



Stochastic proximal-gradient algorithms for penalized mixed models

Gersende Fort¹ · Edouard Ollier^{2,3} · Adeline Samson⁴

Received: 28 April 2017 / Accepted: 2 February 2018 / Published online: 12 February 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

Motivated by penalized likelihood maximization in complex models, we study optimization problems where neither the function to optimize nor its gradient has an explicit expression, but its gradient can be approximated by a Monte Carlo technique. We propose a new algorithm based on a stochastic approximation of the proximal-gradient (PG) algorithm. This new algorithm, named stochastic approximation PG (SAPG) is the combination of a stochastic gradient descent step which—roughly speaking—computes a smoothed approximation of the gradient along the iterations, and a proximal step. The choice of the step size and of the Monte Carlo batch size for the stochastic gradient descent step in SAPG is discussed. Our convergence results cover the cases of biased and unbiased Monte Carlo approximations. While the convergence analysis of some classical Monte Carlo approximation of the gradient is already addressed in the literature (see Atchadé et al. in *J Mach Learn Res* 18(10):1–33, 2017), the convergence analysis of SAPG is new. Practical implementation is discussed, and guidelines to tune the algorithm are given. The two algorithms are compared on a linear mixed effect model as a toy example. A more challenging application is proposed on nonlinear mixed effect models in high dimension with a pharmacokinetic data set including genomic covariates. To our best knowledge, our work provides the first convergence result of a numerical method designed to solve penalized maximum likelihood in a nonlinear mixed effect model.

Keywords Proximal-gradient algorithm · Stochastic gradient · Stochastic EM algorithm · Stochastic approximation · Nonlinear mixed effect models

1 Introduction

Many problems in computational statistics reduce to the maximization of a criterion

$$\operatorname{argmax}_{\theta \in \mathbb{R}^d} F(\theta), \quad \text{where } F := \ell - g, \quad (1)$$

✉ Edouard Ollier
edouard.ollier@univ-st-etienne.fr

Gersende Fort
Gersende.Fort@math.univ-toulouse.fr

Adeline Samson
adeline.leclercq-samson@univ-grenoble-alpes.fr

¹ IMT UMR5219, CNRS, Université de Toulouse, 31062 Toulouse Cedex 9, France

² INSERM, U1059, Dysfonction Vasculaire et Hémostase, Saint Etienne, France

³ U.M.P.A., Ecole Normale Supérieure de Lyon, CNRS UMR 5669, INRIA, Project-Team NUMED, 46 Allée d'Italie, 69364 Lyon Cedex 07, France

⁴ Laboratoire Jean Kuntzmann, UMR CNRS 5224, Université Grenoble-Alpes, 38000 Grenoble, France

and the functions ℓ, g satisfy

H1 The function $g : \mathbb{R}^d \rightarrow [0, +\infty]$ is convex, not identically $+\infty$, and lower semi-continuous.

H2 The function $\ell : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{-\infty\}$ is continuously differentiable on $\Theta := \{\theta \in \mathbb{R}^d : g(\theta) + |\ell(\theta)| < \infty\}$, and its gradient is of the form

$$\nabla \ell(\theta) = \nabla \phi(\theta) + \Psi(\theta) \bar{S}(\theta), \quad (2)$$

with $\bar{S}(\theta) := \int_{\mathcal{Z}} S(z) \pi_{\theta}(z) \nu(dz);$

∇ denotes the gradient operator, and $\pi_{\theta} d\nu$ is a probability distribution on a measurable subset $(\mathcal{Z}, \mathcal{Z})$ of \mathbb{R}^p . The measurable functions $\nabla \phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\Psi : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times q}$ are known but the expectation \bar{S} of the function $S : \mathcal{Z} \rightarrow \mathbb{R}^q$ with respect to $\pi_{\theta} d\nu$ may be intractable. Furthermore, there exists a finite nonnegative constant L such that for all $\theta, \theta' \in \Theta$,

$$\|\nabla \ell(\theta) - \nabla \ell(\theta')\| \leq L \|\theta - \theta'\|; \quad (3)$$

$\|\cdot\|$ is the Euclidean norm.

Examples of functions ℓ satisfying Eq. (2) are given below. We are interested in numerical methods for solving Eq. (1), robust to the case when neither ℓ nor its gradient has an explicit expression.

Such an optimization problem occurs for example when computing a penalized maximum-likelihood estimator in some parametric model indexed by $\theta \in \mathbb{R}^d$: ℓ denotes the log likelihood of the observations Y (the dependence upon Y is omitted) and g is the penalty term.

The optimization problem Eq. (1) covers the computation of the maximum when the parameter θ is restricted to a closed convex subset Θ of \mathbb{R}^d ; in that case, g is the characteristic function of Θ , i.e., $g(\theta) = 0$ for any $\theta \in \Theta$ and $g(\theta) = +\infty$ otherwise. It also covers the case when g is the ridge, the lasso or the elastic net penalty; and more generally, the case when g is the sum of lower semi-continuous nonnegative convex functions.

A first example of such a function ℓ is given by the log likelihood in a latent variable model with complete likelihood from the q -parameter exponential family (see, e.g., Bickel and Doksum 2015; Bartholomew et al. 2011 and the references therein). In that case, ℓ is of the form

$$\theta \mapsto \ell(\theta) := \log \int_Z \exp(\phi(\theta) + \langle S(z), \psi(\theta) \rangle) \nu(dz), \tag{4}$$

where $\langle a, b \rangle$ denotes the scalar of two vectors $a, b \in \mathbb{R}^l$, $\phi: \mathbb{R}^d \rightarrow \mathbb{R}$, $\psi: \mathbb{R}^d \rightarrow \mathbb{R}^q$ and $S: Z \rightarrow \mathbb{R}^q$ are measurable functions, and ν is a σ -finite positive measure on (Z, \mathcal{Z}) . The quantity $\theta \mapsto \phi(\theta) + \langle S(Z), \psi(\theta) \rangle$ is known as the complete log likelihood, and Z is the latent data vector. Under regularity conditions, we have

$$\begin{aligned} \nabla \ell(\theta) &= \nabla \phi(\theta) + J \psi(\theta) \int_Z S(z) \pi_\theta(z) \nu(dz), \\ \text{with } \pi_\theta(z) &:= \frac{\exp(\langle S(z), \psi(\theta) \rangle)}{\int_Z \exp(\langle S(u), \psi(\theta) \rangle) \nu(du)}, \end{aligned} \tag{5}$$

where $J \psi(\theta)$ denotes the transpose of the jacobian matrix of the function ψ at θ .

A second example is given by the log likelihood of N independent observations (Y_1, \dots, Y_N) from a log-linear model for Markov random fields. In this model, ℓ is given by

$$\begin{aligned} \theta \mapsto \ell(\theta) &:= \sum_{k=1}^N \langle S(Y_k), \theta \rangle \\ &\quad - N \log \int_Z \exp(\langle S(z), \theta \rangle) \nu(dz). \end{aligned} \tag{6}$$

The function $\theta \mapsto \int_Z \exp(\langle S(z), \theta \rangle) \nu(dz)$ is known as the partition function. Under regularity conditions, we have

$$\begin{aligned} \nabla \ell(\theta) &= \sum_{k=1}^N S(Y_k) - N \int_Z S(z) \pi_\theta(z) \nu(dz), \\ \text{with } \pi_\theta(z) &:= \frac{\exp(\langle S(z), \theta \rangle)}{\int \exp(\langle S(u), \theta \rangle) \nu(du)}. \end{aligned} \tag{7}$$

In these two examples, the integrals in Eqs. (4)–(7) are intractable except for toy examples: neither the function ℓ nor its gradient is available. Nevertheless, all the integrals in Eqs. (4)–(7) can be approximated by a Monte Carlo sum (see, e.g., Robert and Casella 2004). In the first example, this Monte Carlo approximation consists in imputing the missing variables z ; it is known that such an imputation is far more efficient when the Monte Carlo samples are drawn under $\pi_\theta \nu$, i.e., the a posteriori distribution of the missing variables given the observations (see Eq. 5) than when they are drawn under the a priori distribution. This remark is the essence of the expectation maximization (EM) algorithm (introduced in Dempster et al. 1977), a popular iterative procedure for maximizing the log likelihood ℓ in latent variable models.

In this paper, we are interested in first-order optimization methods to solve Eq. (1), that is methods based on the gradient. In Sect. 2.1, we describe two stochastic first-order descent methods, which are stochastic perturbations of the proximal-gradient (PG) algorithm (introduced in Combettes and Pesquet 2011; see also Beck and Teboulle 2009; Parikh and Boyd 2013 for literature reviews on proximal-gradient algorithms). The two algorithms are the Monte Carlo proximal-gradient algorithm (MCPG) and the stochastic approximation proximal-gradient algorithm (SAPG), which differ in the approximation of the gradient $\nabla \ell$ and more precisely, of the intractable integral $\bar{S}(\theta)$ (see Eq. 2). In MCPG, at each iteration n of the algorithm, this expectation evaluated at the current point θ_n is approximated by a Monte Carlo sum computed from samples $\{Z_{1,n}, \dots, Z_{m_{n+1},n}\}$ approximating $\pi_{\theta_n} \nu$. In SAPG, the approximation is computed as a Monte Carlo sum based on all the points drawn during all the previous iterations of the algorithm $\{Z_{i,j}, i \leq m_{j+1}, j \leq n\}$.

When ℓ is the log likelihood of a latent variable model, we prove in Sect. 2.2 that our algorithms are generalized EM algorithms (see, e.g., McLachlan and Krishnan 2008; Ng et al. 2012) combined with a stochastic E-step: in MCPG and SAPG, the stochastic E-step mimics, respectively, the E-step of the Monte Carlo EM (Wei and Tanner 1990; Levine and Fan 2004) and the E-step of the stochastic approximation EM (see, e.g., Delyon et al. 1999).

