



HAL
open science

Variable selection in sparse GLARMA models

Marina Gomtsyan, Céline Lévy-Leduc, Sarah Ouadah, Laure Sansonnet,
Thomas Blein

► **To cite this version:**

Marina Gomtsyan, Céline Lévy-Leduc, Sarah Ouadah, Laure Sansonnet, Thomas Blein.
Variable selection in sparse GLARMA models. *Statistics*, 2022, 56 (4), pp.755-784.
10.1080/02331888.2022.2090943 . hal-03763853

HAL Id: hal-03763853

<https://hal.science/hal-03763853v1>

Submitted on 31 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

VARIABLE SELECTION IN SPARSE GLARMA MODELS

M. GOMTSYAN, C. LÉVY-LEDUC, S. OUADAH, L. SANSONNET AND T. BLEIN

ABSTRACT. In this paper, we propose a novel and efficient two-stage variable selection approach for sparse GLARMA models, which are pervasive for modeling discrete-valued time series. Our approach consists in iteratively combining the estimation of the autoregressive moving average (ARMA) coefficients of GLARMA models with regularized methods designed for performing variable selection in regression coefficients of Generalized Linear Models (GLM). We first establish the consistency of the ARMA part coefficient estimators in a specific case. Then, we explain how to efficiently implement our approach. Finally, we assess the performance of our methodology using synthetic data, compare it with alternative methods and illustrate it on an example of real-world application. Our approach, which is implemented in the `GlarmaVarSel` R package and available on the CRAN, is very attractive since it benefits from a low computational load and is able to outperform the other methods in terms of coefficient estimation, particularly in recovering the non null regression coefficients.

1. INTRODUCTION

Discrete-valued time series arise in a wide variety of fields ranging from finance to molecular biology and public health. For instance, we can mention the number of transactions in stocks in the finance field, see [3]. In the field of molecular biology, modeling RNA-Seq kinetics data is a challenging issue, see [27] and in the public health context, there is an interest in the modeling of daily asthma presentations in a given hospital, see [24].

The literature on modeling discrete-valued time series is becoming increasingly abundant, see [8] for a review. Different classes of models have been proposed such as the Integer Autoregressive Moving Average (INARMA) models and the generalized state space models.

The Integer Autoregressive process of order 1 (INAR(1)) was first introduced by [19] and the Integer-valued Moving Average (INMA) process is described in [1]. One of the attractive features of INARMA processes is that their autocorrelation structure is similar to the one of autoregressive moving average (ARMA) models. However, it has to be noticed that statistical inference in these models is generally complicated and requires to develop intensive computational approaches such as the efficient MCMC algorithm devised by [22] for INARMA processes of known AR and MA orders. This strategy was extended to unknown AR and MA orders by [11]. For further references on INARMA models, we refer the reader to [28].

The other important class of models for discrete-valued time series is the one of generalized state space models which can have a parameter-driven and an observation-driven

Key words and phrases. GLARMA models; sparse; discrete-valued time series.

version, see [7] for a review. The main difference between these two versions is that in parameter-driven models, the state vector evolves independently of the past history of the observations whereas the state vector depends on the past observations in observation-driven models. More precisely, in parameter-driven models, let (ν_t) be a stationary process, the observations Y_t are thus modeled as follows: conditionally on (ν_t) , Y_t has a Poisson distribution of parameter $\exp(\beta_0^* + \sum_{i=1}^p \beta_i^* x_{t,i} + \nu_t)$, where the $x_{t,i}$'s are the p regressor variables (or covariates). Estimating the parameters in such models has a very high computational load, see [18].

Observation-driven models initially proposed by [4] and further studied in [30] do not have this computational drawback and are thus considered as a promising alternative to parameter-driven models. Different kinds of observation-driven models can be found in the literature: the Generalized Linear Autoregressive Moving Average (GLARMA) models introduced by [7] and further studied in [5], [6], [9] and the (log-)linear Poisson autoregressive models studied in [12], [14] and [13]. Note that GLARMA models cannot be seen as a particular case of the log-linear Poisson autoregressive models.

In the following, we shall consider the GLARMA model introduced in [6] with additional covariates. More precisely, given the past history $\mathcal{F}_{t-1} = \sigma(Y_s, s \leq t-1)$, we assume that

$$Y_t | \mathcal{F}_{t-1} \sim \mathcal{P}(\mu_t^*), \quad (1)$$

where $\mathcal{P}(\mu)$ denotes the Poisson distribution with mean μ . In (1),

$$\mu_t^* = \exp(W_t^*) \text{ with } W_t^* = \beta_0^* + \sum_{i=1}^p \beta_i^* x_{t,i} + Z_t^*, \quad (2)$$

where the $x_{t,i}$'s are the p regressor variables ($p \geq 1$),

$$Z_t^* = \sum_{j=1}^q \gamma_j^* E_{t-j}^* \text{ with } E_t^* = \frac{Y_t - \mu_t^*}{\mu_t^*} = Y_t \exp(-W_t^*) - 1, \quad (3)$$

with $1 \leq q \leq \infty$ and $E_t^* = 0$ for all $t \leq 0$. Here, the E_t^* 's correspond to the working residuals in classical Generalized Linear Models (GLM), which means that we limit ourselves to the case $\lambda = 1$ in the more general definition: $E_t^* = (Y_t - \mu_t^*)\mu_t^{*\lambda}$. Note that in the case where $q = \infty$, (Z_t^*) satisfies the ARMA-like recursions given in Equation (4) of [6]. The model defined by (1), (2) and (3) is thus referred as a GLARMA model.

The main goal of this paper is to introduce a novel variable selection approach in the deterministic part (covariates) of sparse GLARMA models that is in (1), (2) and (3) where the vector of the β_i^* 's is sparse. Sparsity means that many β_i^* 's are null and thus just a few of regressor variables are explanatory. The novel approach that we propose consists in combining a procedure for estimating the ARMA part coefficients to take into account the temporal dependence that may exist in the data with regularized methods designed for GLM as those proposed by [15] and [17]. Our procedure can be useful for modeling RNA-Seq time series data, sometimes referred to as RNA-Seq kinetics data in molecular biology. It allows monitoring the entire gene expression inside a biological sample along a time course. In this application, as explained by [29], non-coding genes are emerging as

potential key regulators of the expression of coding genes, namely the part of the genes coding for proteins. In this framework, only a few among a lot of non-coding genes are likely to be involved for explaining the expression of the coding genes. Hence, designing a variable selection approach for sparse GLARMA models will allow us to identify the relevant non-coding genes. Note that existing variable selection approaches for discrete observations such as [15] indeed do not take into account the temporal dependence that may exist in this kind of data.

The paper is organized as follows. Firstly, in Section 2.1, we describe the classical estimation procedure in GLARMA models and in Section 2.4, establish a consistency result in a specific case. Secondly, we propose a novel two-stage estimation procedure which is described in Section 2.2. It consists in first estimating the ARMA coefficients and then in estimating the regression coefficients by using a regularized approach. The practical implementation of our approach is given in Section 2.3. The R language implementation of the method is provided in the `GlarmaVarSel` package which is available on the CRAN. Thirdly, in Section 3, we provide some numerical experiments to illustrate our method and to compare its performance to alternative approaches on finite sample size data. More precisely, we compared our approach to two different methods: the regularized methods designed for GLM of [15] and the standard estimation procedure in non necessarily sparse GLARMA models implemented in the R `glarma` package. Additionally, in Section 4, we illustrate our method on RNA-Seq time series that follows the temporal evolution of gene expression. Finally, we give the proofs of the theoretical results in Section 5.

2. STATISTICAL INFERENCE

2.1. Classical estimation procedure in GLARMA models. As explained by [6], the parameter $\boldsymbol{\delta}^* = (\boldsymbol{\beta}^{*'}, \boldsymbol{\gamma}^{*'})$, u' denoting the transpose of u , can be estimated by using the following criterion based on the conditional log-likelihood, where $\boldsymbol{\beta}^* = (\beta_0^*, \beta_1^*, \dots, \beta_p^*)'$ is the vector of regressor coefficients defined in (2) and $\boldsymbol{\gamma}^* = (\gamma_1^*, \dots, \gamma_q^*)'$ is the vector of the ARMA part coefficients defined in (3). This criterion consists in maximizing with respect to $\boldsymbol{\delta} = (\boldsymbol{\beta}', \boldsymbol{\gamma}')$, with $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_p)'$ and $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_q)'$:

$$L(\boldsymbol{\delta}) = \sum_{t=1}^n (Y_t W_t(\boldsymbol{\delta}) - \exp(W_t(\boldsymbol{\delta}))). \quad (4)$$

In (4),

$$W_t(\boldsymbol{\delta}) = \boldsymbol{\beta}' x_t + Z_t(\boldsymbol{\delta}) = \beta_0 + \sum_{i=1}^p \beta_i x_{t,i} + \sum_{j=1}^q \gamma_j E_{t-j}(\boldsymbol{\delta}), \quad (5)$$

with $x_t = (x_{t,0}, x_{t,1}, \dots, x_{t,p})'$, $x_{t,0} = 1$ for all t and

$$E_t(\boldsymbol{\delta}) = Y_t \exp(-W_t(\boldsymbol{\delta})) - 1, \text{ if } t > 0 \text{ and } E_t(\boldsymbol{\delta}) = 0, \text{ if } t \leq 0. \quad (6)$$

To obtain $\widehat{\boldsymbol{\delta}}$ defined by

$$\widehat{\boldsymbol{\delta}} = \text{Argmax}_{\boldsymbol{\delta}} L(\boldsymbol{\delta}),$$

the first derivatives of L are considered:

$$\frac{\partial L}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}) = \sum_{t=1}^n (Y_t - \exp(W_t(\boldsymbol{\delta}))) \frac{\partial W_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}), \quad (7)$$

where

$$\frac{\partial W_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}) = \frac{\partial \boldsymbol{\beta}' x_t}{\partial \boldsymbol{\delta}} + \frac{\partial Z_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}),$$

$\boldsymbol{\beta}$, x_t and Z_t being given in (5). The computations of the first derivatives of W_t are detailed in Section 5.1.1.

