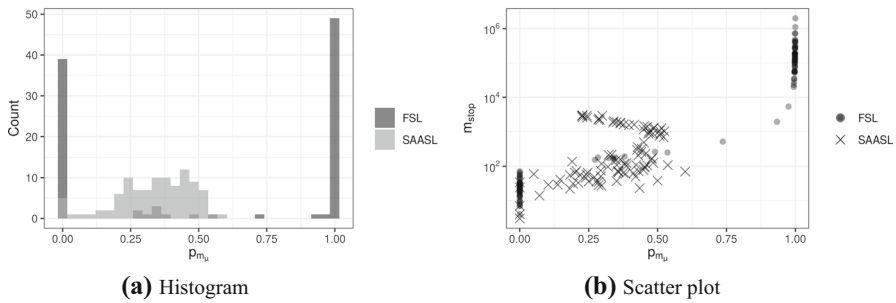


Fig. 6 Distribution of p_{m_μ} in $B = 100$ simulation runs. (a) Histogram of p_{m_μ} . The histogram of the two approaches are overlaid using transparency. (b) Scatter plot of m_{stop} against p_{m_μ} . Points and crosses are displayed with transparency. The y-axis is displayed on a logarithmic scale with base 10. Each tick represents a tenfold increase over the previous one

mates equal or close to zero observed in Fig. 4, the mode close to $p_{m_\mu} = 0$ is expected, as it describes the proportion of simulation runs where μ has not been updated at all. However, as soon as at least one base-learner for μ is recognized as an effective model parameter, the small step-length fixed at 0.1 requires a huge number of updates for the base-learner to actually make an impact on the global loss (hence the large number of maximum iterations allowed for FSL). This results in the second mode also around $p_{m_\mu} = 1$, as the algorithm is mainly occupied with μ in the corresponding runs.

This is illustrated by the scatter plot in Fig. 6b, where p_{m_μ} is plotted against the stopping iteration m_{stop} . Note that the y-axis is displayed with a logarithmic scale and each tick on the y-axis represents a tenfold increase over the previous one. The few points (FSL), whose m_{stop} lie between 10^2 and 10^3 , show a better balance between the updates of μ and σ than other points, i.e., the middle region of p_{m_μ} . But we also observe a bimodal distribution for FSL, i.e., lots of points are equal or close to $p_{m_\mu} = 0$ and 1, with very low and extremely large values for m_{stop} resulting, respectively.

Thus for SAASL, we observe a unimodal distribution of p_{m_μ} in Fig. 6a. The mode smaller than 0.5 indicates SAASL updates σ a little more frequently than μ . Unlike the cyclical approach that enforces an equal number of updates for all distribution parameters, the balance formed by SAASL is more natural. This balance enables the alternate updates between two predictors even though they lie on different scales. Therefore, the information in μ can be fairly discovered in time and it reduces the risk of overlooking the informative base-learners with respect to μ . The number of simulations runs, in which μ is not updated at all ($p_{m_\mu} = 0$), reduces from 39 in FSL to only 5 in SAASL. Moreover, none of the 100 simulations require a substantial amount of updates for μ to get well estimated coefficients (cf. also Fig. 4).

Table 2 displays the information about false positives and false negatives of the two approaches in all 100 simulations with respect to μ and σ . For example, the second and fourth number 77 and 21 in the first line indicate that the informative variable $x_{.2}$ is not included in the final model in 77 out of 100 simulation runs (i.e. false negative), while there are 21 simulations whose final model contains the non-informative variable $x_{.4}$ (i.e. false positive). Similar as Fig. 2 in Sect. 4.1, the conservative small step-length for μ in FSL increases the number of boosting iterations, but reduces the risk of overfitting.

Table 2 The number of simulations with false positives and false negatives for each variable under different modelling methods with respect to the two model parameters. The false negatives part shows the number of simulations in which the informative variables are excluded from the final model, and the false positives part shows how many simulations include the non-informative variables in their final model. Values are taken at the stopping iteration determined by 10-fold CV

		False negatives			False positives	
		$x_{.1}$	$x_{.2}$	$x_{.3}$	$x_{.4}$	$x_{.5}$
μ	FSL	83	77	81	21	20
	SAASL	28	24	28	72	73
σ	FSL	9	1	6	83	82
	SAASL	18	1	9	70	67

Table 3 Comparison of the estimated coefficients

		FSL	ASL	ASL5	SAASL	SAASL05	GAMLSS
(Intercept)	η_{μ}	-772	-169.203	-91.160	-91.160	-91.160	-91.160
	η_{σ}	4.881	4.874	4.912	4.912	4.912	4.912
cBMI	η_{μ}	<0.001	<0.001	-13.925	-13.925	-13.925	-13.926
	η_{σ}	-0.003	-0.003	-0.015	-0.015	-0.015	-0.015
cAge	η_{μ}	-0.038	-0.371	-5.847	-5.847	-5.847	-5.847
	η_{σ}	-0.001	-0.001	0.003	0.003	0.003	0.003
mBMI	η_{μ}	<0.001	<0.001	11.708	11.708	11.708	11.708
	η_{σ}	0.009	0.009	0.009	0.009	0.009	0.009
mAge	η_{μ}	<0.001	<0.001	0.026	0.026	0.026	0.026
	η_{σ}	0.005	0.005	0.005	0.005	0.005	0.005

Less simulations containing the noise variables for μ in FSL than in SAASL confirms this behavior. According to Eq. (11) the ASLs $v_{j^*,\sigma}$ are a sequence of values around 0.05, and (except for the values at early boosting iterations) most of them smaller than 0.1. There are correspondingly slightly more simulations in FSL overfitting the σ -submodel than in SAASL.