Section 3 is devoted to the convergence analysis of MCPG and SAPG. These algorithms can be seen as perturbed proximal-gradient algorithms when the perturbation comes from replacing the exact quantity $\bar{S}(\theta_n)$ by a Monte Carlo approximation S_{n+1} at each iteration of the algorithm. Our convergence analysis covers the case when the points

$\{Z_{1,n}, \dots, Z_{m_{n+1},n}\}$ are sampled from a Markov chain Monte Carlo sampler (MCMC) with target distribution $\pi_{\theta_n} d\nu$ —and therefore, it also covers the case of i.i.d. draws. This implies that the estimator S_{n+1} of $\bar{S}(\theta_n)$ may be biased. There exist many contributions in the literature on the convergence of perturbed proximal-gradient algorithms when ℓ is concave, but except in the works by Atchadé et al. (2017) and Combettes and Pesquet (2015), most of them assume that the error $S_{n+1} - \bar{S}(\theta_n)$ is unbiased and gets small when $n \rightarrow \infty$ (see, e.g., Rosasco et al. 2014; Combettes and Pesquet 2016; Rosasco et al. 2016; Lin et al. 2015). In this paper, we provide sufficient conditions for the almost-sure convergence of MCPG and SAPG under the assumption that ℓ is concave and with no assumptions on the bias of $S_{n+1} - \bar{S}(\theta_n)$. The convergence analysis of MCPG is a special case of Atchadé et al. (2017, Section 4); to our best knowledge, the convergence of SAPG is a new result.

Practical implementation is discussed in Sect. 4. Some guidelines are given in Sect. 4.2 to choose the sequences involved in the stochastic approximation procedures. Then, MCPG and SAPG are compared through a toy example in Sect. 4.3. A more challenging application to penalized inference in a mixed effect model is detailed in Sect. 5. Mixed models are applied to analyze repeated data in a population of subjects. The N independent vectors of observations $(Y_k, k = 1, \dots, N)$ of the N subjects are modeled by

$$Y_k = f(t_k, Z^{(k)}) + \varepsilon_k, \tag{8}$$

with individual latent variable $Z^{(k)}$ independent of the measurement error vector ε_k and f the regression function that depends on the vector of observation times t_k . Mixed models thus enter the class of models given by Eq. (4) with latent variables $Z = (Z^{(1)}, \dots, Z^{(N)})$. When a covariate model is introduced, the number of covariates can be large, but with only a few of them being influential. This is a sparse estimation problem and the selection problem can be treated through the optimization of a penalized version of the log likelihood Eq. (4). In nonlinear mixed models, the optimization problem is not explicit and stochastic penalized versions of EM (Bertrand and Balding 2013; Ollier et al. 2016; Chen et al. 2017) have been proposed. To our best knowledge, stochastic proximal-gradient algorithms have not been proposed for mixed models.

2 Stochastic proximal-gradient based algorithms

In this section, we describe first-order based algorithms for solving Eq. (1) under the assumptions H1 and H2, when the expectation $\bar{S}(\theta)$ in Eq. (2) is intractable.

2.1 The MCPG and SAPG algorithms

Both MCPG and SAPG are iterative algorithms, and each update relies on the combination of a gradient step and a proximal operator. The proximal map (Moreau 1962, see also Bauschke and Combettes 2011; Parikh and Boyd 2013) associated with a convex function g is defined for any $\gamma > 0$ and $\theta \in \mathbb{R}^d$ by

$$\text{Prox}_{\gamma,g}(\theta) := \operatorname{argmin}_{\tau \in \Theta} \left\{ g(\tau) + \frac{1}{2\gamma} \|\theta - \tau\|^2 \right\}. \tag{9}$$

Note that under H1, for any $\gamma > 0$ and $\theta \in \mathbb{R}^d$, there exists a unique point τ minimizing the RHS of Eq. (9). This proximal operator may have an explicit expression. When g is the characteristic function

$$g(\theta) := \begin{cases} 0 & \text{if } \theta \in \Theta \\ +\infty & \text{otherwise,} \end{cases}$$

for some closed convex set $\Theta \subseteq \mathbb{R}^d$, then $g(\theta)$ is the projection of θ on Θ . This projection is explicit for example when Θ is an hyper-rectangle. Another example of explicit proximal operator is the case associated with the so-called elastic net penalty, i.e., $g_{\lambda,\alpha}(\theta) := \lambda \left(\frac{1-\alpha}{2} \sum_{i=1}^d \theta_i^2 + \alpha \sum_{i=1}^d |\theta_i| \right)$ with $\theta = (\theta_1, \dots, \theta_d)$, $\lambda > 0$ and $\alpha \in (0, 1]$, then for any component $i \in \{1, \dots, d\}$,

$$\begin{aligned} & (\text{Prox}_{\gamma,g_{\lambda,\alpha}}(\theta))_i \\ &= \frac{1}{1 + \gamma\lambda(1 - \alpha)} \begin{cases} 0 & \text{if } |\theta_i| \leq \gamma\lambda\alpha, \\ \theta_i - \gamma\lambda\alpha & \text{if } \theta_i \geq \gamma\lambda\alpha, \\ \theta_i + \gamma\lambda\alpha & \text{if } \theta_i \leq -\gamma\lambda\alpha. \end{cases} \end{aligned}$$

The proximal-gradient algorithm for solving the optimization problem Eq. (1) produces a sequence $\{\theta_n, n \geq 0\}$ as follows: given a $(0, 1/L]$ -valued sequence $\{\gamma_n, n \geq 0\}$,

$$\begin{aligned} \theta_{n+1} &= \text{Prox}_{\gamma_{n+1},g}(\theta_n + \gamma_{n+1} \nabla \ell(\theta_n)) \\ &= \text{Prox}_{\gamma_{n+1},g}(\theta_n + \gamma_{n+1} \{\nabla \phi(\theta_n) + \Psi(\theta_n) \bar{S}(\theta_n)\}). \end{aligned} \tag{10}$$

This update scheme can be explained as follows: by H2, we have for any $L \leq \gamma_{n+1}^{-1}$,

$$\begin{aligned} F(\theta) &= \ell(\theta) - g(\theta) \geq \ell(\theta_n) - \langle \nabla \ell(\theta_n), \theta - \theta_n \rangle \\ &\quad - \frac{1}{2\gamma_{n+1}} \|\theta - \theta_n\|^2 - g(\theta). \end{aligned}$$

This minorizing function is equal to $F(\theta_n)$ at the point θ_n ; the maximization (w.r.t. θ) of the RHS yields θ_{n+1} given by Eq. (10). The proximal-gradient algorithm is therefore a minorization–majorization (MM) algorithm and the ascent

property holds: $F(\theta_{n+1}) \geq F(\theta_n)$ for all n . Sufficient conditions for the convergence of the proximal-gradient algorithm Eq. (10) can be derived from the results by Combettes and Wajs (2005) and Parikh and Boyd (2013) or from convergence analysis of MM algorithms (see, e.g., Zangwill 1969; Meyer 1976).

In the case $\bar{S}(\theta)$ can not be computed, we describe two strategies for a Monte Carlo approximation. At iteration $n + 1$, given the current value of the parameter θ_n , m_{n+1} points $\{Z_{1,n}, \dots, Z_{m_{n+1},n}\}$ from the path of a Markov chain with target distribution $\pi_{\theta_n} d\nu$ are sampled. A first strategy consists in replacing $\bar{S}(\theta_n)$ by a Monte Carlo mean:

$$S_{n+1}^{mc} := \frac{1}{m_{n+1}} \sum_{j=1}^{m_{n+1}} S(Z_{j,n}). \tag{11}$$

A second strategy, inspired by stochastic approximation methods (see, e.g., Benveniste et al. 1990; Kushner and Yin 2003), consists in replacing $\bar{S}(\theta_n)$ by a stochastic approximation

$$S_{n+1}^{sa} := (1 - \delta_{n+1})S_n^{sa} + \frac{\delta_{n+1}}{m_{n+1}} \sum_{j=1}^{m_{n+1}} S(Z_{j,n}), \tag{12}$$

where $\{\delta_n, n \geq 0\}$ is a deterministic $[0, 1]$ -valued sequence. These two strategies yield, respectively, the Monte Carlo proximal-gradient (MCPG) algorithm (see Algorithm 1) and the stochastic approximation proximal-gradient (SAPG) algorithm (see Algorithm 2).

input : The initial values $\theta_0 \in \Theta$ and $Z_{m_0,-1} := z_*$, a $(0, 1/L]$ -valued sequence $\{\gamma_n, n \geq 0\}$ and an integer valued sequence $\{m_n, n \geq 0\}$

output : The sequence $\{\theta_n, n \geq 0\}$

```

1 for  $n \geq 0$  do
2   Simulation-step ;
3   sample a path  $Z_{1,n}, \dots, Z_{m_{n+1},n}$  of a Markov chain with
   invariant distribution  $\pi_{\theta_n} d\nu$  and started from  $Z_{m_n,n-1}$ ;
4   Expectation step ;
5   Compute  $S_{n+1}^{mc}$  as in Eq. (11) ;
6   Maximization step ;
7   Set  $\theta_{n+1} = \text{Prox}_{\gamma_{n+1},g} (\theta_n + \gamma_{n+1}\{\nabla\phi(\theta_n) + \Psi(\theta_n)S_{n+1}^{mc}\})$ 

```

Algorithm 1: The Monte Carlo proximal-gradient algorithm for the maximization of $\ell - g$

In Sect. 3, we prove the convergence of MCPG to the maximum points of F when ℓ is concave, for different choices of the sequences $\{\gamma_n, m_n, n \geq 0\}$ including decreasing or constant step sizes $\{\gamma_n, n \geq 0\}$ and, respectively, constant or increasing batch size $\{m_n, n \geq 0\}$. We also establish the convergence of SAPG to the maximum points (in the concave case); only the case of a constant batch size $\{m_n, n \geq 0\}$

input : The initial values $\theta_0 \in \Theta$, $Z_{m_0,-1} := z_*$ and $S_0^{sa} := s_*$, a $(0, 1/L]$ -valued sequence $\{\gamma_n, n \geq 0\}$, a $[0, 1]$ -valued sequence $\{\delta_n, n \geq 0\}$ and an integer valued sequence $\{m_n, n \geq 0\}$

output : The sequence $\{\theta_n, n \geq 0\}$

```

1 for  $n \geq 0$  do
2   Simulation-step ;
3   sample a path  $Z_{1,n}, \dots, Z_{m_{n+1},n}$  of a Markov chain with
   invariant distribution  $\pi_{\theta_n} d\nu$  and started from  $Z_{m_n,n-1}$ ;
4   Expectation step ;
5   Compute  $S_{n+1}^{sa}$  as in Eq. (12) ;
6   Maximization step ;
7   Set  $\theta_{n+1} = \text{Prox}_{\gamma_{n+1},g} (\theta_n + \gamma_{n+1}\{\nabla\phi(\theta_n) + \Psi(\theta_n)S_{n+1}^{sa}\})$ 

```

Algorithm 2: The stochastic approximation proximal-gradient algorithm for the maximization of $\ell - g$

and a decreasing step size $\{\gamma_n, n \geq 0\}$ is studied, since this framework corresponds to the *stochastic approximation* one from which the update rule Eq. (12) is inherited (see details in Delyon et al. 1999). From a numerical point of view, the choice of the sequences $\{\gamma_n, n \geq 0\}$, $\{\delta_n, n \geq 0\}$ and $\{m_n, n \geq 0\}$ is discussed in Sect. 4: guidelines are given in Sect. 4.2 and the behavior of the algorithm is illustrated through a toy example in Sect. 4.3.