Based on Equation (7) which is non linear in $\boldsymbol{\delta}$ and which has to be recursively computed, it is not possible to obtain a closed-form formula for $\widehat{\boldsymbol{\delta}}$. Thus $\widehat{\boldsymbol{\delta}}$ is computed by using the Newton-Raphson algorithm. More precisely, starting from an initial value for $\boldsymbol{\delta}$ denoted by $\boldsymbol{\delta}^{(0)}$, the following recursion for $r \geq 1$ is used:

$$\boldsymbol{\delta}^{(r)} = \boldsymbol{\delta}^{(r-1)} - \frac{\partial^2 L}{\partial \boldsymbol{\delta}' \partial \boldsymbol{\delta}}(\boldsymbol{\delta}^{(r-1)})^{-1} \frac{\partial L}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}^{(r-1)}), \quad (8)$$

where $\frac{\partial^2 L}{\partial \boldsymbol{\delta}' \partial \boldsymbol{\delta}}$ corresponds to the Hessian matrix of L and is defined in (9) given below. Hence, it requires the computation of the first and second derivatives of L . We already explained how to compute the first derivatives of L . As for the second derivatives of L , it can be obtained as follows:

$$\frac{\partial^2 L}{\partial \boldsymbol{\delta}' \partial \boldsymbol{\delta}}(\boldsymbol{\delta}) = \sum_{t=1}^n (Y_t - \exp(W_t(\boldsymbol{\delta}))) \frac{\partial^2 W_t}{\partial \boldsymbol{\delta}' \partial \boldsymbol{\delta}}(\boldsymbol{\delta}) - \sum_{t=1}^n \exp(W_t(\boldsymbol{\delta})) \frac{\partial W_t}{\partial \boldsymbol{\delta}'}(\boldsymbol{\delta}) \frac{\partial W_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}). \quad (9)$$

The computations of the second derivatives of W_t are detailed in Section 5.1.2.

However, in our sparse framework where many components of $\boldsymbol{\beta}^*$ are null, this procedure provides poor estimation results, see Section 3.1.2 for numerical illustration. This is the reason why we devised a novel estimation procedure described in the next section.

2.2. Our estimation procedure. For selecting the most relevant components of $\boldsymbol{\beta}^*$, we propose the following two-stage procedure: Firstly, we estimate $\boldsymbol{\gamma}^*$ by using the Newton-Raphson algorithm described in Section 2.2.1 and secondly, we estimate $\boldsymbol{\beta}^*$ by using the regularized approach detailed in Section 2.2.2.

2.2.1. Estimation of $\boldsymbol{\gamma}^*$. To estimate $\boldsymbol{\gamma}^*$, we propose using

$$\widehat{\boldsymbol{\gamma}} = \text{Argmax}_{\boldsymbol{\gamma}} L(\boldsymbol{\beta}^{(0)'}, \boldsymbol{\gamma}'),$$

where L is defined in (4), $\boldsymbol{\beta}^{(0)} = (\beta_0^{(0)}, \dots, \beta_p^{(0)})'$ is a given initial value for $\boldsymbol{\beta}^*$ and $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_q)'$. Similar to the approach proposed in Section 2.1, we use the Newton-Raphson algorithm to obtain $\widehat{\boldsymbol{\gamma}}$ based on the following recursion for $r \geq 1$ starting from the initial value $\boldsymbol{\gamma}^{(0)} = (\gamma_1^{(0)}, \dots, \gamma_q^{(0)})'$:

$$\boldsymbol{\gamma}^{(r)} = \boldsymbol{\gamma}^{(r-1)} - \frac{\partial^2 L}{\partial \boldsymbol{\gamma}' \partial \boldsymbol{\gamma}}(\boldsymbol{\beta}^{(0)'}, \boldsymbol{\gamma}^{(r-1)'})^{-1} \frac{\partial L}{\partial \boldsymbol{\gamma}}(\boldsymbol{\beta}^{(0)'}, \boldsymbol{\gamma}^{(r-1)}), \quad (10)$$

where the first and second derivatives of L are obtained using the same strategy as the one used for deriving Equations (7) and (9) in Section 2.1.

2.2.2. Variable selection: Estimation of β^* . To perform variable selection in the β_i^* of Model (2) aimed to obtain a sparse estimator of β_i^* , we shall use a methodology inspired by [15] for fitting generalized linear models with ℓ_1 penalties. It consists in penalizing a quadratic approximation to the log-likelihood obtained by a Taylor expansion. Using $\beta^{(0)}$ and $\hat{\gamma}$ defined in Section 2.2.1, the quadratic approximation is obtained as follows:

$$\begin{aligned} \tilde{L}(\beta) &:= L(\beta_0, \dots, \beta_p, \hat{\gamma}) \\ &= \tilde{L}(\beta^{(0)}) + \frac{\partial L}{\partial \beta}(\beta^{(0)}, \hat{\gamma})(\beta - \beta^{(0)}) + \frac{1}{2}(\beta - \beta^{(0)})' \frac{\partial^2 L}{\partial \beta \partial \beta'}(\beta^{(0)}, \hat{\gamma})(\beta - \beta^{(0)}), \end{aligned}$$

where

$$\frac{\partial L}{\partial \beta} = \left(\frac{\partial L}{\partial \beta_0}, \dots, \frac{\partial L}{\partial \beta_p} \right) \text{ and } \frac{\partial^2 L}{\partial \beta \partial \beta'} = \left(\frac{\partial^2 L}{\partial \beta_j \partial \beta_k} \right)_{0 \leq j, k \leq p}.$$

Thus,

$$\tilde{L}(\beta) = \tilde{L}(\beta^{(0)}) + \frac{\partial L}{\partial \beta}(\beta^{(0)}, \hat{\gamma})U(\nu - \nu^{(0)}) - \frac{1}{2}(\nu - \nu^{(0)})' \Lambda(\nu - \nu^{(0)}), \quad (11)$$

where $U\Lambda U'$ is the singular value decomposition of the positive semidefinite symmetric matrix $-\frac{\partial^2 L}{\partial \beta \partial \beta'}(\beta^{(0)}, \hat{\gamma})$ and $\nu - \nu^{(0)} = U'(\beta - \beta^{(0)})$.

In order to obtain a sparse estimator of β^* , we propose using $\hat{\beta}(\lambda)$ defined by

$$\hat{\beta}(\lambda) = \text{Argmin}_{\beta} \left\{ -\tilde{L}_Q(\beta) + \lambda \|\beta\|_1 \right\}, \quad (12)$$

for a positive λ , where $\|\beta\|_1 = \sum_{k=0}^p |\beta_k|$ and $\tilde{L}_Q(\beta)$ denotes the quadratic approximation of the log-likelihood. This quadratic approximation is defined by

$$-\tilde{L}_Q(\beta) = \frac{1}{2} \|\mathcal{Y} - \mathcal{X}\beta\|_2^2, \quad (13)$$

with

$$\mathcal{Y} = \Lambda^{1/2} U' \beta^{(0)} + \Lambda^{-1/2} U' \left(\frac{\partial L}{\partial \beta}(\beta^{(0)}, \hat{\gamma}) \right)', \quad \mathcal{X} = \Lambda^{1/2} U' \quad (14)$$

and $\|\cdot\|_2$ denoting the ℓ_2 norm in \mathbb{R}^{p+1} . Computational details for obtaining the expression (13) of $\tilde{L}_Q(\beta)$ appearing in Criterion (12) are provided in Section 5.2.

To obtain the final estimator $\hat{\beta}$ of β^* , we shall consider two different approaches:

- **Standard stability selection.** It consists in using the stability selection procedure devised by [20] which guarantees the robustness of the selected variables. This approach can be described as follows. The vector \mathcal{Y} defined in (14) is randomly split into several subsamples of size $(p+1)/2$, which corresponds to half of the length of \mathcal{Y} . The number of subsamples is equal to 1000 in our numerical experiments. For each subsample $\mathcal{Y}^{(s)}$ and the corresponding design matrix $\mathcal{X}^{(s)}$, the LASSO criterion (12) is applied with a given λ , where \mathcal{Y} and \mathcal{X} are replaced by $\mathcal{Y}^{(s)}$ and

$\mathcal{X}^{(s)}$, respectively. For each subsampling, the indices i of the non null $\widehat{\beta}_i$ are stored. At the end, we calculate a frequency of index selection, namely the amount of times each index was selected divided by the number of subsamples. For a given threshold, we keep in the final set of selected variables the ones whose indices have a frequency larger than this threshold. Concerning the choice of λ , we shall consider the one obtained by cross-validation (Chapter 7 of [16]) called `ss_cv` in the following and the smallest element of the grid of λ provided by the R `glmnet` package called `ss_min` in the following.

- **Fast stability selection.** It consists in applying the LASSO criterion (12) for several values of λ . For each λ , the indices i of the non null $\widehat{\beta}_i(\lambda)$ are stored. Then, we calculate a frequency of index selection, namely the amount of times each index was selected divided by the number of λ 's considered. For a given threshold, we keep in the final set of selected variables the ones whose indices have a frequency larger than this threshold. This approach is called `fast_ss` in the following.

These approaches will be further investigated in Section 3.

2.3. Practical implementation. In practice, the previous approach can be summarized as follows.

- **Initialization.** We take for $\beta^{(0)}$ the estimator of β^* obtained by fitting a GLM to the observations Y_1, \dots, Y_n thus ignoring the ARMA part of the model in the case where $n > p$. If p is larger than n , then a regularized criterion for GLM models can be used, see for instance [15]. For $\gamma^{(0)}$, we take the null vector.
- **Newton-Raphson algorithm.** We use the recursion defined in (10) with the initialization $(\beta^{(0)}, \gamma^{(0)})$ obtained in the previous step and we stop at the iteration R such that $\|\gamma^{(R)} - \gamma^{(R-1)}\|_\infty < 10^{-6}$.
- **Variable selection.** To obtain a sparse estimator of β^* , we use the criterion (12) where $\beta^{(0)}$ and $\widehat{\gamma}$ appearing in (14) are replaced by $\beta^{(0)}$ and $\gamma^{(R)}$ obtained in the previous steps. We thus get $\widehat{\beta}$ by using one of the three approaches described at the end of Section 2.2.2.

This procedure can be improved by iterating the Newton-Raphson algorithm and Variable selection steps. More precisely, let us denote by $\beta_1^{(0)}$, $\gamma_1^{(R_1)}$ and $\widehat{\beta}_1$ the values of $\beta^{(0)}$, $\gamma^{(R)}$ and $\widehat{\beta}$ obtained in the three steps described above at the first iteration. At the second iteration, $(\beta^{(0)}, \gamma^{(0)})$ appearing in the Newton-Raphson algorithm step is replaced by $(\widehat{\beta}_1, \gamma_1^{(R_1)})$. At the end of this second iteration, $\widehat{\beta}_2$ and $\gamma_2^{(R_2)}$ denote the obtained values of $\widehat{\beta}$ and $\gamma^{(R)}$, respectively. This approach is iterated until the stabilization of $\gamma_k^{(R_k)}$.

2.4. Consistency results. In this section, we shall establish the consistency of the parameter γ_1^* in the case where $q = 1$ from Y_1, \dots, Y_n defined in (1) and (3) where (2) is replaced by

$$\mu_t^* = \exp(W_t^*) \text{ with } W_t^* = \beta_0^* + Z_t^*. \quad (15)$$

We limit ourselves to this framework since in the more general one the consistency is much more tricky to handle and is beyond the scope of this paper. Note that some theoretical

results have already been obtained in this framework (no covariates and $q = 1$) by [5] and [6]. However, here, we provide, on the one hand, a more detailed version of the proof of these results and on the other hand, a proof of the consistency of γ_1^* based on a stochastic equicontinuity result.