Although non-informative variables of μ are excluded from the FSL model, the informative ones are excluded as well. Actually μ was not updated in many simulations at all (cf. Fig. 6a). The false negatives part of Table 2 for μ confirms this. The informative variables $x_{.1}$ to $x_{.3}$ are excluded from the final model in the majority of simulations with FSL but not with SAASL. For σ , the smaller step-length $v_{j^*,\sigma}$ in SAASL selects variables more conservatively and as a consequence slightly more simulations underfit the σ -submodel in SAASL than in FSL, but the difference is far less pronounced.

5 Applications

We apply the proposed algorithms to two datasets. The malnutrition dataset demonstrates the shortcomings of FSL and the pitfalls of using numerical determination of ASL with a fixed search interval, and with the riboflavin dataset we illustrate the variable selection properties of each algorithm.

5.1 Malnutrition of children in India

The first data called `india` from the R package `gamboostLSS` (Hofner et al. 2018; Fahrmeir and Kneib 2011) are sampled from the Standard Demographic and Health Survey between 1998 and 1999 on malnutrition of children in India (Fahrmeir and Kneib 2011). The sample contains 4000 observations and four variables (BMI of the child (`cBMI`), age of the child in months (`cAge`), BMI of the mother (`mBMI`) and age of the mother in years (`mAge`)). The outcome of interest in this case is a numeric z-score for malnutrition ranging from -6 to 6, where the negative values represent malnourished children. To highlight the problems of using a fixed step-length, we work with the original variable `stunting` (corresponding to $100 * z$ -score). The identity and logarithm functions are used as the link functions for μ and σ respectively.

Because this is not a high-dimensional data example, we use the GAMLSS with penalized maximum-likelihood estimation as a gold standard to examine the effectiveness of the adaptive approaches.

Table 3 lists the estimated coefficients of each variable on the predictors η_μ and η_σ at the stopping iteration tuned by 10-folds CV, where the maximum number of iterations is set to 2000. The estimated intercept in η_σ indicates a large variance of the response, with the setting thus being similar to the second simulation above. It is therefore not surprising that FSL selects only one variable (`cAge`) for η_μ , i.e. a large number of updates for the base-learner are required but the given maximal boosting iteration is not large enough. In practice we can certainly increase the maximum number of iterations as well as enlarge the commonly applied step-length 0.1 in order to estimate the coefficients well. But their choices are very subjective and probably result in tedious manual fine-tuning based on trial and error.

The ASL method with the default predefined search interval $[0, 10]$ encounters a similar problem as FSL. Apart from the only selected and underfitted variable `cAge` for μ , the two variables (`cBMI` and `cAge`) for the σ -submodel are also underfitted compared with the results from the gold standard GAMLSS. The reason for this phenomenon lies in the relationship between the variance and step-length discussed in Eq. (5). The log-link or exponential transformation for η_σ in this example data requires a sequence of huge step-lengths, but the default search interval does not fulfill this requirement.

An estimation of ASL by increasing its search interval to $[0, 50000]$, denoted as ASL5 in Table 3, results in coefficients comparable to those of GAMLSS. But choosing a suitable search interval becomes an unavoidable side task for ASL when analyzing this kind of dataset.

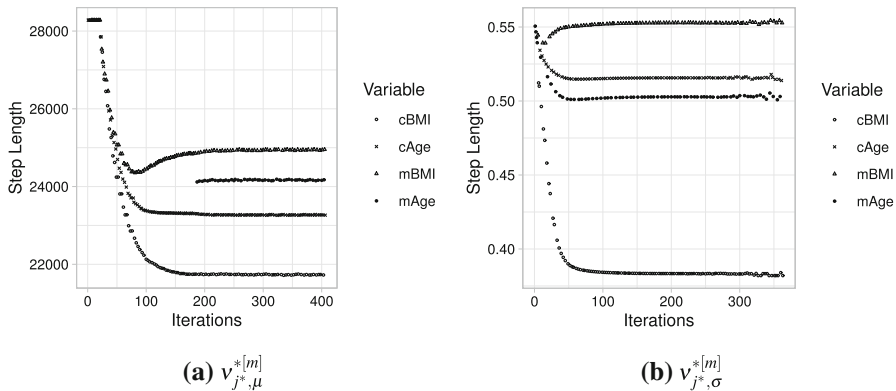


Fig. 7 The optimal step-length of each model parameters against the boosting iterations. Up to the stopping iterations specified by 10-folds CV (here $m_{\text{stop}} = 769$), 406 iterations are used to update μ and 363 iterations are used to update σ

The results of the two semi-analytical approaches hardly differ from the maximum likelihood based GAMLSS. Unlike the numerical determination with a fixed search interval in ASL, the analytical approaches replace this procedure with a direct and precise solution that gets rid of the potential manual intervention (e.g. increasing the search interval). Contrary to the direct influence of the variance on $v_{j^*, \mu}^{*[m]}$ in Eq. (5), the optimal step-length $v_{j^*, \sigma}^{*[m]}$ is dominated by the chosen base-learner, but as the number of learning iterations increases, such effects gradually disappear, and $v_{j^*, \sigma}^{*[m]}$ finally converges to 0.5. Thus, our default search interval $[0, 1]$ is sufficient for $v_{j^*, \mu}^{*[m]}$, and increasing the range of search interval as for $v_{j^*, \mu}^{*[m]}$ in ASL is almost never necessary.