2.2 Case of latent variable models from the exponential family

In this section, we consider the case when ℓ is given by Eq. (4). A classical approach to solve penalized maximum-likelihood problems in latent variables models with complete likelihood from the exponential family is the expectation-maximization (EM) algorithm or a generalization called the generalized EM (GEM) algorithm (Dempster et al. 1977; McLachlan and Krishnan 2008; Ng et al. 2012). Our goal here, is to show that MCPG and SAPG are stochastic perturbations of a GEM algorithm.

The EM algorithm is an iterative algorithm: at each iteration, given the current parameter θ_n , the quantity $Q(\theta|\theta_n)$, defined as the conditional expectation of the complete log likelihood under the a posteriori distribution for the current fit of the parameters, is computed:

$$Q(\theta|\theta') := \phi(\theta) + \langle \bar{S}(\theta'), \psi(\theta) \rangle. \tag{13}$$

The EM sequence $\{\theta_n, n \geq 0\}$ for the maximization of the penalized log likelihood $\ell - g$ is given by (see McLachlan and Krishnan 2008, Section 1.6.1)

$$\theta_{n+1} = \text{argmax}_{\theta \in \Theta} \{ \phi(\theta) + \langle \bar{S}(\theta_n), \psi(\theta) \rangle - g(\theta) \}. \tag{14}$$

When $\bar{S}(\theta)$ is intractable, it was proposed to replace $\bar{S}(\theta_n)$ in this EM-penalized algorithm by an approximation S_{n+1} —see Algorithm 3. When $S_{n+1} = S_{n+1}^{mc}$ (see Eq. 11), this yields

the so-called Monte Carlo EM-penalized algorithm (MCEM-pen), trivially adapted from MCEM proposed by Wei and Tanner (1990) and Levine and Fan (2004). Another popular strategy is to replace $\bar{S}(\theta_n)$ by S_{n+1}^{sa} (see Eq. 12) yielding to the so-called stochastic approximation EM-penalized algorithm (SAEM-pen)—(see Delyon et al. 1999 for the unpenalized version).

```

input : The initial value  $\theta_0 \in \Theta$ 
output : The sequence  $\{\theta_n, n \geq 0\}$ 
1 for  $n \geq 0$  do
2    $E$ -step ;
3   Compute an approximation  $S_{n+1}$  of  $\bar{S}(\theta_n)$  ;
4    $M$ -step ;
5   Set  $\theta_{n+1} = \operatorname{argmax}_{\theta} \{\phi(\theta) + \langle S_{n+1}, \psi(\theta) \rangle - g(\theta)\}$ 
    
```

Algorithm 3: Perturbed EM-penalized algorithms for the maximization of $\ell - g$

When the maximization of Eq. (13) is not explicit, the update of the parameter is modified as follows, yielding the generalized EM-penalized algorithm (GEM-pen):

$$\theta_{n+1} \text{ s.t. } \phi(\theta_{n+1}) + \langle \bar{S}(\theta_n), \psi(\theta_{n+1}) \rangle - g(\theta_{n+1}) \geq \phi(\theta_n) + \langle \bar{S}(\theta_n), \psi(\theta_n) \rangle - g(\theta_n). \tag{15}$$

This update rule still produces a sequence $\{\theta_n, n \geq 0\}$ satisfying the ascent property $F(\theta_{n+1}) \geq F(\theta_n)$ which is the key property for the convergence of EM (see, e.g., Wu 1983). Here again, the approximations defined in Eqs. (11) and (12) can be plugged in the GEM-pen update Eq. (15) when \bar{S} is not explicit.

We show in the following proposition that the sequence $\{\theta_n, n \geq 0\}$ produced by the proximal-gradient algorithm Eq. (10) is a GEM-pen sequence since it satisfies the inequality Eq. (15). As a consequence, MCPG and SAPG are stochastic GEM-pen algorithms.

Proposition 1 *Let g satisfying H1 and ℓ be of the form Eq. (4) with continuously differentiable functions $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^q$ and $S : \mathcal{Z} \rightarrow \mathbb{R}^q$. Set $\Theta := \{g + |\ell| < \infty\}$. Define $\bar{S} : \Theta \rightarrow \mathbb{R}^q$ by $\bar{S}(\theta) := \int_{\mathcal{Z}} S(z) \pi_{\theta}(z) \nu(dz)$ where π_{θ} is given by Eq. (5). Assume that there exists a constant $L > 0$ such that for any $s \in \bar{S}(\Theta)$, and any $\theta, \theta' \in \Theta$,*

$$\|\nabla\phi(\theta) - \nabla\phi(\theta') + (J\psi(\theta) - J\psi(\theta'))s\| \leq L\|\theta - \theta'\|. \tag{16}$$

Let $\{\gamma_n, n \geq 0\}$ be a (deterministic) positive sequence such that $\gamma_n \in (0, 1/L]$ for all $n \geq 0$.

Then the proximal-gradient algorithm Eq. (10) is a GEM-pen algorithm for the maximization of $\ell - g$.

The proof is postponed in ‘‘Appendix A’’. The assumption Eq. (16) holds when Θ is compact and \bar{S} (resp. ϕ and

ψ) are continuous (resp. twice continuously differentiable). Note also that for any $\theta, \theta' \in \Theta$ and $s \in \bar{S}(\Theta)$, we have $(J\psi(\theta) - J\psi(\theta'))s = 0$ if $\bar{S}(\theta) \in \operatorname{Ker}(J\psi(\theta'))$ for any $\theta, \theta' \in \Theta$.

3 Convergence of MCPG and SAPG

The convergence of MCPG and SAPG is established by applying recent results from Atchadé et al. (2017) on the convergence of perturbed proximal-gradient algorithms. Atchadé et al. (2017, Theorem 2) applied to the case $\nabla\ell(\theta)$ is of the form $\nabla\phi(\theta) + \Psi(\theta)\bar{S}(\theta)$, where $\bar{S}(\theta)$ is an intractable expectation and $\nabla\phi, \Psi$ are explicit, yields

Theorem 1 *Assume H1, H2, $\theta \mapsto \ell(\theta)$ is concave, and the set $\mathcal{L} := \operatorname{argmax}_{\theta \in \Theta} F(\theta)$ is a non-empty subset of Θ . Let $\{\theta_n, n \geq 0\}$ be given by*

$$\theta_{n+1} = \operatorname{Prox}_{\gamma_{n+1}, g}(\theta_n + \gamma_{n+1} \{\nabla\phi(\theta_n) + \Psi(\theta_n)S_{n+1}\}),$$

with a $(0, 1/L]$ -valued stepsize sequence $\{\gamma_n, n \geq 0\}$ satisfying $\sum_n \gamma_n = +\infty$. If the series

$$\begin{aligned} & \sum_n \gamma_{n+1} \langle \Psi(\theta_n)(S_{n+1} - \bar{S}(\theta_n)), T_{\gamma_{n+1}}(\theta_n) \rangle, \\ & \sum_n \gamma_{n+1} \Psi(\theta_n)(S_{n+1} - \bar{S}(\theta_n)), \\ & \sum_n \gamma_{n+1}^2 \|\Psi(\theta_n)(S_{n+1} - \bar{S}(\theta_n))\|^2, \end{aligned}$$

converge, where

$$T_{\gamma}(\theta) := \operatorname{Prox}_{\gamma, g}(\theta + \gamma \{\nabla\phi(\theta) + \Psi(\theta)\bar{S}(\theta)\}),$$

then there exists $\theta_{\infty} \in \mathcal{L}$ such that $\lim_n \theta_n = \theta_{\infty}$.

We check the conditions of Theorem 1 in the case S_{n+1} is resp. given by Eq. (11) for the proof of MCPG and by Eq. (12) for the proof of SAPG. Our convergence analysis is restricted to the case ℓ is concave; to our best knowledge, the convergence of the perturbed proximal-gradient algorithms when ℓ is not concave is an open question.

The novelty in this section is Proposition 2 and Theorem 4 which provide resp. a control of the L_2 -norm of the error $S_{n+1}^{sa} - \bar{S}(\theta_n)$ and the convergence of SAPG. These results rely on a rewriting of $(S_{n+1}^{sa} - \bar{S}(\theta_n))$ taking into account that S_{n+1}^{sa} is a weighted sum of the function S evaluated at all the samples $\{Z_{i,j}, i \leq m_{j+1}, j \leq n\}$ drawn from the initialization of the algorithm. This approximation differs from a more classical Monte Carlo approximation (see Theorems 2 and 3 for the convergence of MCPG, which are special cases of the results in Atchadé et al. 2017).