Theorem 2.1. *Assume that Y_1, \dots, Y_n satisfy the model defined by (1), (15) and (3) with $q = 1$ and $\gamma_1^* \in \Gamma$ where Γ is a compact set of \mathbb{R} which does not contain 0. Assume also that (W_t^*) starts with its stationary invariant distribution. Let $\hat{\gamma}_1$ be defined by:*

$$\hat{\gamma}_1 = \text{Argmax}_{\gamma_1 \in \Gamma} L(\beta_0^*, \gamma_1),$$

where

$$L(\beta_0^*, \gamma_1) = \sum_{t=1}^n (Y_t W_t(\beta_0^*, \gamma_1) - \exp(W_t(\beta_0^*, \gamma_1))), \quad (16)$$

with

$$W_t(\beta_0^*, \gamma_1) = \beta_0^* + Z_t(\gamma_1) = \beta_0^* + \gamma_1 E_{t-1}(\gamma_1), \quad (17)$$

$$E_{t-1}(\gamma_1) = Y_{t-1} \exp(-W_{t-1}(\beta_0^*, \gamma_1)) - 1, \text{ if } t > 1 \text{ and } E_{t-1}(\gamma_1) = 0, \text{ if } t \leq 1.$$

Then $\hat{\gamma}_1 \xrightarrow{p} \gamma_1^*$, as n tends to infinity, where \xrightarrow{p} denotes the convergence in probability.

The proof of Theorem 2.1 is based on the following propositions which are proved in Section 5. These propositions are the classical arguments for establishing consistency results of maximum likelihood estimators. Note that we shall explain in the proof of Proposition 2.2 why a stationary invariant distribution for (W_t^*) does exist. The main tools used for proving Propositions 2.2 and 2.4 are the Markov property and the ergodicity of (W_t^*) .

Proposition 2.2. *For all fixed γ_1 , under the assumptions of Theorem 2.1,*

$$\frac{1}{n} L(\beta_0^*, \gamma_1) \xrightarrow{p} \mathcal{L}(\gamma_1) := \mathbb{E} [Y_3 W_3(\beta_0^*, \gamma_1) - \exp(W_3(\beta_0^*, \gamma_1))], \text{ as } n \text{ tends to infinity.} \quad (18)$$

Proposition 2.3. *The function \mathcal{L} defined in (18) has a unique maximum at the true parameter $\gamma_1 = \gamma_1^*$.*

Proposition 2.4. *Under the assumptions of Theorem 2.1*

$$\sup_{\gamma_1 \in \Gamma} \left| \frac{L(\beta_0^*, \gamma_1)}{n} - \mathcal{L}(\gamma_1) \right| \xrightarrow{p} 0, \text{ as } n \text{ tends to infinity,}$$

where $\mathcal{L}(\gamma_1)$ is defined in (18).

3. NUMERICAL EXPERIMENTS

This section aims to investigate the performance of our method, the implementation of which is available in the R package `GLarmaVarSel`. We study it both from a statistical and a numerical point of view, using synthetic data generated by the model defined by (1), (2), and (3).

3.1. Statistical performance.

3.1.1. *Estimation of the parameters when $p = 0$.* In this section, we investigate the statistical performance of our methodology in the model defined by (1), (2) and (3) for n in $\{50, 100, 250, 500, 1000\}$ in the case where $p = 0$, namely when there are no covariates and for q in $\{1, 2, 3\}$. The performance of our approach for estimating β_0^* and the γ_k^* are displayed in Figures 1, 2 and 3. We can see from these figures that the accuracy of the parameter estimations is improved when n increases, which corroborates the consistency of γ_1^* given in Theorem 2.1 in the case $q = 1$.




FIGURE 1. Boxplots for the estimations of $\beta_0^* = 3$ in Model (2) with no regressor and $q = 1$ (left), $q = 2$ (middle) and $q = 3$ (right). The horizontal lines correspond to the value of β_0^* .




FIGURE 2. Boxplots for the estimations of $\gamma_1^* = 0.5$ in Model (2) with no regressor and $q = 1$ (left), $q = 2$ (middle) and $q = 3$ (right). The horizontal lines correspond to the value of γ_1^* .

Moreover, it has to be noticed that in this particular context where there are no covariates ($p = 0$), the performance of our approach in terms of parameters estimation is similar to the one of the package `glarma` described in [10].




FIGURE 3. Boxplots for the estimations of $\gamma_2^* = 1/4$ in Model (2) with no regressor and $q = 2$ (left), $\gamma_2^* = 1/3$ in Model (2) with no regressor and $q = 3$ (middle) and of $\gamma_3^* = 1/4$ in Model (2) with no regressor and $q = 3$ (right). The horizontal lines correspond to the true values of the parameters.

3.1.2. *Estimation of the parameters when $p \geq 1$ and β^* is sparse.* In this section, we assess the performance of our methodology in terms of support recovery, namely the identification of the non null coefficients of β^* , and of the estimation of γ^* . We shall consider Y_1, \dots, Y_n satisfying the model defined by (1), (2) and (3) with covariates chosen in a Fourier basis defined by $x_{t,i} = \cos(2\pi itf/n)$, when $i = 1, \dots, [p/2]$ and $x_{t,i} = \sin(2\pi itf/n)$, when $i = [p/2] + 1, \dots, p$, with $t = 1, \dots, n$ and $f = 0.7$. Note that $[x]$ denotes the integer part of x . Here $n = 1000$ in the first two paragraphs (“Estimation of the support of β^* ” and “Estimation of γ^* ”), $q \in \{1, 2, 3\}$, $p = 100$ and two sparsity levels (5% or 10% of non null coefficients in β^*). More precisely, when the sparsity level is 5% (resp. 10%) all the β_i^* ’s are assumed to be equal to zero except for five (resp. ten) of them: $\beta_1^* = 1.73$, $\beta_3^* = 0.38$, $\beta_{17}^* = 0.29$, $\beta_{33}^* = -0.64$ and $\beta_{44}^* = -0.13$ (resp. $\beta_1^* = 1.73$, $\beta_3^* = 1.2$, $\beta_5^* = 0.67$, $\beta_{10}^* = 0.5$, $\beta_{14}^* = -0.38$, $\beta_{17}^* = 0.29$, $\beta_{30}^* = -0.64$, $\beta_{33}^* = -0.13$, $\beta_{38}^* = -0.1$ and $\beta_{44}^* = -0.07$). Other values of n (150, 200, 500, 1000) will be considered in the third paragraph (“Impact of the value of n ”) to evaluate the impact of n on the performance of our approach.

Estimation of the support of β^*

In this paragraph, we focus on the performance of our approach for retrieving the support of β^* by computing the TPR (True Positive Rates, namely the proportion of non-null coefficients correctly estimated as non null) and FPR (False Positive Rates, namely the proportion of null coefficients estimated as non null). We shall consider the two methods that are proposed in Section 2.2.2: standard stability selection (`ss_cv` and `ss_min`) and fast stability selection (`fast_ss`). For comparison purpose, we shall also consider the standard Lasso approach proposed by [15] in GLM where the parameter λ is either chosen thanks

to the standard cross-validation (`lasso_cv`) or by taking the optimal λ which maximizes the difference between the TPR and FPR (`lasso_best`).

Figures 4, 5 and 6 display the TPR and FPR of the previously mentioned approaches with respect to the threshold defined at the end of Section 2.2.2 when $n = 1000$, the sparsity level is equal to 5% and $q = 1, 2$ and 3 , respectively. We can see from these figures that when the threshold is well tuned, our approaches outperform the classical Lasso even when the parameter λ is chosen in an optimal way. More precisely, the thresholds 0.4, 0.7 and 0.8 achieve a satisfactory trade-off between the TPR and the FPR for `fast_ss`, `ss_cv` and `ss_min`, respectively. The conclusions are similar in the case where the sparsity level is equal to 10%, the corresponding figures (19, 20 and 21) are given in the Appendix. We can observe from these figures that the performance of `fast_ss` are slightly better than `ss_cv` and `ss_min` when the sparsity level is equal to 5% but it is the reverse when the sparsity level is equal to 10%.




FIGURE 4. Error bars of the TPR and FPR associated to the support recovery of β^* for five methods with respect to the thresholds when $n = 1000$, $q = 1$, $p = 100$ and a 5% sparsity level. The error bars of the TPR of `lasso_cv` and `lasso_best` coincide.

We also compare our approach with the method implemented in the `glarma` package of [10] in the case where $q = 1$ and when the sparsity level is equal to 5%. Since this method is not devised for performing variable selection, we consider that a given component of β^*

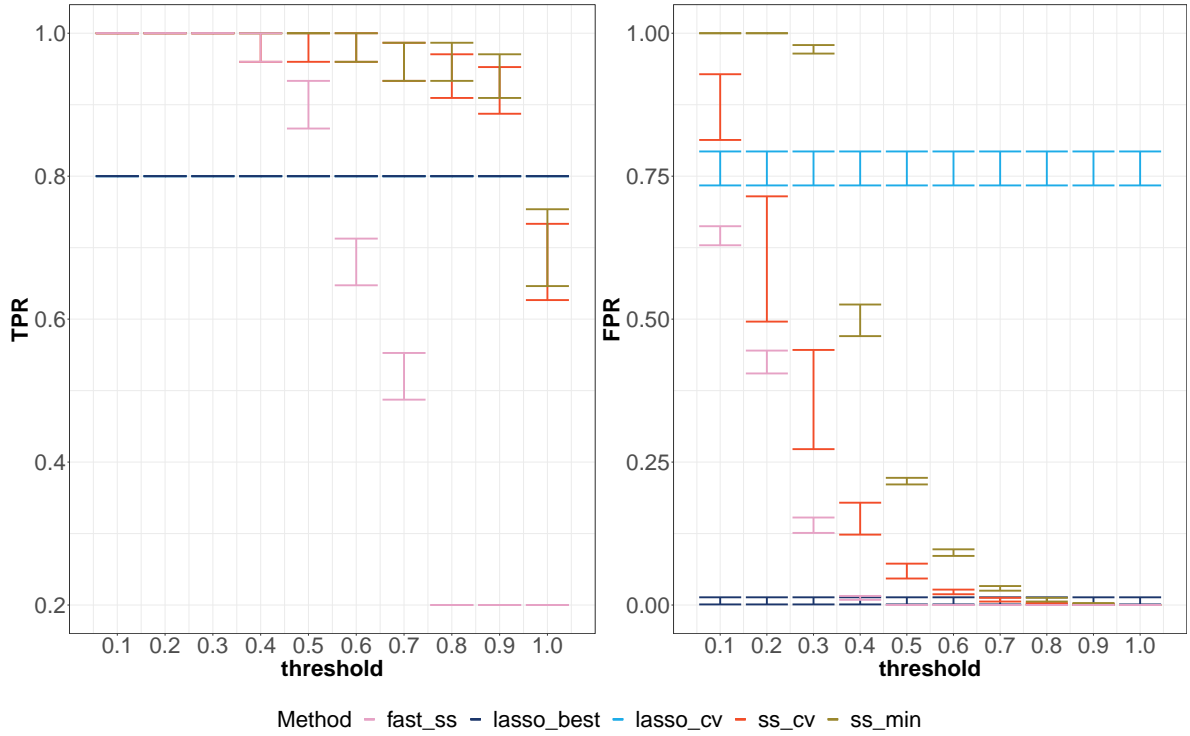


FIGURE 5. Error bars of the TPR and FPR associated to the support recovery of β^* for five methods with respect to the thresholds when $n = 1000$, $q = 2$, $p = 100$ and a 5% sparsity level. The error bars of the TPR of `lasso_cv` and `lasso_best` coincide.

is estimated by 0 if its estimation obtained by the `glarma` package is smaller than a given threshold. The results are displayed in Figure 7 for different thresholds ranging from 10^{-9} to 0.1. We can see from this figure that for the best choice of the threshold the results of the variable selection provided by the `glarma` package underperform our method.