Theoretically, the ASL with a sufficiently large search interval (ASL5 in this example) and SAASL should result in the same values as discussed in the previous theoretical section. Due to the calculation accuracy of computers and the numerical optimization steps, their outputs are very similar but can differ slightly for the malnutrition data.

Figure 7 presents the optimal step-lengths $v_{j^*, \mu}^{*[m]}$ and $v_{j^*, \sigma}^{*[m]}$ using SAASL for each variable up to 769 boosting iterations specified by 10-folds CV for one simulation run. Apparently, the optimal step-lengths for μ over the entire learning process are over 20000, which is far larger than the fixed step-length 0.1 and the upper boundary 10 of the predefined search interval in ASL. Without knowing this information, it is not trivial to determine the search interval for $v_{j^*, \mu}^{*[m]}$. And we thus (after acquiring this graphic) re-estimated the example data with ASL5.

Additionally, Fig. 7b illustrates the optimal step-length for σ . After several boosting iterations the optimal values of each covariate converge to their own stable regions (ranging from about 0.38 to 0.56). As discussed above, the optimal step-lengths for σ should be some values around 0.5, and this graphic confirms this statement.

As this example is not high-dimensional and does not necessarily require variable selection, we can use GAMLSS with penalized maximum likelihood estimation for

Table 4 Number of chosen variables for η_μ and η_σ . The diagonal depicts the number per method, the off-diagonal elements overlapping variables

	η_μ				η_σ			
	FSL	ASL	SAASL	SAASL05	FSL	ASL	SAASL	SAASL05
FSL	13	9	9	11	16	9	9	12
ASL	9	20	20	18	9	17	17	15
SAASL	9	20	20	18	9	17	17	15
SAASL05	11	18	18	24	12	15	15	24

Table 5 Comparison of the out-of-sample MSE

	FSL	ASL	SAASL	SAASL05	glmnet
MSE	2.611	1.111	1.111	1.193	0.946

comparison. The fact that its results are very similar to those of the semi-analytical approaches indicates that results from SAASL and SAASL05 are reliable. The only alternative to achieve balance between predictors would be using a cyclical algorithm (with the downsides discussed in the introduction). Rescaling the response variable or standardizing the negative partial derivatives could reduce the scaling problem to some extent, but would not eliminate the need to increase the step-length or reduce the imbalance between predictors.

5.2 Riboflavin dataset

This data set describes the riboflavin (also known as vitamin B_2) production by *Bacillus subtilis*, containing 71 observations and 4088 predictors (gene expressions) (Bühlmann et al. 2014; Dezeure et al. 2015). The log-transformed riboflavin production rate, which is close to a Gaussian distribution, is regarded as the response. This data set is chosen to demonstrate the capability of the boosting algorithm to deal with situations in which the number of covariates exceeds the number of observations. Please note that a comparison to the original GAMLSS algorithm is not possible in this case, since the algorithm is not able to deal with more model parameters than available observations. In order to compare the out-of-sample MSE of each algorithm, we select 10 observations randomly as the validation set.

Table 4 summarize the selected informative variables for μ and σ separately at the stopping iteration tuned by 5-fold CV, the corresponding coefficients are listed in Appendix D. The results in both tables demonstrate the intersection of the selected variables, for example FSL selects 13 informative variables in total, and 9 of them are also chosen by ASL and SAASL, and there are 11 variables common with SAASL05. In general, for both μ and σ , more variables are included in the adaptive approaches and the difference in the selected variables mainly lies between the adaptive and fixed approach. Because the optimal step-length $v_{j^*,\mu}^{*[m]}$ lies in the predefined search interval $[0, 10]$ (and is actually smaller than 1, i.e. the adaptive step-length $v_{j^*,\mu}^{[m]} < 0.1$),

and $v_{j^*,\sigma}^{*[m]}$ lies also in a narrower predefined search interval $[0, 1]$, ASL and SAASL have the same results. Moreover, as the adaptive step-length is smaller than the fixed step-length 0.1, the adaptive approaches make conservative (small) updates, leading to more boosting iterations. Several of the gene expressions for μ and σ are selected by all algorithms and are thus consistently included in the set of informative covariates. Actually almost all gene expressions chosen by FSL are also recognized as informative variables by all other methods.

To compare the performance of each algorithm, Table 5 lists the out-of-sample MSE. In contrast to the fixed approach, the three adaptive approaches perform in general well, where the performance of SAASL05 is slightly worse than the other two. In addition, Table 5 demonstrates also the result of Lasso estimator from the R package `glmnet` (Friedman et al. 2010) suggested by Bühlmann et al. (2014). The mean squared prediction error of `glmnet` is the smallest among the five approaches, but the difference with the adaptive approaches is relatively small.

As `glmnet` cannot model the scale parameter σ , only the estimated coefficients of the μ -submodel are provided in Appendix D. Out of the 21 genes selected by `glmnet`, 7 and 9 of them are common with the ASL/SAASL and SAASL05, respectively. The signs (positive/negative) of the estimated coefficients of these common covariates from `glmnet` match the adaptive approaches. This comparison indicates that the boosted GAMLSS with adaptive step-length is an applicable and competitive approach for high-dimensional data analysis.

6 Conclusions and outlook

The step-length is often not treated as an important tuning parameter in many boosting algorithms, as long as it is set to a small value. However, if complex models like GAMLSS with several predictors for the different distribution parameters are estimated, the different scales of the distribution parameters can lead to imbalanced updates and resulting bad performances if one common small fixed step-length is used, as we show in this paper.