We allow the simulation step of MCPG and SAPG to rely on a Markov chain Monte Carlo sampling: at iteration $(n + 1)$, the conditional distribution of $Z_{j+1,n}$ given the past is $P_{\theta_n}(Z_{j,n}, \cdot)$ where P_{θ} is a Markov transition kernel having $\pi_{\theta}d\nu$ as its unique invariant distribution. The control of the quantities $S_{n+1} - \bar{S}(\theta_n)$ requires some ergodic properties on the kernels $\{P_{\theta_n}, n \geq 0\}$ along the path $\{\theta_n, n \geq 0\}$ produced by the algorithm. These properties have to be uniform in θ , a property often called the “containment condition” (see, e.g., the literature on the convergence of adaptive MCMC samplers, for example Andrieu and Moulines 2006; Roberts and Rosenthal 2007; Fort et al. 2011a). There are therefore three main strategies to prove the containment condition. In the first strategy, Θ is assumed to be bounded, and a uniform ergodic assumption on the kernels $\{P_{\theta}, \theta \in \Theta\}$ is assumed. In the second one, there is no boundedness assumption on Θ but the property $\mathbb{P}(\limsup_n \|\theta_n\| < \infty) = 1$ has to be established *prior* the proof of convergence; a kind of local boundedness condition on the sequence $\{\theta_n, n \geq 0\}$ is then applied—see, e.g., Andrieu and Moulines (2006) and Fort et al. (2011a). The last strategy consists in showing that $\mathbb{P}(\sup_n \rho_n \|\theta_n\| < \infty) = 1$ for some deterministic sequence $\{\rho_n, n \geq 0\}$ vanishing to zero when $n \rightarrow \infty$ at a rate compatible with the decaying ergodicity rate—see, e.g., Saksman and Vihola (2010). The last two strategies are really technical and require from the reader a strong background on controlled Markov chain theory; for pedagogical purposes, we therefore decided to state our results in the first context: we will assume that Θ is bounded.

By allowing MCMC approximations, we propose a theory which covers the case of a biased approximation, called below the *biased case*: conditionally to the past

$$\mathcal{F}_n := \sigma(Z_{i,j}, i \leq m_{j+1}, j \leq n - 1), \quad (17)$$

the expectation of S_{n+1} is not $\bar{S}(\theta_n)$: $\mathbb{E}[S_{n+1}|\mathcal{F}_n] \neq \bar{S}(\theta_n)$. As soon as the samplers $\{P_{\theta}, \theta \in \Theta\}$ are ergodic enough (for example, under (H4a) and (H4b)), the bias vanishes when the number of Monte Carlo points m_n tends to infinity. Therefore, the proof for the biased case when the sequence $\{m_n, n \geq 0\}$ is constant is the most technical situation since the bias does not decay. It relies on a specific decomposition of the error $S_{n+1} - \bar{S}(\theta_n)$ into a martingale increment with bounded L^2 -moments, and a remainder term which vanishes when $n \rightarrow \infty$ even when the batch size m_n is constant. Such a behavior of the remainder term is a consequence of regularity properties on the functions $\nabla\phi, \Psi, \bar{S}$ (see H3c), on the proximity operator (see H3d) and on the kernels $\{P_{\theta}, \theta \in \Theta\}$ (see H4c).

Our theory also covers the *unbiased case*, i.e., when

$$\mathbb{E}[S_{n+1}|\mathcal{F}_n] = \bar{S}(\theta_n).$$

We therefore establish the convergence of MCPG and SAPG by strengthening the conditions H1 and H2 with

- H3** (a) ℓ is concave and the set $\mathcal{L} := \operatorname{argmax}_{\Theta} F$ is a non-empty subset of Θ .
 (b) Θ is bounded.
 (c) There exists a constant L such that for any $\theta, \theta' \in \Theta$,

$$\begin{aligned} & \|\nabla\phi(\theta) - \nabla\phi(\theta')\| + \|\Psi(\theta) - \Psi(\theta')\| \\ & + \|\bar{S}(\theta) - \bar{S}(\theta')\| \leq L\|\theta - \theta'\|, \end{aligned}$$

where for a matrix A , $\|A\|$ denotes the operator norm associated with the Euclidean vector norm.

- (d) $\sup_{\gamma \in (0, 1/L]} \sup_{\theta \in \Theta} \gamma^{-1} \|\operatorname{Prox}_{\gamma, g}(\theta) - \theta\| < \infty$.

Note that the assumptions (H3b)–(H3c) imply Eq. (3) and $\sup_{\theta \in \Theta} (\|\nabla\phi(\theta)\| + \|\Psi(\theta)\| + \|\bar{S}(\theta)\|) < \infty$. When Θ is a compact convex set, then (H3d) holds for the elastic net penalty, the Lasso or the fused Lasso penalty. Atchadé et al. (2017, Proposition 11) give general conditions for (H3d) to hold.

Before stating the ergodicity conditions on the kernels $\{P_{\theta}, \theta \in \Theta\}$, let us recall some basic properties on Markov kernels. A Markov kernel P on the measurable set (Z, \mathcal{Z}) is an application on $Z \times \mathcal{Z}$, taking values in $[0, 1]$ such that for any $x \in Z$, $P(x, \cdot)$ is a probability measure on \mathcal{Z} ; and for any $A \in \mathcal{Z}$, $x \mapsto P(x, A)$ is measurable. Furthermore, if P is a Markov kernel, P^k denotes the k th iterate of P defined by induction as

$$P^0(x, A) := \mathbb{1}_A(x),$$

$$P^k(x, A) := \int P^{k-1}(x, dz)P(z, A), \quad k \geq 1.$$

Finally, the kernel P acts on the probability measures: for any probability measure ξ on \mathcal{Z} , ξP is a probability measure defined by

$$\xi P(A) := \int \xi(dz)P(z, A), \quad A \in \mathcal{Z};$$

and P acts on the positive measurable functions: for a measurable function $f: Z \rightarrow \mathbb{R}_+$, Pf is a measurable function defined by

$$Pf(z) := \int f(y)P(z, dy).$$

We refer the reader to Meyn and Tweedie (2009) for the definitions and basic properties on Markov chains. Given a measurable function $W: Z \rightarrow [1, +\infty)$, define the W -norm of a signed measure ν on \mathcal{Z} and the W -norm of a function $f: Z \rightarrow \mathbb{R}^d$:

$$\|f\|_W := \sup_Z \frac{\|f\|}{W}, \quad \|\nu\|_W := \sup_{f: \|f\|_W \leq 1} \left| \int f d\nu \right|;$$

these norms generalize resp. the supremum norm of a function and the total variation norm of a measure.

Our results are derived under the following conditions on the kernels:

H4 (a) There exist $\lambda \in (0, 1], b < \infty$ and a measurable function $W : Z \rightarrow [1, +\infty)$ such that

$$|S|_{\sqrt{W}} < \infty, \quad \sup_{\theta \in \Theta} P_\theta W \leq \lambda W + b.$$

(b) There exist constants $C < \infty$ and $\rho \in (0, 1)$ such that for any $z \in Z$ and $n \geq 0$,

$$\sup_{\theta \in \Theta} \|P_\theta^n(z, \cdot) - \pi_\theta\|_W \leq C \rho^n W(z).$$

(c) There exists a constant C such that for any $\theta, \theta' \in \Theta$,

$$\begin{aligned} \|\pi_\theta - \pi_{\theta'}\|_{\sqrt{W}} + \sup_{z \in Z} \frac{\|P_\theta(z, \cdot) - P_{\theta'}(z, \cdot)\|_{\sqrt{W}}}{\sqrt{W}(z)} \\ \leq C \|\theta - \theta'\|. \end{aligned}$$

Sufficient conditions for the uniform-in- θ ergodic behavior (H4b) are given, e.g., in Fort et al. (2011b, Lemma 2.3): this lemma shows how to deduce such a control from a minorization condition and a drift inequality on the Markov kernels. Examples of MCMC kernels P_θ satisfying these assumptions can be found in Andrieu and Moulines (2006, Proposition 12) and Saksman and Vihola (2010, Proposition 15) for the adaptive Hastings–Metropolis algorithm, in Fort et al. (2011b, Proposition 3.1) for an interactive tempering sampler, in Schreck et al. (2013, Proposition 3.2) for the equi-energy sampler, and in Fort et al. (2015, Proposition 3.1) for a Wang–Landau type sampler.

Theorem 2 establishes the convergence of MCPG when the number of points in the Monte Carlo sum S_{n+1}^{mc} is constant over iterations and the step size sequence $\{\gamma_n, n \geq 0\}$ vanishes at a convenient rate. It is proved in Atchadé et al. (2017, Theorem 4).

Theorem 2 Assume H1, H2, (H3a–c) and (H4a–b). Let $\{\theta_n, n \geq 0\}$ be the sequence given by Algorithm 1 with a $(0, 1/L]$ -valued sequence $\{\gamma_n, n \geq 0\}$ such that $\sum_n \gamma_n = +\infty$ and $\sum_n \gamma_n^2 < \infty$, and with a constant sequence $\{m_n, n \geq 0\}$.

In the biased case, assume also (H3d) and (H4c) and $\sum_n |\gamma_{n+1} - \gamma_n| < \infty$.

Then, with probability one, there exists $\theta_\infty \in \mathcal{L}$ such that $\lim_n \theta_n = \theta_\infty$.

Theorem 3 establishes the convergence of MCPG when the number of points in the Monte Carlo sum S_{n+1}^{mc} is increasing; it allows a constant stepsize sequence $\{\gamma_n, n \geq 0\}$. It is proved in Atchadé et al. (2017, Theorem 6).

Theorem 3 Assume H1, H2, (H3a–c) and (H4a–b). Let $\{\theta_n, n \geq 0\}$ be the sequence given by Algorithm 1 with a $(0, 1/L]$ -valued sequence $\{\gamma_n, n \geq 0\}$ and an integer valued sequence $\{m_n, n \geq 0\}$ such that $\sum_n \gamma_n = +\infty$ and $\sum_n \gamma_n^2/m_n < \infty$.