Estimation of γ^*

Figures 8, 9 and 10 display the boxplots for the estimations of γ^* in Model (2) with a 5% sparsity level and $q = 1, 2, 3$ obtained by `ss_cv`, `fast_ss` and `ss_min`, respectively. The threshold chosen for each of these methods is the one achieving a satisfactory trade-off between the TPR and the FPR, namely 0.7, 0.4 and 0.8. We can see from these figures that all these approaches provide accurate estimations of γ^* from the second iteration. The conclusions are similar in the case where the sparsity level is equal to 10%, the corresponding figures 22, 23 and 24 are given in the Appendix.

Impact of the value of n

In this paragraph, we study the impact of the value of n on the TPR and the FPR associated to the support recovery of β^* and on the estimation of γ^* for `ss_min`, the other approaches providing similar results.




FIGURE 6. Error bars of the TPR and FPR associated to the support recovery of β^* for five methods with respect to the thresholds when $n = 1000$, $q = 3$, $p = 100$ and a 5% sparsity level. The error bars of the TPR of `lasso_cv` and `lasso_best` coincide.

Based on Figures 11 and 12, we chose a threshold equal to 0.7 for both sparsity levels (5% and 10%) which provides a good trade-off between TPR and FPR for all values of n . We can see from Figure 13 that `ss_min` with this threshold outperforms `lasso_cv` when the sparsity level is equal to 5% and all the values of n considered. In the case where the sparsity level is equal to 10%, `lasso_cv` has a slightly larger TPR for $n = 150$ and $n = 200$. However, the FPR of `ss_min` is much smaller.

Figure 14 displays the boxplots for the estimations of γ^* in Model (2) for $q = 1$, $p = 100$, different values of n (150, 200, 500, 1000) and sparsity levels (5% and 10%) obtained by `ss_min` with a threshold of 0.7 for six iterations. We can see from this figure that this approach provides accurate estimations of γ_1^* from Iteration 2 especially when n is larger than 200.




FIGURE 7. Error bars of the TPR and FPR associated to the support recovery of β^* obtained with the `glarma` package for different thresholds when $n = 1000$, $q = 1$, $p = 100$ and a 5% sparsity level.

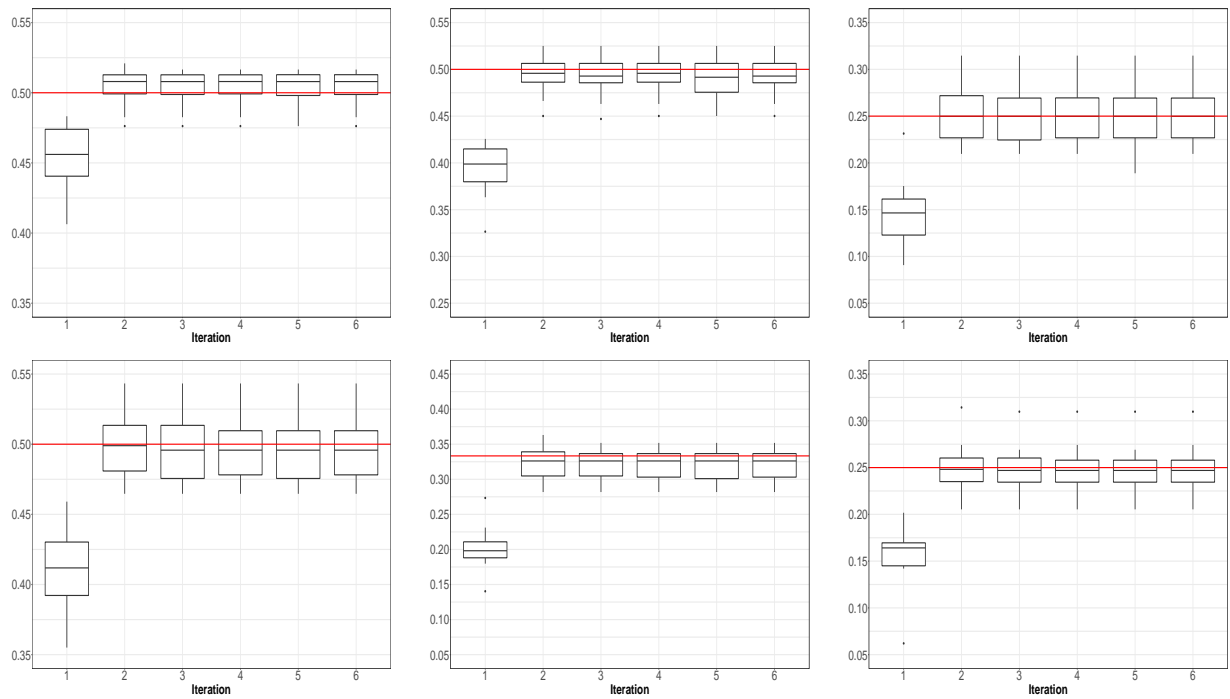


FIGURE 8. Boxplots for the estimations of γ^* in Model (2) with a 5% sparsity level and $q = 1, 2, 3$ obtained by `ss_cv`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_1^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.

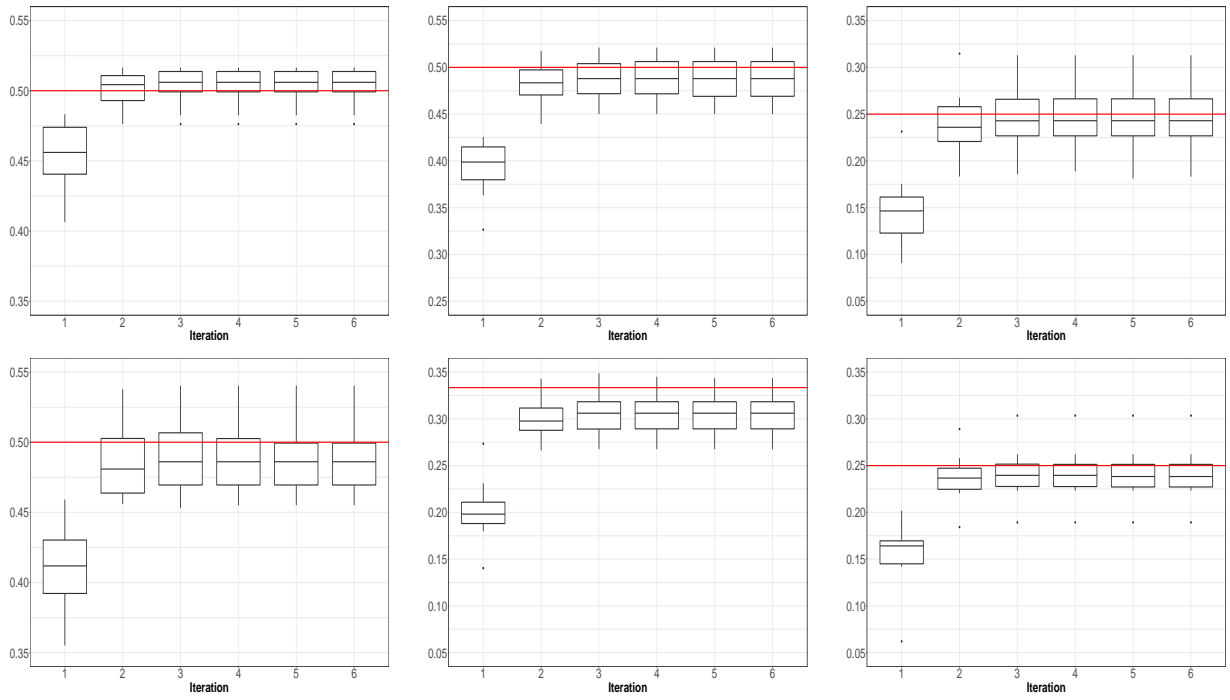


FIGURE 9. Boxplots for the estimations of γ^* in Model (2) with a 5% sparsity level and $q = 1, 2, 3$ obtained by `fast_ss`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_1^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.




FIGURE 10. Boxplots for the estimations of γ^* in Model (2) with a 5% sparsity level and $q = 1, 2, 3$ obtained by `ss_min`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_1^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.




FIGURE 11. Error bars of the TPR and FPR associated to the support recovery of β^* for `ss_min` with respect to the thresholds for different values of n , $q = 1$, $p = 100$ and a 5% sparsity level.




FIGURE 12. Error bars of the TPR and FPR associated to the support recovery of β^* for `ss_min` with respect to the thresholds for different values of n , $q = 1$, $p = 100$ and a 10% sparsity level.




FIGURE 13. Error bars of the TPR and FPR associated to the support recovery of β^* for `ss_min` and `lasso_cv` for different values of n , $q = 1$, $p = 100$ and different sparsity levels.




FIGURE 14. Boxplots for the estimations of γ^* in Model (2) for $q = 1$, $p = 100$, different values of n and sparsity levels (left: 5%, right: 10%) obtained by `ss_min` with a threshold of 0.7 for different iterations (`iter`).

3.2. Numerical performance. Figure 15 displays the means of the computational times for `ss_min` and `fast_ss`. The timings were obtained on a workstation with 8GB of RAM and Dual-Core Intel Core i5 (2.7GHz) CPU. The performance of `ss_cv` are not displayed since they are similar to the one of `ss_min`. We can see from this figure that it takes around 1 minute to process observations Y_1, \dots, Y_n satisfying Model (1) for a given threshold and one iteration, when $n = 1000$ and $p = 100$. Moreover, we can observe that the computational burden of `fast_ss` is slightly smaller than the one of `ss_min`.




FIGURE 15. Means of the computational times in seconds for `ss_min` and `fast_ss` in the case where $p = 100$, and different values of n and q , a given threshold and one iteration.

4. APPLICATION TO THE ANALYSIS OF RNA-SEQ KINETICS DATA

4.1. Biological context and modeling. RNA sequencing (RNA-Seq) allows identifying and counting the numbers of RNA fragments present in a biological sample. By linking these RNA fragments to genes, one can determine the expression level of genes as integer counts. Over the past decades, advances in RNA-Seq analysis have revealed that many eukaryotic genomes were transcribed outside of protein-coding genes. These thousands of new transcripts have been named non-coding RNAs (ncRNAs, [2]) as opposed to coding RNAs, which code for proteins. Among these ncRNAs, long non-coding RNAs (lncRNAs) are a heterogeneous group of RNA molecules greater than 200 nucleotides, transcribed from non-coding genes, that regulate genome expression. This application aims at identifying the lncRNAs, the expression of which affects the expression of coding genes, by using the temporal evolution of the expression of both coding genes and lncRNAs.