The main contribution of this article is the proposal to use separate adaptive step-lengths for each distribution parameter in a non-cyclical boosting algorithm for GAMLSS. In addition to the resulting balance in updates between different distribution parameters, a balance between over- and underfitting is obtained by taking only a proportion (shrinkage parameter) such as 10% of the determined optimal step-length as the adaptive step-length. The optimal step-length can be found by optimization procedures such as a line search. We illustrated with an example the importance of updating the search interval for the search if necessary to find the optimal solution.

For the Gaussian location and scale models, we derived an analytical solution for the adaptive step-length for the mean parameter μ , which avoids numerical optimization and specification of a search interval. For the scale parameter σ , we obtained an approximate solution of 0.5 (or 0.05 with 10% proportion), which gives a better motivated default value than 0.1 relative to the step-length for μ , and discussed a combination with a one-dimensional line search in the semi-analytical approach.

In simulations and empirical applications, we showed favorable behavior compared to using a fixed step-length FSL. We showed highly competitive results of our adaptive approaches compared to a standard GAMLSS with respect to estimation accuracy for the low-dimensional case, while the adaptive boosting approach has the advantages of shrinkage and variable selection, which makes it also applicable to the high-dimensional case of more covariates than observations. Overall, the semi-analytical method for adaptive step-length selection performed best among the considered methods.

In this paper we focused on the Gaussian location and scale models to derive analytical or semi-analytical solutions for the optimal step-length, but in most cases, a line search has to be conducted for all distribution parameters. In the future, if possible it is worth investigating analytical adaptive step-lengths for other distributions, because analytical or approximate adaptive step-lengths increase the numerical efficiency and also reveal the relationships between the optimal step-lengths for different parameters and model parameters (as well as properties of commonly used but probably less than ideal step-length settings).

We are confident that the adaptive step-length concept is relevant way beyond the Gaussian specification, so further work should contain the study on the stability and effectiveness of the implementation of adaptive step-length to other common distributions or zero-inflated count models. Further work should also include the implementation of further (e.g. non-linear, spatial etc.) effects (Hothorn et al. 2011) into the model, and test the influence of the adaptive step-length on such effects. Moreover, we discovered correlations between the optimal step-length $v_{j^*}^{*[m]}$ of a variable and the coefficient of this variable in the σ -submodel through our application of the algorithm. Future work should also investigate the relationship among the optimal step-lengths of different parameters and the relationship of these step-lengths to the model coefficients.

A basic R package ASL based on this article is available online at <https://github.com/BoyaoZhang/ASL>. This package contains the source code of Algorithm 3 and the function of the corresponding cross-validation. Some simple examples can also be found in this package. This package is originated from the R package `gamboostLSS`, we hope to implement the functions of ASL into the latter in the future.

Acknowledgements This work was supported by the Freigeist-Fellowships of Volkswagen Stiftung, project “Bayesian Boosting - A new approach to data science, unifying two statistical philosophies”. Boyao Zhang performed the present work in partial fulfilment of the requirements for obtaining the degree “Dr. rer. pol.” at the Georg-August-Universität Göttingen.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendices

A Derive the analytical ASL for the Gaussian distribution

Take the negative log-likelihood as the loss function, the loss for Gaussian distribution can be displayed as

$$\begin{aligned} \rho(\mathbf{y}, \{\eta_\mu, \eta_\sigma\}) &= -\log \left[\frac{1}{(\sqrt{2\pi})^n} \cdot \det(\text{diag}(\exp(-\eta_\sigma(\mathbf{X}))) \cdot \exp\left(-\frac{1}{2}(\mathbf{y} - \eta_\mu(\mathbf{X}))^T \cdot \right. \right. \\ &\quad \left. \left. \cdot \text{diag}(\exp(-2\eta_\sigma(\mathbf{X})) \cdot (\mathbf{y} - \eta_\mu(\mathbf{X}))) \right) \right] \\ &= \frac{n}{2} \log(2\pi) + \mathbf{1}_n^T \eta_\sigma(\mathbf{X}) + \frac{1}{2} (\mathbf{y} - \eta_\mu(\mathbf{X}))^T \text{diag}(\exp(-2\eta_\sigma(\mathbf{X}))) (\mathbf{y} - \eta_\mu(\mathbf{X})). \end{aligned}$$

The negative partial derivatives for both distribution parameters in iteration m are then

$$\mathbf{u}_\mu^{[m]} = - \frac{\partial \rho(\mathbf{y}, \{\hat{\eta}_\mu^{[m-1]}, \hat{\eta}_\sigma^{[m-1]}\})}{\partial \hat{\eta}_\mu} \tag{12}$$

$$= \text{diag}(\exp(-2\hat{\eta}_\sigma^{[m-1]}(\mathbf{X}))) (\mathbf{y} - \hat{\eta}_\mu^{[m-1]}(\mathbf{X})), \tag{13}$$

$$\mathbf{u}_\sigma^{[m]} = - \frac{\partial \rho(\mathbf{y}, \{\hat{\eta}_\mu^{[m-1]}, \hat{\eta}_\sigma^{[m-1]}\})}{\partial \hat{\eta}_\sigma} \tag{14}$$

$$= -\mathbf{1}_n + \text{diag}\left(\left(\mathbf{y} - \hat{\eta}_\mu^{[m-1]}(\mathbf{X})\right)^T\right). \tag{15}$$

$$\cdot \text{diag}(\exp(-2\hat{\eta}_\sigma^{[m-1]}(\mathbf{X}))) (\mathbf{y} - \hat{\eta}_\mu^{[m-1]}(\mathbf{X})). \tag{16}$$