In the biased case, assume also $\sum_n \gamma_n/m_n < \infty$.

Then, with probability one, there exists $\theta_\infty \in \mathcal{L}$ such that $\lim_n \theta_n = \theta_\infty$.

MCPG and SAPG differ in their approximation of $\bar{S}(\theta_n)$ at each iteration. We provide below a control of this error for a constant or a polynomially increasing batch size $\{m_n, n \geq 0\}$, and polynomially decreasing stepsize sequences $\{\gamma_n, n \geq 0\}$ and $\{\delta_n, n \geq 0\}$.

Proposition 2 Let $\gamma_\star, \delta_\star, m_\star$ be positive constants and $\beta \in [0, 1), \alpha \geq \beta, c \geq 0$. Set $\gamma_n = \gamma_\star n^{-\alpha}, \delta_n = \delta_\star n^{-\beta}$ and $m_n = m_\star n^c$. Assume H1 to H4. Then

$$\begin{aligned} \mathbb{E} \left[\|S_{n+1}^{mc} - \bar{S}(\theta_n)\|^2 \right] &= O(n^{-c}), \\ \mathbb{E} \left[\|S_{n+1}^{sa} - \bar{S}(\theta_n)\|^2 \right] &= O\left(n^{-\{2(\alpha-\beta) \wedge (\beta+c)\}}\right). \end{aligned}$$

The proof is given in ‘‘Appendix C’’. This proposition shows that when applying MCPG with a constant batch size ($c = 0$), the error $S_{n+1}^{mc} - \bar{S}(\theta_n)$ does not vanish; this is not the case for SAPG, since even when $c = 0$, the error $S_{n+1}^{sa} - \bar{S}(\theta_n)$ vanishes as soon as $\alpha > \beta > 0$. Since the case ‘‘constant batch size’’ is the usual choice of the practitioners in order to reduce the computational cost of the algorithm, this proposition supports the use of SAPG instead of MCPG.

We finally study the convergence of SAPG without assuming that the batch size sequence $\{m_n, n \geq 0\}$ is constant, which implies the following assumption on the sequences $\{\gamma_n, \delta_n, m_n, n \geq 0\}$.

H5 The step size sequences $\{\gamma_n, n \geq 0\}, \{\delta_n, n \geq 0\}$ and the batch size sequence $\{m_n, n \geq 0\}$ satisfy

$$(a) \quad \gamma_n \in (0, 1/L], \delta_n \in (0, 1), m_n \in \mathbb{N}, \sum_n \gamma_n = +\infty, \sum_n \gamma_n^2 < \infty,$$

$$\begin{aligned} \sum_n \left(\gamma_{n-1} \gamma_n + \gamma_{n-1}^2 + |\gamma_n - \gamma_{n-1}| \right) D_n < \infty, \\ \sum_n \gamma_n^2 \delta_n^2 (1 + D_{n+1})^2 m_n^{-1} < \infty, \end{aligned}$$

$$\text{where } D_n := \sum_{k \geq n} \left(\prod_{j=n}^k (1 - \delta_j) \right).$$

(b) Furthermore,

$$\begin{aligned} \sum_n \gamma_{n+1} |m_{n+1}^{-1} \delta_{n+1} - m_n^{-1} \delta_n| < \infty, \\ \sum_n \gamma_{n+1} |m_{n+1}^{-1} \delta_{n+1} D_{n+2} - m_n^{-1} \delta_n D_{n+1}| < \infty, \end{aligned}$$

$$\sum_n \left(\gamma_{n-1} \gamma_n + \gamma_{n-1}^2 + |\gamma_n - \gamma_{n-1}| \right) \dots \times m_{n-1}^{-1} \delta_{n-1} (1 + D_n) < \infty.$$

Let us comment this assumption in the case the batch size sequence $\{m_n, n \geq 0\}$ is constant. This situation corresponds to the “stochastic approximation regime” where the number of draws at each iteration is $m_n = 1$ (or say, $m_n = m$ for any n), and it also corresponds to what is usually done by practitioners in order to reduce the computational cost. When $\delta_n = \delta_\star \in (0, 1)$ for any $n \geq 0$, then $D_n = \delta_\star^{-1}$ for any $n \geq 0$. This implies that the condition H5 is satisfied with polynomially decreasing sequences $\gamma_n \sim \gamma_\star/n^\alpha$ with $\alpha \in (1/2, 1]$ (and $m_n = m$ for any n).

When $\delta_n \sim \delta_\star n^{-\beta}$ for $\beta \in (0, 1)$, then $D_n = O(n^\beta)$ (see Lemma 3). Hence, using Lemma 3, (H5a) and (H5b) are satisfied with $\gamma_n \sim \gamma_\star n^{-\alpha}$ where $\beta < (1 + \beta)/2 < \alpha \leq 1$, and $m_n = m$ for any n .

We cannot have $\delta_n = \delta_\star n^{-1}$ since it implies $D_n = +\infty$ for any $n \geq 0$.

Theorem 4 Assume H1, H2, H3 and (H4a–b). Let $\{\theta_n, n \geq 0\}$ be the sequence given by Algorithm 2 and applied with sequences $\{\gamma_n, \delta_n, m_n, n \geq 0\}$ verifying (H5a).

In the biased case, assume also (H4c) and (H5b).

Then with probability one, there exists $\theta_\infty \in \mathcal{L}$ such that $\lim_n \theta_n = \theta_\infty$.

Proof The proof is in Section D. □

4 Numerical illustration in the convex case

In this section, we illustrate the behavior of the algorithms MCPG and SAPG on a toy example. We first introduce the example and then give some guidelines for a specific choice of the sequences $\{\delta_n, n \geq 0\}, \{\gamma_n, n \geq 0\}$. Finally, the algorithms are compared more systematically on repeated simulations.

4.1 A toy example

The example is a mixed model, where the regression function is linear in the latent variable Z . More precisely, we observe data (Y_1, \dots, Y_N) from N subjects, each individual data being a vector of size J : $Y_k := (Y_{k1}, \dots, Y_{kJ})$. For the subject $k, k = 1, \dots, N, Y_{kj}$ is the j th measurement at time $t_{kj}, j = 1, \dots, J$. It is assumed that $\{Y_k, k = 1, \dots, N\}$ are independent and for all $k = 1, \dots, N,$

$$Y_{kj} | Z^{(k)} \stackrel{\text{ind}}{\sim} \mathcal{N} \left(\left\langle Z^{(k)}, \bar{t}_{kj} \right\rangle, 1 \right), \tag{18}$$

$$\bar{t}_{kj} := \begin{bmatrix} 1 \\ t_{kj} \end{bmatrix} \quad j = 1, \dots, J;$$

that is, a linear regression model with individual random intercept and slope, the \mathbb{R}^2 -valued vector being denoted by $Z^{(k)}$. The latent variable is $Z = (Z^{(1)}, \dots, Z^{(N)})$. Furthermore,

$$Z^{(k)} \stackrel{\text{ind}}{\sim} \mathcal{N}_2(X_k \theta, I_2); \tag{19}$$

here, $\theta \in \mathbb{R}^{2(D+1)}$ is an unknown parameter and the design matrix $X_k \in \mathbb{R}^{2 \times 2(D+1)}$ is known

$$X_k := \begin{bmatrix} 1 & X_{k1} & \dots & X_{kD} & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 1 & X_{k1} & \dots & X_{kD} \end{bmatrix}. \tag{20}$$

The optimization problem of the form Eq. (1) that we consider is the log likelihood $\ell(\theta)$ penalized by a lasso penalty: the objective is the selection of the influential covariates

$$(X_{k1}, \dots, X_{kD})$$

on the two components of $Z^{(k)}$. We thus penalize all the elements except θ_1 and θ_{D+2} which correspond to the two intercepts; hence, we set

$$g(\theta) := \lambda \sum_{r \neq \{1, D+2\}} |\theta_r|.$$

The above model is a latent variable model with complete log likelihood equal to—up to an additive constant

$$-\frac{1}{2} \sum_{k=1}^N \left\{ \sum_{j=1}^J \left(Y_{kj} - \langle Z^{(k)}, \bar{t}_{kj} \rangle \right)^2 + \|Z^{(k)} - X_k \theta\|^2 \right\}.$$

It is of the form $\phi(\theta) + \langle S(z), \psi(\theta) \rangle$ by setting (with $(\cdot)'$ denoting the transpose of a matrix)

$$\phi(\theta) := -\frac{1}{2} \theta' \left(\sum_{k=1}^N X_k' X_k \right) \theta - \frac{1}{2} \sum_{k=1}^N \sum_{j=1}^J Y_{kj}^2,$$

$$\psi(\theta) := \begin{bmatrix} 1 \\ \theta \end{bmatrix} \in \mathbb{R}^{1+2(D+1)},$$

$$S(z^{(1)}, \dots, z^{(N)})$$

$$:= -\frac{1}{2} \sum_{k=1}^N \begin{bmatrix} z^{(k)'} (I + T_k) z^{(k)} - 2 \langle z^{(k)}, \bar{Y}_k \rangle \\ -2 X_k' z^{(k)} \end{bmatrix},$$

$$T_k := \sum_{j=1}^J \bar{t}_{kj} \bar{t}_{kj}',$$

$$\bar{Y}_k := \sum_{j=1}^J Y_{kj} \bar{t}_{kj}.$$

The a posteriori distribution π_θ is a Gaussian distribution on \mathbb{R}^{2N} , equal to the product of N Gaussian distributions on \mathbb{R}^2 :

$$\pi_\theta(z^{(1)}, \dots, z^{(N)}) := \prod_{k=1}^N \mathcal{N}_2 \left((I + T_k)^{-1} (\bar{Y}_k + X_k \theta), (I + T_k)^{-1} \right) [z^{(k)}]. \tag{21}$$