Here, we applied the methodology proposed in the paper to 9000 RNA-seq kinetics (or time series) of coding genes, each having a length $n = 15$, to find which lncRNAs among $p = 95$ affect their values. Note that the Poisson modeling is adapted since the expression

of coding genes are integer-valued. More precisely, for each coding gene, the time series is described by its expression (values) at 15 temporal points. In Model (1), (2), and (3) the expression of a given coding gene at time t is denoted by Y_t with $t = 1, 2, \dots, n = 15$ and the expression of the j th lncRNAs at time t is denoted by $x_{j,t}$ with $j = 1, 2, \dots, p = 95$. Our goal is to find which lncRNAs affect the values of (Y_t) which boils down to finding which β_k^* are non null.

4.2. Additional numerical experiments. In order to tune the threshold of `ss_cv` in the specific context of this application ($n = 15$ and $p = 95$), we ran additional numerical experiments. We used the $x_{j,t}$ corresponding to the expression data of the lncRNAs for generating the Y_t 's following the model described in (1), (2), and (3) with $q = 1$, $\gamma_1^* = 0.5$ and 5 non null β_k^* 's. We can see from Figure 16 that `ss_cv` outperforms `lasso_cv` even in this framework where n is much smaller than p and that the best threshold for our approach is 0.4. We shall thus use this value in the following.

4.3. Results. In Figure 17 the results are displayed only for 10 coding genes out of 9000. Our approach selected 46 out of 95 lncRNAs as being relevant for explaining the expression of these 10 coding genes.

In Figure 17, if a coefficient β_k^* is estimated as non null, meaning that the associated lncRNA affects the values of a given coding gene, there is a dot in the plot. If the influence is negative, the dot is blue and if it is positive, the dot is red. The brighter the color of the dot, the larger is the influence.




FIGURE 16. Error bars of the difference between TPR and FPR associated to the support recovery of β^* for `ss_cv` with 4 iterations and `lasso_cv` with respect to the thresholds when $n = 15$, $q = 1$, $p = 95$ and 5 non null coefficients.




FIGURE 17. Estimation of β^* with `ss_cv` for explaining the values of 10 coding genes (Y_t) by some of the lncRNAs ($x_{t,i}$).

Moreover, Figure 18 displays the estimation of γ_1^* obtained for the 10 series associated to the coding genes. Since n is very small and the model has to estimate many parameters, it is unrealistic to expect better results by taking a value of q larger than 1. After 4 iterations for all 10 coding genes, all but one estimates of γ_1^* converge to a value in the interval from -1 to 0.1 .




FIGURE 18. Estimation of γ_1^* with `ss_cv` for explaining the values of 10 coding genes (Y_t) by some of the lncRNAs ($x_{t,i}$).

5. PROOFS

5.1. **Computation of the first and second derivatives of W_t defined in (5).** The computations given below are similar to those provided in [6] but are specific to the parametrization $\boldsymbol{\delta} = (\boldsymbol{\beta}', \boldsymbol{\gamma}')$ considered in this paper.

5.1.1. *Computation of the first derivatives of W_t .* By the definition of W_t given in (5), we get

$$\frac{\partial W_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}) = \frac{\partial \boldsymbol{\beta}' x_t}{\partial \boldsymbol{\delta}} + \frac{\partial Z_t}{\partial \boldsymbol{\delta}}(\boldsymbol{\delta}),$$

where $\boldsymbol{\beta}$, x_t and Z_t are defined in (5). More precisely, for all $k \in \{0, \dots, p\}$, $\ell \in \{1, \dots, q\}$ and $t \in \{1, \dots, n\}$, by (6),

$$\begin{aligned} \frac{\partial W_t}{\partial \beta_k} &= x_{t,k} + \frac{\partial Z_t}{\partial \beta_k} = x_{t,k} + \sum_{j=1}^{q \wedge (t-1)} \gamma_j \frac{\partial E_{t-j}}{\partial \beta_k} \\ &= x_{t,k} - \sum_{j=1}^{q \wedge (t-1)} \gamma_j Y_{t-j} \frac{\partial W_{t-j}}{\partial \beta_k} \exp(-W_{t-j}) = x_{t,k} - \sum_{j=1}^{q \wedge (t-1)} \gamma_j (1 + E_{t-j}) \frac{\partial W_{t-j}}{\partial \beta_k}, \end{aligned} \quad (19)$$

$$\begin{aligned} \frac{\partial W_t}{\partial \gamma_\ell} &= E_{t-\ell} + \sum_{j=1}^{q \wedge (t-1)} \gamma_j \frac{\partial E_{t-j}}{\partial \gamma_\ell} \\ &= E_{t-\ell} - \sum_{j=1}^{q \wedge (t-1)} \gamma_j Y_{t-j} \frac{\partial W_{t-j}}{\partial \gamma_\ell} \exp(-W_{t-j}) = E_{t-\ell} - \sum_{j=1}^{q \wedge (t-1)} \gamma_j (1 + E_{t-j}) \frac{\partial W_{t-j}}{\partial \gamma_\ell}, \end{aligned} \quad (20)$$

where we used that $E_t = 0$, $\forall t \leq 0$.

The first derivatives of W_t are thus obtained from the following recursive expressions. For all $k \in \{0, \dots, p\}$

$$\begin{aligned} \frac{\partial W_1}{\partial \beta_k} &= x_{1,k}, \\ \frac{\partial W_2}{\partial \beta_k} &= x_{2,k} - \gamma_1 (1 + E_1) \frac{\partial W_1}{\partial \beta_k}, \end{aligned}$$

where

$$W_1 = \boldsymbol{\beta}' x_1 \text{ and } E_1 = Y_1 \exp(-W_1) - 1. \quad (21)$$

Moreover,

$$\frac{\partial W_3}{\partial \beta_k} = x_{3,k} - \gamma_1 (1 + E_2) \frac{\partial W_2}{\partial \beta_k} - \gamma_2 (1 + E_1) \frac{\partial W_1}{\partial \beta_k},$$

where

$$W_2 = \boldsymbol{\beta}' x_2 + \gamma_1 E_1, \quad E_2 = Y_2 \exp(-W_2) - 1, \quad (22)$$

and so on. In the same way, for all $\ell \in \{1, \dots, q\}$

$$\begin{aligned}\frac{\partial W_1}{\partial \gamma_\ell} &= 0, \\ \frac{\partial W_2}{\partial \gamma_\ell} &= E_{2-\ell}, \\ \frac{\partial W_3}{\partial \gamma_\ell} &= E_{3-\ell} - \gamma_1(1 + E_2) \frac{\partial W_2}{\partial \gamma_\ell}\end{aligned}$$

and so on, where $E_t = 0, \forall t \leq 0$ and E_1, E_2 are defined in (21) and (22), respectively.

5.1.2. *Computation of the second derivatives of W_t .* Using (19) and (20), we get that for all $j, k \in \{0, \dots, p\}, \ell, m \in \{1, \dots, q\}$ and $t \in \{1, \dots, n\}$,

$$\begin{aligned}\frac{\partial^2 W_t}{\partial \beta_j \partial \beta_k} &= - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i(1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \beta_j \partial \beta_k} - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i \frac{\partial E_{t-i}}{\partial \beta_j} \frac{\partial W_{t-i}}{\partial \beta_k} \\ &= - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i(1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \beta_j \partial \beta_k} + \sum_{i=1}^{q^\wedge(t-1)} \gamma_i(1 + E_{t-i}) \frac{\partial W_{t-i}}{\partial \beta_j} \frac{\partial W_{t-i}}{\partial \beta_k}, \\ \frac{\partial^2 W_t}{\partial \beta_k \partial \gamma_\ell} &= -(1 + E_{t-\ell}) \frac{\partial W_{t-\ell}}{\partial \beta_k} - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i \left\{ \frac{\partial W_{t-i}}{\partial \beta_k} \frac{\partial E_{t-i}}{\partial \gamma_\ell} + (1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \beta_k \partial \gamma_\ell} \right\} \\ &= -(1 + E_{t-\ell}) \frac{\partial W_{t-\ell}}{\partial \beta_k} - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i \left\{ -(1 + E_{t-i}) \frac{\partial W_{t-i}}{\partial \beta_k} \frac{\partial W_{t-i}}{\partial \gamma_\ell} + (1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \beta_k \partial \gamma_\ell} \right\}, \\ \frac{\partial^2 W_t}{\partial \gamma_\ell \partial \gamma_m} &= \frac{\partial E_{t-\ell}}{\partial \gamma_m} - (1 + E_{t-m}) \frac{\partial W_{t-m}}{\partial \gamma_\ell} - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i \left\{ \frac{\partial W_{t-i}}{\partial \gamma_\ell} \frac{\partial E_{t-i}}{\partial \gamma_m} + (1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \gamma_\ell \partial \gamma_m} \right\} \\ &= -(1 + E_{t-\ell}) \frac{\partial W_{t-\ell}}{\partial \gamma_m} - (1 + E_{t-m}) \frac{\partial W_{t-m}}{\partial \gamma_\ell} \\ &\quad - \sum_{i=1}^{q^\wedge(t-1)} \gamma_i \left\{ -(1 + E_{t-i}) \frac{\partial W_{t-i}}{\partial \gamma_\ell} \frac{\partial W_{t-i}}{\partial \gamma_m} + (1 + E_{t-i}) \frac{\partial^2 W_{t-i}}{\partial \gamma_\ell \partial \gamma_m} \right\}.\end{aligned}$$

To compute the second derivatives of W_t , we shall use the following recursive expressions for all $j, k \in \{0, \dots, p\}$

$$\begin{aligned}\frac{\partial^2 W_1}{\partial \beta_j \partial \beta_k} &= 0, \\ \frac{\partial^2 W_2}{\partial \beta_j \partial \beta_k} &= \gamma_1(1 + E_1) x_{1,j} x_{1,k},\end{aligned}$$

where E_1 is defined in (21) and so on. Moreover, for all $k \in \{0, \dots, p\}$ and $\ell \in \{1, \dots, q\}$

$$\begin{aligned}\frac{\partial^2 W_1}{\partial \beta_k \partial \gamma_\ell} &= 0, \\ \frac{\partial^2 W_2}{\partial \beta_k \partial \gamma_\ell} &= -(1 + E_{2-\ell}) \frac{\partial W_{2-\ell}}{\partial \beta_k},\end{aligned}$$

where $E_t = 0$ for all $t \leq 0$ and the first derivatives of W_t are computed in (19). Note also that

$$\begin{aligned}\frac{\partial^2 W_1}{\partial \gamma_\ell \partial \gamma_m} &= 0, \\ \frac{\partial^2 W_2}{\partial \gamma_\ell \partial \gamma_m} &= 0\end{aligned}$$

and so on.