Both $\mathbf{u}_\theta^{[m]}$, $\theta \in \{\mu, \sigma\}$ can be regressed on the simple linear base-learner $h_{j^*, \theta}^{[m]}(\mathbf{x}_{.j^*})$, where j^* denotes the best-fitting variable.

$$\mathbf{u}_\mu^{[m]} = \hat{h}_{j^*, \mu}^{[m]}(\mathbf{x}_{.j^*}) + \hat{\epsilon}_\mu^{[m]} \tag{17}$$

$$\mathbf{u}_\sigma^{[m]} = \hat{h}_{j^*, \sigma}^{[m]}(\mathbf{x}_{.j^*}) + \hat{\epsilon}_\sigma^{[m]}, \tag{18}$$

where $\hat{\epsilon}_\mu^{[m]}$ and $\hat{\epsilon}_\sigma^{[m]}$ denote the residuals in simple linear regression models.

A.1 Optimal step-length for μ

The analytical optimal step-length for μ in iteration m is obtained by minimizing the empirical risk,

$$v_{j^*, \mu}^{*[m]} = \arg \min_v \sum_{i=1}^n \rho(y_i, \{\hat{\eta}_\mu^{[m]}(\mathbf{x}_{i.}), \hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_{i.})\})$$

$$\begin{aligned}
 &= \arg \min_v \sum_{i=1}^n \rho \left(y_i, \{ \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) + v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}), \hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i) \} \right) \\
 &= \arg \min_v \sum_{i=1}^n - \log \left[\frac{1}{\sqrt{2\pi} \exp(\hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i))} \exp \left(- \frac{(y_i - \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2 \exp(2\hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i))} \right) \right] \\
 &= \arg \min_v \sum_{i=1}^n \left[\frac{1}{2} \log(2\pi) + \log(\hat{\sigma}_i^{[m-1]}) + \frac{(y_i - \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2\hat{\sigma}_i^{2[m-1]}} \right] \\
 &= \arg \min_v \sum_{i=1}^n \frac{(y_i - \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2\hat{\sigma}_i^{2[m-1]}}.
 \end{aligned}$$

Note that the expression $\hat{\sigma}_i^{2[m-1]}$ represents the square of the standard deviation in the previous boosting iteration, i.e. $\hat{\sigma}_i^{2[m-1]} = (\hat{\sigma}_i^{[m-1]})^2$. And according to the model specification $\hat{\sigma}_i^{[m-1]} = \exp(\hat{\eta}_\sigma^{[m-1]}(\mathbf{x}_i))$.

It can be shown, that the expression is a convex function, so the optimal value $v_\mu^{*[m]}$ is accessed by letting the first order derivative equal zero,

$$\begin{aligned}
 &\frac{\partial}{\partial v} \sum_{i=1}^n \frac{(y_i - \hat{\eta}_\mu^{[m-1]}(\mathbf{x}_i) - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2\hat{\sigma}_i^{2[m-1]}} \\
 Eq.(13) \quad &\frac{\partial}{\partial v} \sum_{i=1}^n \frac{(u_{\mu,i}^{[m]} \hat{\sigma}_i^{2[m-1]} - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2\hat{\sigma}_i^{2[m-1]}} \\
 &= \frac{\partial}{\partial v} \sum_{i=1}^n \left(\frac{1}{2} u_{\mu,i}^{2[m]} \hat{\sigma}_i^{2[m-1]} - v \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) u_{\mu,i}^{[m]} + \frac{v^2 (\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{2\hat{\sigma}_i^{2[m-1]}} \right) \\
 &= \sum_{i=1}^n \left(-\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) + v \frac{(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{\hat{\sigma}_i^{2[m-1]}} \right) \stackrel{!}{=} 0 \\
 \Leftrightarrow v &= \frac{\sum_{i=1}^n \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) u_{\mu,i}^{[m]}}{\sum_{i=1}^n \frac{(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{\hat{\sigma}_i^{2[m-1]}}} \\
 Eq.(17) \quad &\frac{\sum_{i=1}^n \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) (\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) + \hat{\epsilon}_{\mu,i}^{[m]})}{\sum_{i=1}^n \frac{(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{\hat{\sigma}_i^{2[m-1]}}} \\
 &= \frac{\sum_{i=1}^n (\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2 + \sum_{i=1}^n \hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}) \hat{\epsilon}_{\mu,i}}{\sum_{i=1}^n \frac{(\hat{h}_{j^*, \mu}^{[m]}(x_{ij^*}))^2}{\hat{\sigma}_i^{2[m-1]}}}
 \end{aligned}$$

$$= \frac{\sum_{i=1}^n \left(\hat{h}_{j^*,\mu}^{[m]}(x_{ij^*})\right)^2}{\sum_{i=1}^n \frac{\left(\hat{h}_{j^*,\mu}^{[m]}(x_{ij^*})\right)^2}{\hat{\sigma}_i^{2[m-1]}}}$$

where $\sum_{i=1}^n \hat{h}_{j^*,\mu}^{[m]}(x_{ij^*})\hat{\epsilon}_{\mu,i} = 0$, because the residuals are uncorrelated with the fitted values. □