Hence, $\bar{S}(\theta)$ is explicit and given by

$$\bar{S}(\theta) = -\frac{1}{2} \sum_{k=1}^N \left[\text{Trace}((I + T_k)\Sigma_k) - 2\bar{Y}'_k(I + T_k)^{-1}(\bar{Y}_k + X_k\theta) - 2X'_k(I + T_k)^{-1}(\bar{Y}_k + X_k\theta) \right] \tag{22}$$

with

$$\Sigma_k := (I + T_k)^{-1} + (I + T_k)^{-1} (\bar{Y}_k + X_k\theta) (\bar{Y}_k + X_k\theta)' (I + T_k)^{-1}. \tag{23}$$

Finally, note that in this example, the function ℓ is explicit and given by (up to an additive constant)

$$\ell(\theta) = -\frac{1}{2}\theta' \left(\sum_{k=1}^N X'_k X_k \right) \theta + \frac{1}{2} \sum_{k=1}^N (\bar{Y}_k + X_k\theta)' (I + T_k)^{-1} (\bar{Y}_k + X_k\theta).$$

Thus ℓ is a concave function. Furthermore, in this toy example, $\theta \mapsto \nabla \ell(\theta)$ is linear so that the Lipschitz constant L is explicit and equal to

$$L = \left\| -\sum_{k=1}^N X'_k X_k + \sum_{k=1}^N X'_k (I + T_k)^{-1} X_k \right\|, \tag{24}$$

where for a matrix A , $\|A\|_2$ denotes the spectral norm. Finally, we assumed that $\Theta = \{\theta \in \mathbb{R}^{2(D+1)} \mid \|\theta\| < 10^4\}$ to fulfill the theoretical boundedness assumption. The MCMC algorithm includes a projection step on Θ if necessary. But in practice, it never happens.

A data set is simulated using this model with $N = 40$, $J = 8$, $D = 300$ and $t_{kj} \in \{0.25, 4, 6, 8, 10, 12, 14, 16\}$, $\forall k \in \{1, \dots, N\}$. The design components (X_{k1}, \dots, X_{kD}) (see Eq. 20) are drawn from a centered Gaussian distribution with covariance matrix Γ defined by $\Gamma_{rr'} = 0.5^{|r-r'|}$ ($r, r' = 1, \dots, 300$). To sample the observations, we use a parameter vector θ^* defined as follows: $\theta_1^* = \theta_{D+2}^* = 1$; the other components are set to zero, except 12 components randomly selected (6 among the components $\{2, \dots, D + 1\}$ and 6 among the components $\{D + 3, \dots, 2D + 2\}$) and chosen uniformly in $[0.5, 1.5]$ —see the last row in Fig. 7.

4.2 Guidelines for the implementation

In this section, we give some guidelines on the choice of the sequences $\{\delta_n, n \geq 0\}$ and $\{\gamma_n, n \geq 0\}$. We illustrate the results on single runs of each algorithm. We use the same random draws for all the algorithms to avoid potential differences due to the randomness of the simulations. Similar results have been observed when simulations are replicated. We refer to Sect. 4.3 for replicated simulations.

Classical sequences $\{\delta_n, n \geq 0\}$ and $\{\gamma_n, n \geq 0\}$ are of the form:

$$\gamma_{n+1} = \begin{cases} \gamma_\star & \text{if } n \leq n_\alpha, \\ \gamma_\star(n - n_\alpha)^{-\alpha} & \text{if } n > n_\alpha, \end{cases} \tag{25}$$

$$\delta_{n+1} = \begin{cases} \delta_\star & \text{if } n \leq n_\beta, \\ \delta_\star(n - n_\beta)^{-\beta} & \text{if } n > n_\beta. \end{cases} \tag{26}$$

Impact of γ_\star and δ_\star on the transient phase The theoretical study on the asymptotic behavior of SAPG and MCPG is derived under the assumption that $\gamma_n \leq 1/L$: when $\alpha > 0$, this property holds for any n large enough. In this section, we illustrate the role of γ_n, δ_n for small values of n that is, in the transient phase of the algorithm. In Fig. 1, we display the behavior of MCPG and SAPG for two different values of the initial point $\theta_{n=0}$: on the left, it corresponds to a standard initialization ($\theta_{n=0} = (0, \dots, 0)$), while on the right, it corresponds to a poor initialization—which mimics what may happen in practice for challenging numerical applications.

On both plots, we indicate by a vertical line the smallest n such that $\gamma_n \leq 1/L$ —remember that in this example, L is explicit (see Eq. 24). The plots show the estimation of component #245, as a function of the number of iterations n . In all cases, $n_\alpha = n_\beta = 0, \alpha = 0.75, m_n = 60$, and for SAPG, $\beta = 0.5$. The dotted blue curve displays a run of SAPG when $(\gamma_\star, \delta_\star) = (0.009, 0.2)$; the dashed-dotted yellow curve displays a run of SAPG when $(\gamma_\star, \delta_\star) = (0.009, 0.5)$; the dashed red curve displays a run of SAPG when $(\gamma_\star, \delta_\star) = (0.009, 0.8)$; the green solid curve displays a run of MCPG when $\gamma_\star = 0.009$.

The stability of MCPG during the transient phase depends crucially on the first values of the sequence $\{\gamma_n, n \geq 0\}$. Then when n is large enough so that $\gamma_n \leq 1/L$ (after the vertical line), MCPG is more stable and gets smoother. For SAPG, a small value of δ_\star implies an important impact of the initial point $\theta_{n=0}$. When this initial point is poorly chosen, a small value of δ_\star delays the convergence of SAPG. A value of δ_\star around 0.5 is a good compromise.

Role of α and β : Figure 2 displays the behavior of SAPG for different values of α and β with $(\gamma_\star, \delta_\star) = (0.015, 0.5), n_\alpha = n_\beta = 0$ and $m_n = 60$. The plots show that the larger the parameter α is, the longer the transient phase is. We then recommend to set α close to 0.6. The parameter β




Fig. 1 Estimation of the component #245 of the vector θ along 10,000 iterations (x -axis is in \log_{10} scale) for MCPG and SAPG. MCPG is represented in green solid line. A run of SAPG is displayed in dashed red line in the case $\delta_\star = 0.8$; in dashed-dotted yellow line in the case $\delta_\star = 0.5$ and in dotted blue line in the case $\delta_\star = 0.2$. For all runs,

$\gamma_\star = 0.009, n_\alpha = n_\beta = 0, (\alpha, \beta) = (0.75, 0.5)$ and $m_n = 60$. The vertical black dashed line corresponds to the smallest n such that $\gamma_n \leq 1/L$. Left: standard initialization of θ ($\theta_{n=0}$ is the null vector). Right: poor initialization of θ . The penalty term λ was set to 50 for all the runs. (Color figure online)




Fig. 2 Estimation of the component #245 of the vector θ along 10,000 iterations (x -axis is in \log_{10} scale). Runs of SAPG with $(\gamma_\star, \delta_\star) = (0.015, 0.5), n_\alpha = n_\beta = 0$ and $m_n = 60$ and different values of (α, β) . (Left) $\alpha = 1$ and $\beta = 0.9$ (green solid line), $\beta = 0.6$ (red dashed line), $\beta = 0.3$ (yellow dotted line), $\beta = 0$ (blue dash-dotted line). (Middle)

$\alpha = 0.8$ and $\beta = 0.5$ (green solid line), $\beta = 0.3$ (red dashed line), $\beta = 0.1$ (yellow dotted line), $\beta = 0$ (blue dash-dotted line). (Right) $\alpha = 0.6$ and $\beta = 0.2$ (green solid line), $\beta = 0.1$ (yellow dotted line), $\beta = 0$ (blue dash-dotted line). The penalty term λ was set to 50 for all the runs. (Color figure online)

seems to have an impact only when α is close to 1. Therefore, we recommend to set δ_n constant during the transient phase ($n_\beta > 0$) and then to decrease it rapidly in the convergence phase.

Random stepsize sequence $\{\gamma_n, n \geq 0\}$: The convergence of the SAPG algorithm can suffer from the scale difference of the parameters, when run with the same stepsize sequence $\{\gamma_n, n \geq 0\}$ applied to each component of θ_n .

Ideally each component of θ_n should have a specific γ_n value adapted to its scale. But it can be time-consuming to find, by hand-tuning, a sequence that ensures a fast and stable convergence of the algorithm. As an alternative, we suggest to use a matrix-valued random sequence $\{\Gamma^n, n \geq 0\}$ and replace the update rule of SAPG by

$$(\theta_{n+1})_i = \text{Prox}_{\Gamma_{ii}^{n+1} g} \left((\theta_n)_i + \Gamma_{ii}^{n+1} (\nabla \phi(\theta_n) + \Psi(\theta_n) S_{n+1}^{\text{sa}})_i \right).$$

We propose to define the matrix Γ_{n+1} as a diagonal matrix with entries Γ_{ii}^{n+1} depending on H_{ii}^n , where H^n is an approximation of the hessian of $-\ell(\theta)$. (We give an example of such an approximation in Sect. 5.) Through numerical experiments, we observed that asymptotically, H^n converges. Hence, to ensure a stepsize sequence decaying like $O(n^{-\alpha})$ asymptotically, we propose the following definition of the random sequence:

$$\Gamma_{ii}^{n+1} = \begin{cases} 1/|H_{ii}^n| & \text{if } n \leq n_\alpha, \\ ((n - n_\alpha)^\alpha |H_{ii}^n|)^{-1} & \text{if } n > n_\alpha. \end{cases} \tag{27}$$