5.2. Computational details for obtaining Criterion (12). By (11),

$$\tilde{L}(\boldsymbol{\beta}) = \tilde{L}(\boldsymbol{\beta}^{(0)}) + \frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U(\boldsymbol{\nu} - \boldsymbol{\nu}^{(0)}) - \frac{1}{2}(\boldsymbol{\nu} - \boldsymbol{\nu}^{(0)})' \Lambda (\boldsymbol{\nu} - \boldsymbol{\nu}^{(0)}),$$

where $\boldsymbol{\nu} - \boldsymbol{\nu}^{(0)} = U'(\boldsymbol{\beta} - \boldsymbol{\beta}^{(0)})$. Hence,

$$\begin{aligned}\tilde{L}(\boldsymbol{\beta}) &= \tilde{L}(\boldsymbol{\beta}^{(0)}) + \sum_{k=0}^p \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U \right)_k (\nu_k - \nu_k^{(0)}) - \frac{1}{2} \sum_{k=0}^p \lambda_k (\nu_k - \nu_k^{(0)})^2 \\ &= \tilde{L}(\boldsymbol{\beta}^{(0)}) - \frac{1}{2} \sum_{k=0}^p \lambda_k \left(\nu_k - \nu_k^{(0)} - \frac{1}{\lambda_k} \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U \right)_k \right)^2 + \sum_{k=0}^p \frac{1}{2\lambda_k} \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U \right)_k^2,\end{aligned}$$

where the λ_k 's are the diagonal terms of Λ .

Since the only term depending on $\boldsymbol{\beta}$ is the second one in the last expression of $\tilde{L}(\boldsymbol{\beta})$, we define $\tilde{L}_Q(\boldsymbol{\beta})$ appearing in Criterion (12) as follows:

$$\begin{aligned}-\tilde{L}_Q(\boldsymbol{\beta}) &= \frac{1}{2} \sum_{k=0}^p \lambda_k \left(\nu_k - \nu_k^{(0)} - \frac{1}{\lambda_k} \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U \right)_k \right)^2 \\ &= \frac{1}{2} \left\| \Lambda^{1/2} \left(\boldsymbol{\nu} - \boldsymbol{\nu}^{(0)} - \Lambda^{-1} \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}})U \right)' \right) \right\|_2^2 \\ &= \frac{1}{2} \left\| \Lambda^{1/2} U'(\boldsymbol{\beta} - \boldsymbol{\beta}^{(0)}) - \Lambda^{-1/2} U' \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}}) \right)' \right\|_2^2 \\ &= \frac{1}{2} \left\| \Lambda^{1/2} U'(\boldsymbol{\beta}^{(0)} - \boldsymbol{\beta}) + \Lambda^{-1/2} U' \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \hat{\boldsymbol{\gamma}}) \right)' \right\|_2^2 \\ &= \frac{1}{2} \|\mathcal{Y} - \mathcal{X}\boldsymbol{\beta}\|_2^2,\end{aligned}$$

where

$$\mathcal{Y} = \Lambda^{1/2}U'\boldsymbol{\beta}^{(0)} + \Lambda^{-1/2}U' \left(\frac{\partial L}{\partial \boldsymbol{\beta}}(\boldsymbol{\beta}^{(0)}, \widehat{\boldsymbol{\gamma}}) \right)', \quad \mathcal{X} = \Lambda^{1/2}U'.$$

5.3. Proofs of Propositions 2.2, 2.3 and 2.4 and of Lemma 5.1. This section contains the proofs of Propositions 2.2, 2.3 and 2.4.

5.3.1. *Proof of Proposition 2.2.* We first establish the following lemma for proving Proposition 2.2.

Lemma 5.1. *(W_t^*) is an aperiodic Markov process satisfying Doeblin's condition.*

Proof of Lemma 5.1. By (15) and (3), we observe that:

$$W_t^* = (\beta_0^* - \gamma_1^*) + \gamma_1^* Y_{t-1} \exp(-W_{t-1}^*). \quad (23)$$

Thus, $\mathcal{F}_{t-2} = \mathcal{F}_{t-1}^{W^*} := \sigma(W_s, s \leq t-1)$. By (1), the distribution of Y_{t-1} conditionally to \mathcal{F}_{t-2} is $\mathcal{P}(\exp(W_{t-1}^*))$. Hence, the distribution of W_t^* conditionally to $\mathcal{F}_{t-1}^{W^*}$ is the same as distribution of W_t^* conditionally to W_{t-1}^* , which means that (W_t^*) has the Markov property.

Let us now prove that (W_t^*) is strongly aperiodic which implies that it is aperiodic.

$\mathbb{P}(W_t^* = \beta_0^* - \gamma_1^* | W_{t-1}^* = \beta_0^* - \gamma_1^*) = \mathbb{P}(Y_{t-1} = 0 | W_{t-1}^* = \beta_0^* - \gamma_1^*) = \exp(-\exp(\beta_0^* - \gamma_1^*)) > 0$, where the first equality comes from (23) and the last equality comes from (1) since $\mathcal{F}_{t-2} = \mathcal{F}_{t-1}^{W^*}$.

To prove that (W_t^*) satisfies Doeblin's condition namely that there exists a probability measure ν with the property that, for some $m \geq 1$, $\varepsilon > 0$ and $\delta > 0$,

$$\nu(B) > \varepsilon \implies \mathbb{P}(W_{t+m-1} \in B, W_{t+m-2} \in B, \dots, W_{t+1} \in B, W_t \in B | W_{t-1} = x) \geq \delta, \quad (24)$$

for all x in the state space X of W_t^* and B in the Borel sets of X , we refer the reader to the proof of Proposition 2 in [5]. □

Proof of Proposition 2.2. For proving Proposition 2.2, we shall use Theorems 1.3.3 and 1.3.5 of [26]. In order to apply these theorems it is enough to prove that (W_t^*) is a strictly stationary and ergodic process since $Y_t W_t(\beta_0^*, \gamma_1) - \exp(W_t(\beta_0^*, \gamma_1))$ is a measurable function of $W_{t+1}^*, W_t^*, \dots, W_2^*$. Note that the latter fact comes from (15) and (3) for Y_t and from (5) with $q = 1$ and $p = 0$ for W_t .

In order to prove that (W_t^*) is a strictly stationary and ergodic process, we have first to prove that (W_t^*) is an aperiodic Markov process satisfying Doeblin's condition, see Lemma 5.1. The statement of Lemma 5.1 corresponds to Assertion (iv) of Theorem 16.0.2 of [21] which is equivalent to Assertion (i) of this theorem, and implies that (W_t^*) is uniformly ergodic.

Hence, by Definition (16.6) of uniform ergodicity given in [21], there exists a unique stationary invariant measure for (W_t^*) , see also the paragraph below Equation (1.3) of [23] for an additional justification. Combining that existence of a unique stationary invariant measure for (W_t^*) with the following arguments shows that (W_t^*) is a strictly stationary process and also an ergodic Markov process.

By Theorem 3.6.3, Corollary 3.6.1 and Definition 3.6.6 of [25], if the process (W_t^*) is started with its unique stationary invariant distribution, (W_t^*) is a strictly stationary process.

By Definition 3.6.8 of [25], the existence of a unique stationary invariant measure for (W_t^*) means that (W_t^*) is an ergodic Markov process, see also the paragraph below (b) [23, p. 717].

Finally, by Theorem 3.6.5 of [25], since (W_t^*) is an ergodic Markov process and a strictly stationary process, (W_t^*) is an ergodic and strictly stationary process in the sense of the assumption of Theorem 1.3.5 of [26]. \square

5.3.2. *Proof of Proposition 2.3.* Note that for all γ_1 ,

$$\begin{aligned} \mathcal{L}(\gamma_1) &= \mathbb{E} [Y_3 W_3(\beta_0^*, \gamma_1) - \exp(W_3(\beta_0^*, \gamma_1))] = \mathbb{E} [\mathbb{E} [Y_3 W_3(\beta_0^*, \gamma_1) - \exp(W_3(\beta_0^*, \gamma_1)) | \mathcal{F}_2]] \\ &= \mathbb{E} [\exp(W_3^*) W_3(\beta_0^*, \gamma_1) - \exp(W_3(\beta_0^*, \gamma_1))] \\ &= \mathbb{E} [\exp(W_3^*) (W_3(\beta_0^*, \gamma_1) - W_3^* + W_3^* - \exp(W_3(\beta_0^*, \gamma_1) - W_3^*))] \\ &\leq \mathbb{E} [\exp(W_3^*) (W_3^* - 1)] = \mathcal{L}(\gamma_1^*), \end{aligned}$$

where the inequality comes from the following inequality $x - \exp(x) \leq -1$, for all $x \in \mathbb{R}$. This inequality is an equality only when $x = 0$ which means that $\gamma_1 = \gamma_1^*$.

5.3.3. *Proof of Proposition 2.4.* The proof of this proposition comes from Proposition 2.2 and the stochastic equicontinuity of $n^{-1}L(\beta_0^*, \gamma_1)$. Thus, it is enough to prove that there exists a positive δ such that

$$\sup_{|\gamma_1 - \gamma_2| \leq \delta} \left| \frac{L(\beta_0^*, \gamma_1)}{n} - \frac{L(\beta_0^*, \gamma_2)}{n} \right| \xrightarrow{p} 0, \text{ as } n \text{ tends to infinity.}$$

Observe that, by (16),

$$\begin{aligned} \left| \frac{L(\beta_0^*, \gamma_1)}{n} - \frac{L(\beta_0^*, \gamma_2)}{n} \right| &\leq \frac{1}{n} \sum_{t=1}^n Y_t |W_t(\beta_0^*, \gamma_1) - W_t(\beta_0^*, \gamma_2)| \\ &\quad + \frac{1}{n} \sum_{t=1}^n |\exp(W_t(\beta_0^*, \gamma_1)) - \exp(W_t(\beta_0^*, \gamma_2))|. \end{aligned}$$