A.2 Optimal step-length for σ

The analytical optimal step-length for σ in iteration m is obtained by minimizing the empirical risk,

$$\begin{aligned} v_{\sigma}^{*[m]} &= \arg \min_v \sum_{i=1}^n \rho \left(y_i, \{ \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}), \hat{\eta}_{\sigma}^{[m]}(\mathbf{x}_{i\cdot}) \} \right) \\ &= \arg \min_v \sum_{i=1}^n \rho \left(y_i, \{ \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}), \hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \} \right) \\ &= \arg \min_v \sum_{i=1}^n -\log \left[\frac{1}{\sqrt{2\pi} \exp \left(\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)} \right. \\ &\quad \cdot \exp \left(-\frac{\left(y_i - \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)^2}{2 \exp \left(2\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + 2v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)} \right) \left. \right] \\ &= \arg \min_v \sum_{i=1}^n \frac{1}{2} \log(2\pi) + \sum_{i=1}^n \left(\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) + \\ &\quad + \sum_{i=1}^n \frac{\left(y_i - \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)^2}{2 \exp \left(2\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + 2v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)} \\ &= \arg \min_v \sum_{i=1}^n \left(\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) \\ &\quad + \sum_{i=1}^n \frac{\left(y_i - \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)^2}{2 \exp \left(2\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + 2v\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)}. \end{aligned}$$

It can be shown, that the second order derivative of the expression is positive and thus the expression a convex function. Letting the first order derivative equal zero, we get

$$\begin{aligned}
 & \frac{\partial}{\partial v} \left[\sum_{i=1}^n \left(\hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) + \sum_{i=1}^n \frac{\left(y_i - \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)^2}{2 \exp \left(2 \hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) + 2v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)} \right] \\
 &= \sum_{i=1}^n \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \\
 &\quad - \sum_{i=1}^n \left(y_i - \hat{\eta}_{\mu}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)^2 \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \exp \left(-2 \hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) - 2v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) \\
 &\stackrel{Eq.(16)}{=} \sum_{i=1}^n \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \\
 &\quad - \sum_{i=1}^n \frac{u_{\sigma,i}^{[m]} + 1}{\exp \left(-2 \hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) \right)} \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \exp \left(-2 \hat{\eta}_{\sigma}^{[m-1]}(\mathbf{x}_{i\cdot}) - 2v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) \\
 &= \sum_{i=1}^n \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) - \sum_{i=1}^n \left(u_{\sigma,i}^{[m]} + 1 \right) \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \exp \left(-2v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right) \\
 &\stackrel{Eq.(18)}{=} \sum_{i=1}^n \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) - \sum_{i=1}^n \frac{\left(\hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) + \hat{\epsilon}_{\sigma,i}^{[m]} + 1 \right) \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*})}{\exp \left(2v \hat{h}_{j^*,\sigma}^{[m]}(x_{ij^*}) \right)} \stackrel{!}{=} 0
 \end{aligned}$$

□

B Additional simulation graphics

In this appendix, we present the results for some of the simulated examples in Sect. 4.1. Boxplot of the estimated coefficients are showed in Figs 8 and 9. Figure 10 illustrates the negative log-likelihood. The summary of stopping iterations m_{stop} is demonstrated in Fig. 11.




Fig. 8 Boxplot of the estimated coefficients of η_μ in 100 simulation runs. Values are taken at the stopping iterations determined by 10-folds cross-validation. The results are separated according to fixed and adaptive approaches with respect to different non-informative variables settings, i.e. $p_{n-\text{inf}} = 0, 50, 250$ and 500 . The horizontal red lines indicate the true coefficients. The shrinkage of the coefficients towards zero can be observed from this graphic




Fig. 9 Boxplot of the estimated coefficients of η_σ in 100 simulation runs. Values are taken at the stopping iterations tuned by 10-folds cross-validation. The results are separated according to fixed and adaptive approaches with respect to different non-informative variables settings, i.e. $p_{n-\text{inf}} = 0, 50, 250$ and 500 . The horizontal red lines indicate the true coefficients. The shrinkage of the coefficients towards zero can be observed from this graphic




Fig. 10 Summary of the negative log-likelihood of 100 simulation runs with different estimating approaches with respect to various non-informative variables settings. Values are taken at the stopping iteration determined by 10-folds cross-validation

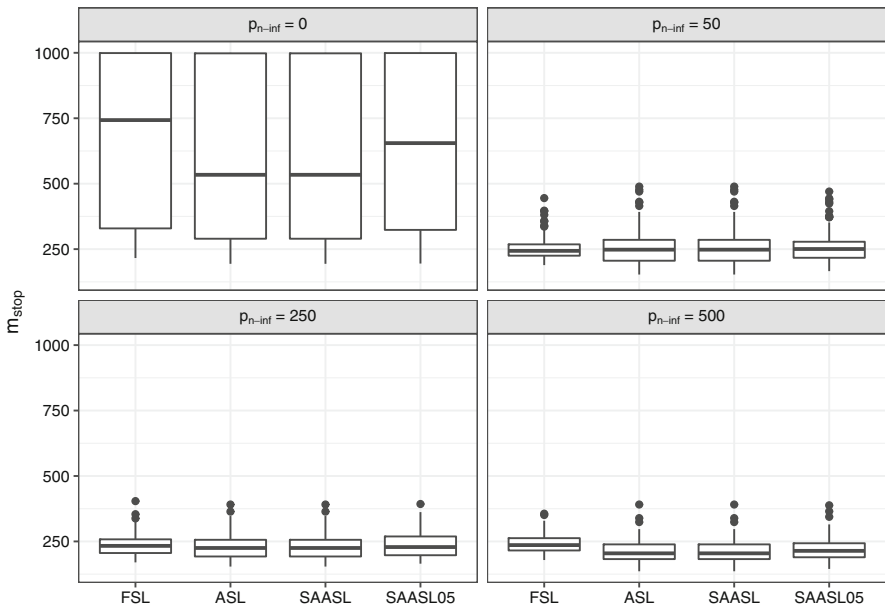


Fig. 11 m_{stop} tuned by 10-fold CV with different estimating methods with respect to different non-informative variables settings. The predefined maximal learning iteration is 1000