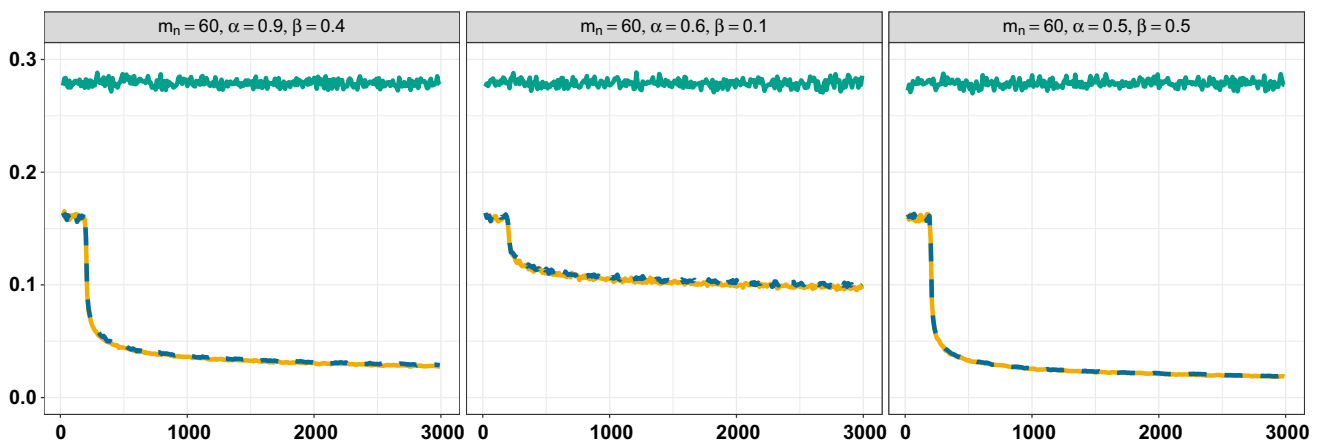


Fig. 3 Evolution of the Monte Carlo approximation of $\|S_{n+1} - \bar{S}(\theta_n)\|_2$ with iterations n for algorithms MCPG (solid green), SAEM-pen (solid yellow), SAPG (dashed blue), implemented with $(\alpha, \beta) = (0.9, 0.4)$

[left], $(\alpha, \beta) = (0.6, 0.1)$ [center] and $(\alpha, \beta) = (0.5, 0.5)$ [right]; for MCPG and SAPG, the batch size is fixed $m_n = 60$. (Color figure online)

4.3 Long-time behavior of the algorithm

In this section, we illustrate numerically the theoretical results on the long term convergence of the algorithms MCPG, SAPG and SAEM-pen (i.e., Algorithm 3 applied with $S_{n+1} = S_{n+1}^{sa}$) and EM-pen on the toy model. In this example, the exact algorithm EM-pen (see Eq. 14) applies: the quantity $\bar{S}(\theta)$ is an explicit expectation under a Gaussian distribution π_θ . Therefore, we use this example (i) to illustrate the convergence of the three stochastic methods to the same limit point as EM-pen, (ii) to compare the two approximations S_{n+1}^{mc} and S_{n+1}^{sa} of $\bar{S}(\theta_n)$ in a GEM-pen approach and (iii) to study the effect of relaxing the M-step by comparing the GEM-pen and EM-pen approaches namely SAPG and SAEM-pen.

The sequences $\{\gamma_n, n \geq 0\}$ and $\{\delta_n, n \geq 0\}$ are defined as follows: $(\gamma_*, \delta_*) = (0.004, 0.5)$, and $n_\alpha = n_\beta = 0$; three different pairs (α, β) are considered: $(\alpha, \beta) = (0.9, 0.4)$, $(\alpha, \beta) = (0.6, 0.1)$, and $(\alpha, \beta) = (0.5, 0.5)$. The algorithms are implemented with a fixed batch size $m_n = 60$. 100 independent runs of each algorithm are performed. For the penalty term, we set $\lambda = 50$. In MCPG, SAPG and SAEM-pen, the simulation step at iteration $(n + 1)$ relies on exact sampling from π_{θ_n} —see Eq. (21); therefore, in this toy example, the Monte Carlo approximation of $\bar{S}(\theta_n)$ is unbiased.

In Fig. 3, for the three algorithms MCPG, SAPG and SAEM-pen, the evolution of an approximation of $\|S_{n+1} - \bar{S}(\theta_n)\|_2$ with iterations n is plotted, where, for a random variable U , $\|U\|_2 := \sqrt{\mathbb{E}[\|U\|^2]}$. This L_2 -norm is approximated by a Monte Carlo sum computed from 100 independent realizations of S_{n+1} ; here, $\bar{S}(\theta_n)$ is explicit (see Eq. 22). SAEM-pen and SAPG behave similarly; the L_2 -

norm converges to 0, and the convergence is slower when $(\alpha, \beta) = (0.6, 0.1)$ —this plot illustrates the result stated in Proposition 2, Sect. 3. This convergence does not hold for MCPG because the size m_n of the Monte Carlo approximation is kept fixed.

We compared the limiting vectors $\lim_n \theta_n$ obtained by each algorithm, over the 100 independent runs. They are all equal, and the limiting vector is also the limiting value θ_∞ of the EM-pen algorithm. In order to discuss the rate of convergence, we show the behavior of the algorithms when estimating the component #245 of the regression coefficients; this component was chosen among the non-null component of θ_∞ . Figure 4 shows the boxplot of 100 estimations of the component #245 of the vector θ_n , when $n = 5, 25, 50, 500$, for the algorithms MCPG, SAPG and SAEM-pen with $(\alpha, \beta) = (0.9, 0.4)$. Here, SAPG and MCPG behave similarly, with a smaller variability among the 100 runs than SAEM-pen. SAEM-pen converges faster than SAPG and MCPG which was expected since they correspond, respectively, to stochastic perturbations of EM-pen and GEM-pen algorithms. Figure 5 shows the boxplot of 100 estimations by MCPG, SAPG and SAEM-pen of the component #245 after $n = 500$ iterations with different values for the parameters α and β . We observe that the three algorithms give similar final estimates for the three conditions on parameters α and β . This is due to the fact that with $n_\alpha = n_\beta = 200$, the algorithms have already attained the convergence phase when $n = 200$. This allows the algorithms to quickly converge toward the limit points when $n > 200$.

Figure 6 shows the convergence of a Monte Carlo approximation of $n \mapsto \mathbb{E}[F(\theta_n)]$ based on 100 independent estimations θ_n obtained by three different algorithms: EM-pen, MCPG, SAPG and SAEM-pen run with $(\alpha, \beta) = (0.9, 0.4)$ and $m_n = 60$. Here again, all the algorithms




Fig. 4 Estimation of the component #245 of the vector θ_n when $n = 5, 25, 50, 500$. MCPG (green, left), SAPG (blue, middle) and SAEM-pen (yellow, right) are implemented with $(\alpha, \beta) = (0.9, 0.4)$; for MCPG and SAPG, the batch size is fixed $m_n = 60$. Each boxplot is computed from 100 independent runs. Black dashed line corresponds to the value obtained with EM-pen algorithm at iteration 500. (Color figure online)




Fig. 6 Monte Carlo approximation of $\mathbb{E}[F(\theta_n)]$ (based on 100 independent samples) along the iterations n , for algorithms EM-pen (solid black), MCPG (solid green), SAEM-pen (dash-dotted yellow), SAPG (dashed blue), implemented with $(\alpha, \beta) = (0.9, 0.4)$ and $m_n = 60$. (Color figure online)

converge to the same value and EM-pen and SAEM-pen converge faster than MCPG and SAPG. We observe that the path of SAPG is far more smooth than the path of MCPG.

Finally, Fig. 7 shows the support of the vector $\lim_n \theta_n$ (where the component θ_1 and θ_{302} are removed) estimated by MCPG, SAPG, SAEM-pen and EM-pen (the estimated support is the same for the four algorithms). The frequency, among 100 independent runs, for each component to be in the support of the limit value $\lim_n \theta_n$, is displayed. Algorithms are implemented with $(\alpha, \beta) = (0.9, 0.4)$ and $m_n = 60$. For all algorithms, we observe that most of the non-null components of $\lim_n \theta_n$ are non-null components of θ^* . Note also

that the stochastic algorithms MCPG, SAPG and SAEM-pen converge to the same vector as EM-pen.

5 Inference in nonlinear mixed models for pharmacokinetic data

In this section, SAPG is applied to solve a more challenging problem. The objective is to illustrate the algorithm in cases that are not covered by the theory. The application is in pharmacokinetic analysis, with nonlinear mixed effect models




Fig. 5 Estimation of the component #245 of the vector θ_n when $n = 500$. MCPG (green), SAPG (blue) and SAEM-pen (yellow) are implemented with $(\alpha, \beta) = (0.9, 0.4)$ [left], $(\alpha, \beta) = (0.6, 0.1)$ [center] and $(\alpha, \beta) = (0.5, 0.5)$ [right]; for MCPG and SAPG, the batch

size is fixed $m_n = 60$. Each boxplot is computed from 100 independent runs. Black dashed line corresponds to the value obtained with EM-pen algorithm at iteration 500. (Color figure online)




Fig. 7 (Top) Support of $\lim_n \theta_n$ estimated by all the algorithms MCPG, SAPG, SAEM-pen and EM-pen over 100 runs for $(\alpha, \beta) = (0.9, 0.4)$ and $m_n = 60$. (Bottom) The support of θ^* used to produce the observations. On both rows, the components 1 and $D + 1$ are not displayed

(NLMEM); in this application, the penalized maximum-likelihood inference is usually solved by the SAEM-pen algorithm, possibly combined with an approximation of the M-step when it is non-explicit. This section also provides a numerical comparison of SAPG and SAEM-pen. Both algorithms have a simulation step; in this more challenging application, it will rely on a Markov chain Monte Carlo (MCMC) sampler—see Sect. 5.1. Therefore, for both algorithms, $\bar{S}(\theta)$ is approximated by a biased Monte Carlo sum.

We start with a presentation of the statistical analysis and its translation into an optimization problem; we then propose a modification of the SAPG by allowing a random choice of the stepsize sequence $\{\gamma_n, n \geq 0\}$, to improve the numerical properties of the algorithm. We conclude the section by a comparison of the methods on a pharmacokinetic real data set.