Let us first focus on bounding the following expression for $t \geq 2$ (since $W_1(\beta_0^*, \gamma) = \beta_0^*$, for all γ). By (17)

$$\begin{aligned}
|W_t(\beta_0^*, \gamma_1) - W_t(\beta_0^*, \gamma_2)| &= |Z_t(\gamma_1) - Z_t(\gamma_2)| = |\gamma_1 E_{t-1}(\gamma_1) - \gamma_2 E_{t-1}(\gamma_2)| \\
&= |\gamma_1 [Y_{t-1} \exp(-W_{t-1}(\beta_0^*, \gamma_1)) - 1] - \gamma_2 [Y_{t-1} \exp(-W_{t-1}(\beta_0^*, \gamma_2)) - 1]| \\
&= |Y_{t-1} e^{-\beta_0^*} [\gamma_1 \exp(-Z_{t-1}(\gamma_1)) - \gamma_2 \exp(-Z_{t-1}(\gamma_2))] + \gamma_2 - \gamma_1| \\
&\leq Y_{t-1} e^{-\beta_0^*} [|\gamma_1 - \gamma_2| \exp(-Z_{t-1}(\gamma_1)) + |\gamma_2| |\exp(-Z_{t-1}(\gamma_1)) - \exp(-Z_{t-1}(\gamma_2))|] + |\gamma_2 - \gamma_1| \\
&\leq Y_{t-1} e^{-\beta_0^*} |\gamma_1 - \gamma_2| \exp(-Z_{t-1}(\gamma_1)) \\
&+ Y_{t-1} e^{-\beta_0^*} |\gamma_2| \exp(-Z_{t-1}(\gamma_1)) |Z_{t-1}(\gamma_1) - Z_{t-1}(\gamma_2)| \exp(|Z_{t-1}(\gamma_1) - Z_{t-1}(\gamma_2)|) \\
&+ |\gamma_2 - \gamma_1|,
\end{aligned}$$

where we used in the last inequality that for all x and y in \mathbb{R} ,

$$|e^x - e^y| = e^x |1 - e^{y-x}| \leq e^x |y - x| e^{|y-x|}. \quad (25)$$

Observing that

$$\exp(-Z_t(\gamma_1)) = \exp(-\gamma_1 [Y_{t-1} e^{-\beta_0^*} \exp(-Z_{t-1}(\gamma_1)) - 1]), \quad (26)$$

and $|Z_2(\gamma_1) - Z_2(\gamma_2)| \leq \delta [Y_1 e^{-\beta_0^*} + 1]$ we get, for γ_1 and γ_2 such that $|\gamma_1 - \gamma_2| \leq \delta$, that

$$|W_t(\beta_0^*, \gamma_1) - W_t(\beta_0^*, \gamma_2)| \leq \delta F(Y_{t-1}, Y_{t-2}, \dots, Y_1), \quad (27)$$

where F is a measurable function. By (25),

$$\begin{aligned}
&|\exp(W_t(\beta_0^*, \gamma_1)) - \exp(W_t(\beta_0^*, \gamma_2))| \\
&\leq \exp(W_t(\beta_0^*, \gamma_1)) |W_t(\beta_0^*, \gamma_1) - W_t(\beta_0^*, \gamma_2)| \exp(|W_t(\beta_0^*, \gamma_1) - W_t(\beta_0^*, \gamma_2)|) \\
&\leq \delta G(Y_{t-1}, Y_{t-2}, \dots, Y_1)
\end{aligned}$$

where the last inequality comes from (27), (26) and (17) and where G is a measurable function. Thus, we get that

$$\left| \frac{L(\beta_0^*, \gamma_1)}{n} - \frac{L(\beta_0^*, \gamma_2)}{n} \right| \leq \frac{\delta}{n} \sum_{t=1}^n H(Y_t, Y_{t-1}, \dots, Y_1),$$

which gives the result by using similar arguments as those given in the proof of Proposition 2.2 namely that (Y_t) is strictly stationary and ergodic. By Theorem 1.3.3 of [26], $H(Y_t, Y_{t-1}, \dots, Y_1)$ is strictly stationary and ergodic since (Y_t) has these properties. Thus, $\mathbb{E}[|H(Y_t, Y_{t-1}, \dots, Y_1)|] < \infty$, which concludes the proof by Theorem 1.3.5 of [26].

6. CONCLUSION

In this paper we propose a novel and efficient two-stage variable selection approach for sparse GLARMA models, which are pervasive for modeling discrete-valued time series. It consists in first estimating the ARMA coefficients and then in estimating the regression coefficients by using a regularized approach. In the course of this study we have shown that our method has two main features which make it very attractive. Firstly, our approach showed very good statistical performance since it is able to outperform the other methods

in recovering the non null regression coefficients. Secondly, its low computational load makes its use possible on relatively large data.

APPENDIX

This appendix contains additional results for the support recovery of β^* and for the estimation of γ^* discussed in Section 3.1.2. In addition to the error bar plots for 10% sparsity, Tables 1 and 2 summarize the information for all the experiments that we conducted.




FIGURE 19. Error bars of the TPR and FPR associated to the support recovery of β^* for five methods with respect to the thresholds when $n = 1000$, $q = 1$, $p = 100$ and a 10% sparsity level.




FIGURE 20. Error bars of the TPR and FPR associated to the support recovery of β^* for five methods with respect to the thresholds when $n = 1000$, $q = 2$, $p = 100$ and a 10% sparsity level.

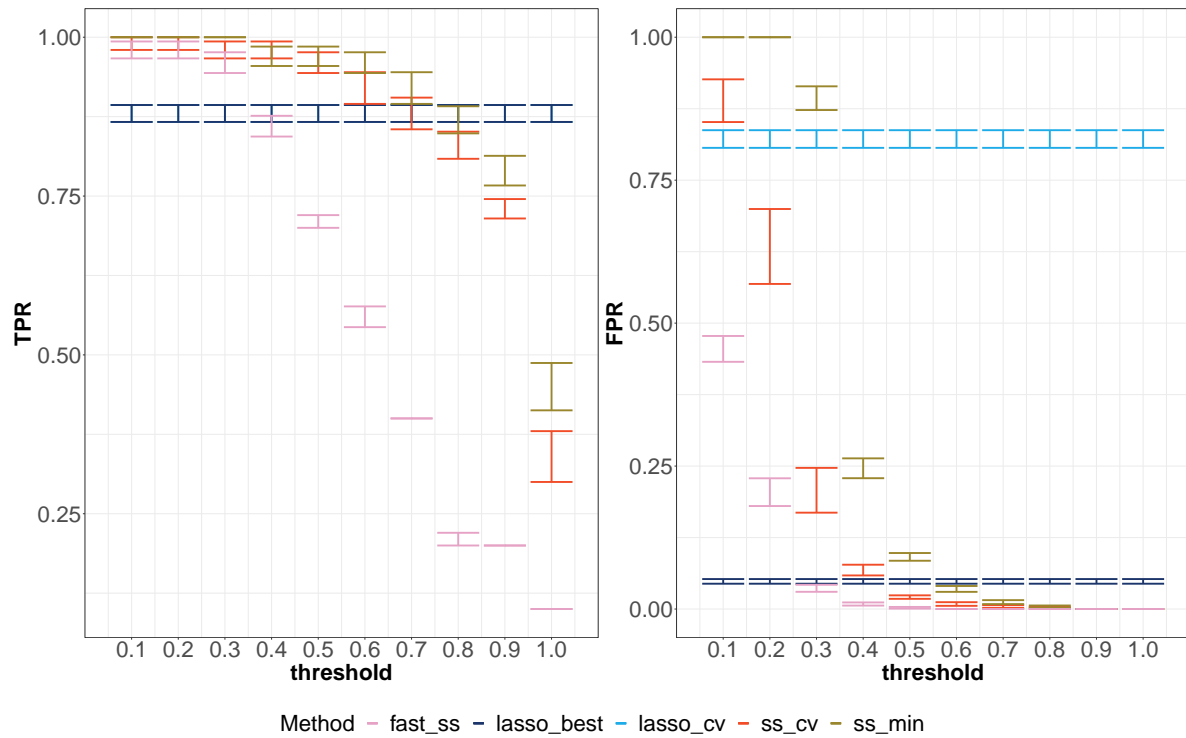


FIGURE 21. Error bars of the TPR and FPR giving the corresponding final sets of selected variables for five methods with respect to the thresholds when $n = 1000$, $q = 3$, $p = 100$ and a 10% sparsity level.




FIGURE 22. Boxplots for the estimations of γ^* in Model (2) with a 10% sparsity level and $q = 1, 2, 3$ obtained by `ss.cv`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_2^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.




FIGURE 23. Boxplots for the estimations of γ^* in Model (2) with a 10% sparsity level and $q = 1, 2, 3$ obtained by `fast_ss`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_1^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.

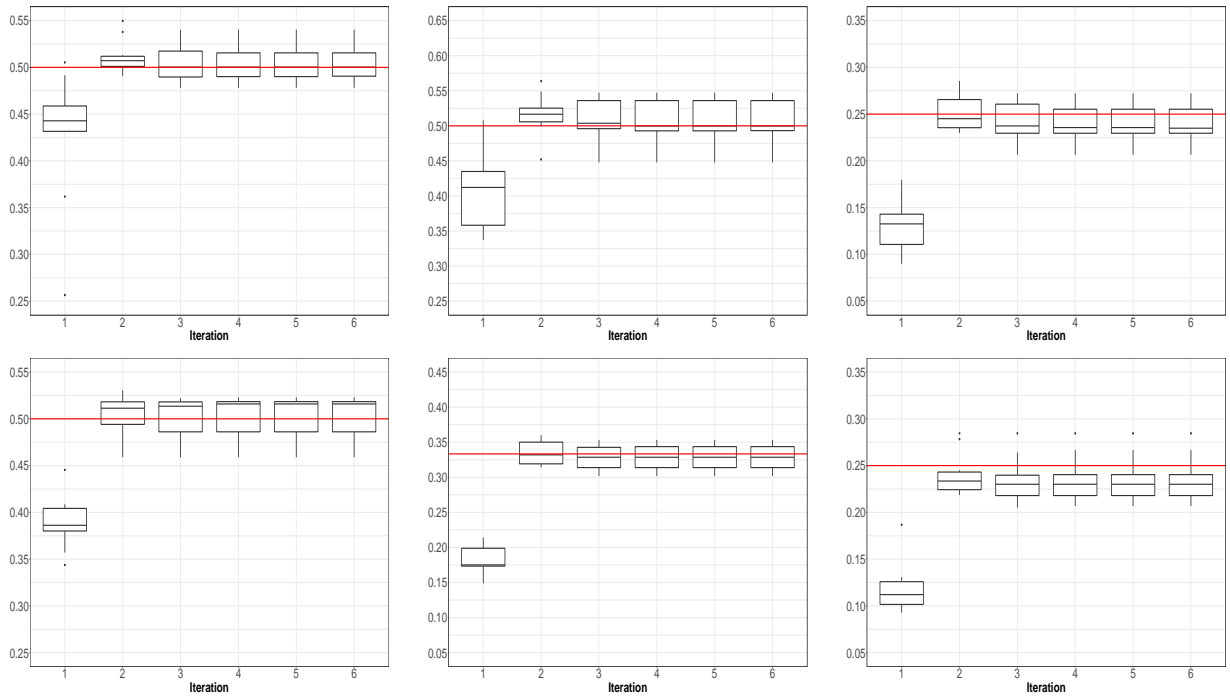


FIGURE 24. Boxplots for the estimations of γ^* in Model (2) with a 10% sparsity level and $q = 1, 2, 3$ obtained by `ss_min`. Top: $q = 1$ and $\gamma_1^* = 0.5$ (left), $q = 2$ and $\gamma_2^* = 0.5$ (middle), $q = 2$ and $\gamma_2^* = 0.25$ (right). Bottom: $q = 3$ and $\gamma_1^* = 0.5$ (left), $q = 3$ and $\gamma_2^* = 1/3$ (middle), $q = 3$ and $\gamma_3^* = 0.25$ (right). The horizontal lines correspond to the values of the γ_i^* 's.