Table 6 The average MSE of the estimated coefficients for both model parameters μ and σ w.r.t. three estimation approaches. The MSE for each coefficient is calculated not only from 100 simulation runs (Total) at their stopping iterations but also from the true positive subsets (TP), i.e., the simulations from which a coefficient is selected by all three approaches

	μ				σ							
	$\hat{\beta}_1$ Total	TP	$\hat{\beta}_2$ Total	TP	$\hat{\beta}_3$ Total	TP	$\hat{\beta}_1$ Total	TP	$\hat{\beta}_2$ Total	TP	$\hat{\beta}_3$ Total	TP
FSL	13.7	80.9	26.5	112.2	13.8	76.0	0.84	0.81	1.41	1.42	0.84	0.82
GAMLSS	113.8	328.8	145.6	355.8	116.8	339.4	0.82	0.79	1.44	1.44	0.81	0.79
SAASL	71.3	250.2	95.8	271.7	73.3	238.5	0.85	0.81	1.39	1.40	0.85	0.82

C Additional simulation table

The additional Table 6 summaries the average MSE of the estimated coefficients for both Gaussian distribution parameters in Sect. 4.2.

D Estimated coefficients of riboflavin dataset

In this appendix, we provide the estimated coefficients with fixed and adaptive approaches for riboflavin data in Sect. 5.2. Tables 7 and 8 concern about the μ -submodel and σ -submodel, respectively.

Table 7 The estimated coefficients of the μ -submodel with fixed and adaptive approaches. Values are taken at the m_{stop} tuned by 5-folds CV

	Variable	FSL	ASL	SAASL	SAASL05	glmnet
1	(Intercept)	-7.03	-7.04	-7.04	-7.03	0.72
2	ARGF_at	-0.08	-0.02	-0.02	-0.02	
3	IOLE_at	-0.32			-0.02	
4	LYSC_at					-0.06
5	RPLO_at				-0.02	-0.08
6	SPOIISA_at	0.35	0.19	0.19	0.23	0.02
7	XKDC_at	0.18			0.09	0.19
8	XKDO_at		0.02	0.02	0.03	
9	XKDS_at	0.11	0.08	0.08	0.08	0.06
10	XLYA_at					0.05
11	XTMA_at		0.03	0.03	0.01	
12	XTRA_at					0.01
13	YCDH_at					-0.03
14	YCGM_at	-0.09	-0.06	-0.06	-0.06	-0.01
15	YCGN_at		-0.04	-0.04	-0.04	

Table 7 continued

	Variable	FSL	ASL	SAASL	SAASL05	glmnet
16	YCGO_at	-0.07				-0.14
17	YCGP_at				-0.03	
18	YCKE_at	0.15	0.12	0.12	0.14	0.15
19	YCLB_at	0.29				
20	YCSG_at		-0.08	-0.08	-0.18	
21	YDAO_at				-0.03	
22	YDAR_at	-0.24	-0.16	-0.16	-0.16	
23	YDDK_at					-0.04
24	YEBC_at					-0.55
25	YHAI_at		0.12	0.12	0.11	
26	YHFU_at		-0.02	-0.02	-0.03	-0.01
27	YJCJ_at	0.11	0.04	0.04	0.04	
28	YKBA_at					0.01
29	YKUH_at		0.05	0.05	0.06	
30	YOAB_at					-0.34
31	YORB_i_at		0.03	0.03	0.05	0.10
32	YOZH_i_at		0.02	0.02		
33	YPGA_at					-0.05
34	YTGB_at					-0.09
35	YWQD_at				-0.02	
36	YXJA_at		-0.01	-0.01	-0.01	
37	YXLC_at		-0.03	-0.03		
38	YXLD_at	-0.12	-0.14	-0.14	-0.16	-0.14
39	YXLE_at	-0.06	-0.01	-0.01	-0.01	

Table 8 The estimated coefficients of σ -submodel with fixed and adaptive approaches. Values are taken at the m_{stop} tuned by 5-folds CV

	Variables	FSL	ASL	SAASL	SAASL05
1	(Intercept)	-1.41	-1.29	-1.29	-1.53
2	COTJC_at		-0.18	-0.18	
3	DEGA_at	-0.13	-0.61	-0.61	-0.89
4	EXPZ_at		0.21	0.21	0.05
5	LEVD_at	0.24	0.15	0.15	0.19
6	NTH_at	-0.09	-0.11	-0.11	
7	PHRI_r_at	0.06			0.03
8	TRUA_at	-0.09	-0.74	-0.74	-0.71
9	XLYA_at	-0.06			
10	XPF_at				-0.05
11	YACN_at	0.20			

Table 8 continued

	Variables	FSL	ASL	SAASL	SAASL05
12	YCNK_at	0.66	0.06	0.06	0.27
13	YFIG_at	-0.08			-0.09
14	YFMD_at	-0.25	-0.43	-0.43	-0.35
15	YHBD_at		-0.21	-0.21	-0.19
16	YHEN_at				-0.05
17	YHFS_at				0.06
18	YITQ_at				-0.24
19	YJFB_at		-0.11	-0.11	-0.09
20	YKRS_at		0.29	0.29	0.35
21	YKVV_at		0.06	0.06	0.28
22	YPGA_at	0.07	0.05	0.05	0.12
23	YSBA_at		-0.55	-0.55	-0.28
24	YSBB_at	-0.15	-0.23	-0.23	-0.20
25	YTFP_at	-0.24			
26	YTQI_at	0.10	0.16	0.16	0.11
27	YURR_at		-0.07	-0.07	-0.11
28	YWQA_at	-0.11			-0.20
29	YYAE_at				-0.06
30	YYBT_at	-0.08			-0.03

References

- Brent RP (2013) Algorithms for minimization without derivatives. Courier Corporation
- Bühlmann P (2006) Boosting for high-dimensional linear models. *Ann Stat* 34(2):559–583
- Bühlmann P, Hothorn T (2007) Boosting algorithms: regularization, prediction and model fitting. *Stat Sci* 22(4):477–505
- Bühlmann P, Yu B (2003) Boosting with the L2 loss. *J Am Stat Assoc* 98(462):324–339. <https://doi.org/10.1198/016214503000125>
- Bühlmann P, Kalisch M, Meier L (2014) High-dimensional statistics with a view toward applications in biology. *Annu Rev Stat Appl* 1(1):255–278. <https://doi.org/10.1146/annurev-statistics-022513-115545>
- Dezeure R, Bühlmann P, Meier L, Meinshausen N (2015) High-dimensional inference: Confidence intervals, p-values and R-software HDI. *Stat Sci* 30(4):533–558
- Fahrmeir L, Kneib T (2011) Bayesian smoothing and regression for longitudinal, spatial and event history data. Oxford University Press, Oxford
- Friedman J, Hastie T, Tibshirani R (2000) Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *Ann Stat* 28(2):337–407
- Friedman J, Hastie T, Tibshirani R (2010) Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* 33(1):1–22
- Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Ann Stat* 29(5):1189–1232
- Hastie T, Taylor J, Tibshirani R, Walther G (2007) Forward stagewise regression and the monotone lasso. *Electr J Stat* 1:1–29
- Hepp T, Schmid M, Gefeller O, Waldmann E, Mayr A (2016) Approaches to regularized regression—a comparison between gradient boosting and the Lasso. *Methods Inf Med* 55(5):422–430
- Hepp T, Schmid M, Mayr A (2019) Significance tests for boosted location and scale models with linear base-learners. *Int J Biostat* 15(1):20180110

- Hofner B, Mayr A, Robinzonov N, Schmid M (2014) Model-based boosting in R: a hands-on tutorial using the R package mboost. *Comput Stat* 29(1–2):3–35
- Hofner B, Boccuto L, Göker M (2015) Controlling false discoveries in high-dimensional situations: boosting with stability selection. *BMC Bioinf* 16(1):144
- Hofner B, Mayr A, Schmid M (2016) gamboostlss: an R package for model building and variable selection in the GAMLSS framework. *J Stat Softw Articles* 74(1):1–31
- Hofner B, Mayr A, Fenske N, Schmid M (2018) gamboostLSS: Boosting methods for GAMLSS models. R package version 2.0-1
- Hothorn T, Buehlmann P, Kneib T, Schmid M, Hofner B (2010) Model-based boosting 2.0. *J Mach Learn Res* 11:2109–2113
- Hothorn T, Müller J, Schröder B, Kneib T, Brandl R (2011) Decomposing environmental, spatial, and spatiotemporal components of species distributions. *Ecol Monogr* 81(2):329–347. <https://doi.org/10.1890/10-0602.1>
- Hothorn T, Buehlmann P, Kneib T, Schmid M, Hofner B (2018) mboost: Model-based boosting. R package version 2.9-1
- Kohavi R et al (1995) A study of cross-validation and bootstrap for accuracy estimation and model selection. *Ijcai* 14(2):1137–1145
- Mayr A, Fenske N, Hofner B, Kneib T, Schmid M (2012) Generalized additive models for location, scale and shape for high dimensional data—a flexible approach based on boosting. *J R Stat Soc Ser C Appl Stat* 61(3):403–427. <https://doi.org/10.1111/j.1467-9876.2011.01033.x>
- Meinshausen N, Bühlmann P (2010) Stability selection. *J R Stat Soc Ser B Stat Methodol* 72(4):417–473. <https://doi.org/10.1111/j.1467-9868.2010.00740.x>
- Rigby RA, Stasinopoulos DM (2005) Generalized additive models for location, scale and shape. *J R Stat Soc Ser C Appl Stat* 54(3):507–554. <https://doi.org/10.1111/j.1467-9876.2005.00510.x>
- Ripley BD (2004) Selecting amongst large classes of models. *Methods Models Stat*. https://doi.org/10.1142/9781860945410_0007
- Stasinopoulos D, Rigby R (2007) Generalized additive models for location scale and shape (GAMLSS) in R. *J Stat Softw Articles* 23(7):1–46
- Stasinopoulos D, Rigby R, Heller G, Voudouris V, De Bastiani F (2017) Flexible regression and smoothing: using GAMLSS in R. Chapman and Hall/CRC
- Thomas J, Mayr A, Bischl B, Schmid M, Smith A, Hofner B (2018) Gradient boosting for distributional regression: faster tuning and improved variable selection via noncyclical updates. *Stat Comput* 28(3):673–687
- Zhang T, Yu B (2005) Boosting with early stopping: convergence and consistency. *Ann Stat* 33(4):1538–1579

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.