n	q	sparsity (in %)	<code>ss_cv</code>	<code>ss_min</code>	<code>fast_ss</code>	<code>lasso_cv</code>	<code>lasso_best</code>
1000	1	5	1(0)	1(0)	1(0)	0.8(0)	0.8(0)
1000	2	5	0.94(0.09)	0.96(0.08)	0.98(0.06)	0.8(0)	0.8(0)
1000	3	5	0.94(0.1)	1(0)	1(0)	0.8(0)	0.8(0)
500	1	5	0.96(0.08)	1(0)	1(0)	0.8(0)	0.8(0)
500	2	5	0.9(0.14)	0.96(0.08)	1(0)	0.8(0)	0.8(0)
500	3	5	0.92(0.1)	0.98(0.06)	1(0)	0.8(0)	0.8(0)
200	1	5	0.88(0.17)	0.94(0.1)	0.98(0.06)	0.8(0)	0.8(0)
200	2	5	0.98(0.06)	1(0)	1(0)	0.8(0)	0.8(0)
200	3	5	0.92(0.14)	0.94(0.13)	1(0)	0.8(0)	0.8(0)
150	1	5	0.96(0.13)	0.94(0.13)	1(0)	0.8(0)	0.8(0)
150	2	5	0.84(0.16)	0.94(0.09)	1(0)	0.76(0.08)	0.78(0.06)
150	3	5	0.82(0.15)	0.92(0.1)	0.96(0.08)	0.76(0.08)	0.76(0.08)
1000	1	10	0.92(0.08)	0.95(0.05)	0.94(0.05)	0.89(0.03)	0.88(0.04)
1000	2	10	0.91(0.09)	0.96(0.05)	0.93(0.05)	0.9(0)	0.9(0)
1000	3	10	0.88(0.08)	0.92(0.08)	0.96(0.05)	0.88(0.04)	0.88(0.04)
500	1	10	0.88(0.09)	0.93(0.05)	0.97(0.04)	0.9(0)	0.89(0.03)
200	1	10	0.8(0.09)	0.83(0.07)	0.94(0.07)	0.87(0.07)	0.86(0.07)
150	1	10	0.88(0.08)	0.88(0.06)	0.97(0.05)	0.87(0.05)	0.86(0.05)

TABLE 1. Mean of TPR and corresponding standard deviation given in parenthesis associated to the support recovery of β^* for five methods, for different values of n , q , sparsity levels and $p = 100$. For 5% sparsity the thresholds of `ss_cv` and `ss_min` are 0.8 and the threshold of `fast_ss` is 0.4. For 10% sparsity the thresholds of `ss_cv` and `ss_min` are 0.7 and the threshold of `fast_ss` is 0.3.

n	q	sparsity (in %)	ss_cv	ss_min	fast_ss	lasso_cv	lasso_best
1000	1	5	0.001(0.003)	0.005(0.005)	0.003(0.005)	0.42(0.16)	0(0)
1000	2	5	0.002(0.004)	0.01(0.01)	0.013(0.01)	0.76(0.09)	0.007(0.02)
1000	3	5	0.003(0.005)	0.01(0.008)	0.04(0.01)	0.83(0.09)	0.01(0.02)
500	1	5	0.002(0.004)	0.01(0.01)	0.04(0.02)	0.55(0.13)	0.005(0.01)
500	2	5	0(0)	0.009(0.007)	0.07(0.02)	0.76(0.12)	0.02(0.04)
500	3	5	0.002(0.004)	0.01(0.01)	0.12(0.01)	0.77(0.12)	0.02(0.03)
200	1	5	0.002(0.004)	0.014(0.014)	0.19(0.04)	0.46(0.22)	0.05(0.06)
200	2	5	0.003(0.005)	0.02(0.01)	0.24(0.08)	0.39(0.12)	0.06(0.04)
200	3	5	0.008(0.009)	0.02(0.01)	0.23(0.06)	0.39(0.14)	0.05(0.04)
150	1	5	0.001(0.003)	0.02(0.018)	0.22(0.05)	0.2(0.1)	0.05(0.06)
150	2	5	0.006(0.01)	0.04(0.04)	0.28(0.08)	0.2(0.09)	0.08(0.05)
150	3	5	0.003(0.005)	0.03(0.007)	0.3(0.1)	0.2(0.08)	0.05(0.04)
1000	1	10	0(0)	0.008(0.01)	0.008(0.01)	0.51(0.14)	0.04(0.02)
1000	2	10	0.004(0.005)	0.008(0.008)	0.03(0.02)	0.78(0.13)	0.06(0.01)
1000	3	10	0.004(0.007)	0.012(0.01)	0.04(0.02)	0.82(0.05)	0.05(0.01)
500	1	10	0.004(0.005)	0.019(0.016)	0.06(0.01)	0.67(0.19)	0.05(0.02)
200	1	10	0.008(0.008)	0.09(0.08)	0.34(0.23)	0.5(0.2)	0.07(0.04)
150	1	10	0.01(0.01)	0.11(0.07)	0.38(0.16)	0.33(0.16)	0.07(0.02)

TABLE 2. Mean of FPR and corresponding standard deviation given in parenthesis associated to the support recovery of β^* for five methods, for different values of n , q , sparsity levels and $p = 100$. For 5% sparsity the thresholds of `ss_cv` and `ss_min` are 0.8 and the threshold of `fast_ss` is 0.4. For 10% sparsity the thresholds of `ss_cv` and `ss_min` are 0.7 and the threshold of `fast_ss` is 0.3.

REFERENCES

- [1] Al-Osh, M. and A. A. Alzaid (1988). Integer-valued moving average (INMA) process. Statistical Papers **29**(1), 281–300.
- [2] Ariel, F., N. Romero-Barrios, T. Jégu, M. Benhamed, and M. Crespi (2015). Battles and hijacks: noncoding transcription in plants. Trends Plant Sci **20**(6), 362–371.
- [3] Brännäs, K. and A. M. M. S. Quoreshi (2010). Integer-valued moving average modelling of the number of transactions in stocks. Applied Financial Economics **20**(18), 1429–1440.
- [4] Cox, D. R., G. Gudmundsson, G. Lindgren, L. Bondesson, E. Harsaae, P. Laake, K. Juselius, and S. L. Lauritzen (1981). Statistical analysis of time series: Some recent developments [with discussion and reply]. Scandinavian Journal of Statistics **8**(2), 93–115.
- [5] Davis, R. A., W. T. M. Dunsmuir, and S. B. Strett (2003). Observation-driven models for Poisson counts. Biometrika **90**(4), 777–790.
- [6] Davis, R. A., W. T. M. Dunsmuir, and S. B. Strett (2005). Maximum likelihood estimation for an observation driven model for Poisson counts. Methodology and Computing in Applied Probability **7**(2), 149–159.
- [7] Davis, R. A., W. T. M. Dunsmuir, and Y. Wang (1999). Modeling time series of count data. Statistics Textbooks and Monographs **158**, 63–114.
- [8] Davis, R. A., S. H. Holan, R. Lund, and N. Ravishanker (Eds.) (2016). Handbook of discrete-valued time series. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL.
- [9] Dunsmuir, W. T. M. (2015). Generalized Linear Autoregressive Moving Average Models, Chapter 3, pp. 51–76. CRC Press.
- [10] Dunsmuir, W. T. M. and D. Scott (2015). The glarma package for observation-driven time series regression of counts. Journal of Statistical Software, Articles **67**(7), 1–36.
- [11] Enciso-Mora, V., P. Neal, and T. Subba Rao (2009). Efficient order selection algorithms for integer-valued ARMA processes. Journal of Time Series Analysis **30**(1), 1–18.
- [12] Fokianos, K., A. Rahbek, and D. Tjøstheim (2009). Poisson autoregression. Journal of the American Statistical Association **104**(488), 1430–1439.
- [13] Fokianos, K. and D. Tjøstheim (2012). Nonlinear poisson autoregression. Annals of the Institute of Statistical Mathematics **64**(6), 1205–1225.
- [14] Fokianos, K. and D. Tjøstheim (2011). Log-linear poisson autoregression. Journal of Multivariate Analysis **102**(3), 563 – 578.
- [15] Friedman, J., T. Hastie, and R. Tibshirani (2010). Regularization paths for generalized linear models via coordinate descent. Journal of Statistical Software **33**(1), 1–22.
- [16] Hastie, T., R. Tibshirani, and J. Friedman (2009). The elements of statistical learning: data mining, inference, and prediction. Springer Science & Business Media.
- [17] Hastie, T., R. Tibshirani, and M. Wainwright (2015). Statistical learning with sparsity: the lasso and generalizations. Chapman and Hall/CRC.
- [18] Jung, R. C. and R. Liesenfeld (2001). Estimating time series models for count data using efficient importance sampling. ASTA Advances in Statistical Analysis **4**(85), 387–407.
- [19] McKenzie, E. (1985). Some simple models for discrete variate time series. Journal of the American Water Resources Association **21**(4), 645–650.
- [20] Meinshausen, N. and P. Bühlmann (2010). Stability selection. Journal of the Royal Statistical Society: Series B (Statistical Methodology) **72**(4), 417–473.
- [21] Meyn, S. and R. Tweedie (1993). Markov Chains and Stochastic Stability. Springer-Verlag, London.
- [22] Neal, P. and T. Subba Rao (2007). MCMC for integer-valued ARMA processes. Journal of Time Series Analysis **28**(1), 92–110.
- [23] Sandrić, N. (2017). A note on the Birkhoff ergodic theorem. Results in Mathematics **72**(1), 715–730.

- [24] Souza, J. B. d., V. A. Reisen, J. M. Santos, and G. C. Franco (2014). Principal components and generalized linear modeling in the correlation between hospital admissions and air pollution. *Revista de Saúde Pública* **48**, 451 – 458.
- [25] Stout, W. (1974). *Almost sure convergence*. Probability and mathematical statistics. Academic Press.
- [26] Taniguchi, M. and Y. Kakizawa (2012). *Asymptotic theory of statistical inference for time series*. Springer Science & Business Media.
- [27] Thorne, T. (2018). Approximate inference of gene regulatory network models from RNA-Seq time series data. *BMC Bioinformatics* **19**(1), 127.
- [28] Weiss, C. (2018). *An Introduction to Discrete-Valued Time Series*. John Wiley & Sons Ltd.
- [29] Wu, H., L. Yang, and L.-L. Chen (2017). The diversity of long noncoding rnas and their generation. *Trends in genetics* **33**(8), 540–552.
- [30] Zeger, S. L. and B. Qaqish (1988). Markov regression models for time series: A quasi-likelihood approach. *Biometrics* **44**(4), 1019–1031.

UNIVERSITÉ PARIS-SACLAY, AGROPARISTECH, INRAE, UMR MIA PARIS-SACLAY, 91120 PALAISEAU, FRANCE

INSTITUTE OF PLANT SCIENCES PARIS-SACLAY, CNRS, INRAE, UNIVERSITÉ EVRY, UNIVERSITÉ PARIS-SACLAY, UNIVERSITÉ DE PARIS, 91405 ORSAY, FRANCE.