5.1 The nonlinear mixed effect model

Pharmacokinetic data are observed along time for N patients. Let Y_k be the vector of the J drug concentrations observed at time t_{kj} ($j \in \{1, \dots, J\}$) for the k th patient ($k \in \{1, \dots, N\}$). The kinetic of the drug concentration is described by a nonlinear pharmacokinetic regression model f , which is a function of time t and unobserved pharmacokinetic parameters $Z^{(k)}$. These parameters are typically the rates of absorption or elimination of the drug by the body. An example is detailed below. The variability among patients is modeled by the randomness of the hidden variables $Z^{(k)}$. These pharmacokinetic parameters may be influenced by covariates, such as age, gender but also genomic variables. Among these high dimension factors, only few of them are correlated with $Z^{(k)}$. Their selection can thus be performed by optimizing the likelihood with a sparsity inducing penalty, an optimization problem that enters problem Eq. (1). However, the likelihood is generally not concave, that is, through this example, we explore beyond the framework in which we are able to prove the convergence of MCPG and SAPG (see Sect. 3).

Let us now detail the model and the optimization problem. The mixed model is defined as

$$Y_{kj} = f(t_{kj}, Z^{(k)}) + \epsilon_{kj}, \quad \epsilon_{kj} \sim \mathcal{N}(0, \sigma^2) \text{ (iid)}, \quad (28)$$

where the measurement errors ϵ_{kj} are centered, independent and identically normally distributed with variance σ^2 . Indi-

vidual parameters $Z^{(k)}$ for the k th subject is a R -dimensional random vector, independent of ϵ_{kj} . In a high dimension context, the $Z^{(k)}$'s depend on covariates (typically genomics variables) gathered in a matrix design $X_k \in \mathbb{R}^{R \times (D+1)R}$. The distribution of $Z^{(k)}$ is usually assumed to be normal with independent components

$$Z^{(k)} \stackrel{\text{ind}}{\sim} \mathcal{N}_R(X_k \mu, \Omega) \quad (29)$$

where $\mu \in \mathbb{R}^{(D+1)R}$ is the mean parameter vector and Ω is the covariance matrix of the random parameters $Z^{(k)}$, assumed to be diagonal. The unknown parameters are $\theta = (\mu, \Omega_{11}, \dots, \Omega_{RR}, \sigma^2) \in \mathbb{R}^{R(D+1)} \times (0, +\infty)^{R+1}$.

A typical function f is the two-compartmental pharmacokinetic model with first-order absorption, describing the distribution of a drug administered orally. The drug is absorbed from the gut and reaches the blood circulation where it can spread in peripheral tissues. This model corresponds to $f = \frac{A_c}{V_c}$ with A_c defined as

$$\begin{aligned} \frac{dA_d}{dt} &= -k_a A_d, \\ \frac{dA_c}{dt} &= k_a A_d + \frac{Q}{V_p} A_p - \frac{Q}{V_c} A_c - \frac{Cl}{V_c} A_c, \\ \frac{dA_p}{dt} &= \frac{Q}{V_c} A_c - \frac{Q}{V_p} A_p, \end{aligned} \quad (30)$$

with $A_d(0) = \text{Dose}$, $A_c(0) = 0$, $A_p(0) = 0$ and where A_d, A_c, A_p are the amount of drug in the depot, central and peripheral compartments, respectively; V_c, V_p are the volume of the central compartment and the peripheral compartment, respectively; Q and Cl are the intercompartment and global elimination clearances, respectively. To assure positiveness of the parameters, the hidden vector is

$$z = (\log(V_c), \log(V_p), \log(Q), \log(Cl), \log(k_a)).$$

It is easy to show that the model described by Eqs. (28)–(29) belongs to the curved exponential family (see Eq. 4) with minimal sufficient statistics:

$$\begin{aligned} S_{1k}(z) &= z^{(k)}, \quad S_2(z) = \sum_{k=1}^N z^{(k)} z^{(k)'}, \\ S_3(z) &= \sum_{k=1}^N \sum_{j=1}^J (Y_{kj} - f(t_{kj}, z^{(k)}))^2; \end{aligned}$$

$$\begin{aligned} \psi_{1k}(\theta) &= (X_k \mu)' \Omega^{-1}, \quad \psi_2(\theta) = -\frac{1}{2} \Omega^{-1}, \\ \psi_3(\theta) &= -\frac{1}{2\sigma^2}, \end{aligned}$$

and $S(z) := \text{Vect}(S_{11}(z), \dots, S_{1N}(z), S_2(z), S_3(z))$, $\psi := \text{Vect}(\psi_{11}, \dots, \psi_{1N}, \psi_2, \psi_3)$. The function ϕ is given by $\phi(\theta) = -JN \log(\sigma) - \frac{N}{2} \log(|\Omega|) - \frac{1}{2} \sum_k (X_k \mu)' \Omega^{-1} (X_k \mu)$. The selection of genomic variables that influence all coordinates of $Z^{(k)}$ could be obtained by optimizing the log likelihood penalized by the function $g(\theta) = \lambda \|\mu\|_1$, the L_1 norm of μ with λ a regularization parameter.

However, this estimator is not invariant under a scaling transformation (i.e., $\tilde{Z}^{(k)} = bZ^{(k)}$, $\tilde{\mu} = b\mu$ and $\tilde{\Omega}_{rr}^{1/2} = b\Omega_{rr}^{1/2}$) (see, e.g., Lehmann and Casella 2006). In our high dimension experiments, the scale of the hidden variables has a non-negligible influence on the selection of the support. To be more precise, let us denote, for $r \in \{1, \dots, R\}$,

$$\mu_{(r)} := (\mu_{(r-1)(D+1)+1}, \dots, \mu_{r(D+1)})$$

the coordinates corresponding to the r th pharmacokinetic parameter of function f . When the variance Ω_{rr} of the random parameters $Z_r^{(k)}$ is low, the algorithms tend to select too many covariates. This phenomenon is strengthened with a small number of subjects as random effect variances are more difficult to estimate. A solution is to consider the following penalty

$$\lambda \sum_{r=1}^R \Omega_{rr}^{-\frac{1}{2}} \|\mu_{(r)}\|_1,$$

that makes the estimator invariant under scaling transformation. It was initially proposed by Städler et al. (2010) to estimate the regression coefficients and the residual error's variance in a mixture of penalized regression models. However, the resulting optimization problem is difficult to solve directly because the variance of the random effect Ω_{rr} appears in the penalty term. Therefore, we propose a new parameterization

$$\tilde{\mu}_{(r)} := \mu_{(r)} \Omega_{rr}^{-\frac{1}{2}}, \quad \Sigma_{rr} := \Omega_{rr}^{-\frac{1}{2}}$$

and $\tilde{\theta} := \{\tilde{\mu}, \Sigma_{11}, \dots, \Sigma_{RR}, \sigma^2\} \in \mathbb{R}^{R(D+1)} \times (0, +\infty)^{R+1}$. Then, the optimization problem is the following:

$$\underset{\tilde{\theta}}{\text{Argmax}} \left(\ell(\tilde{\theta}) - g(\tilde{\theta}) \right), \quad \text{with } g(\tilde{\theta}) = \lambda \|\tilde{\mu}\|_1. \tag{31}$$

This problem can be solved using MCPG, SAPG or SAEM-pen algorithms. Indeed, the complete log likelihood is now—up to an additive constant—

$$\begin{aligned} \log p(Y, Z; \tilde{\theta}) &= -JN \log(\sigma) \\ &\quad - \frac{1}{2} \sum_{k=1}^N \sum_{j=1}^J \frac{(Y_{kj} - f(t_{kj}, Z^{(k)}))^2}{\sigma^2} \\ &\quad + N \log(|\Sigma|) - \frac{1}{2} \sum_{k=1}^N \|\Sigma Z^{(k)} - X_k \tilde{\mu}\|^2 \end{aligned}$$

It is again a complete likelihood from the exponential family, with the statistic S unchanged and the functions ϕ and ψ given by—up to an additive constant—

$$\begin{aligned} \phi(\tilde{\theta}) &= -JN \log(\sigma) + N \log(|\Sigma|) - \frac{1}{2} \sum_{k=1}^N \|X_k \tilde{\mu}\|^2, \\ \psi_{1k}(\tilde{\theta}) &= \Sigma (X_k \tilde{\mu})^t, \quad \psi_2(\tilde{\theta}) = -\frac{1}{2} \Sigma^2, \quad \psi_3(\tilde{\theta}) = -\frac{1}{2\sigma^2}. \end{aligned}$$

With these definitions of ϕ , ψ and g , the M-step of SAEM-pen amounts to compute the optimum of a convex function, which is solved numerically by a call to a cyclical coordinate descent implemented in the R package glmnet (Friedman et al. 2010).

MCMC sampler In the context of nonlinear mixed models, simulation from $\pi_{\theta_n} dv$ can not be performed directly like in the toy example. We then use a MCMC sampler based on a Metropolis Hastings algorithm to perform the simulation step. Two proposal kernels are successively used during the iterations of the Metropolis Hastings algorithm. The first kernel corresponds to the prior distribution of $\Sigma Z^{(k)}$ that is the Gaussian distribution $\mathcal{N}(X_k \tilde{\mu}_n, I)$. The second kernel corresponds to a succession of R uni-dimensional random walk in order to update successively each component of $Z^{(k)}$. The variance of each random walk is automatically tuned to reach a target acceptance ratio following the principle of an adaptive MCMC algorithm (Andrieu and Thoms 2008).

Adaptive random stepsize sequences In the context of NLMEM, numerical experiments reveal that choosing a deterministic sequence $\{\gamma_n, n \geq 0\}$ that achieve a fast convergence of SAPG algorithm could be difficult. Indeed, parameters to estimate are of different scales. For example, random effect and residual variances are constrained to be positive. Some of them are close to zero; some are not. As explained in Sect. 4.2, an alternative is to implement a matrix-valued random sequence $\{\Gamma^n, n \geq 0\}$. The gradient and the hessian of the likelihood $\ell(\theta)$ can be approximated by stochastic approximation using the Louis principle (see McLachlan and Krishnan 2008, Chapter 4). Let us denote H_n the stochastic approximation of the hessian obtained at iteration n as explained by Samson et al. (2007). Note that no supplementary random samples are required to obtain this approximation. Along the iterations, each diagonal entry of

the matrix H^n converges: this limiting value can be seen as a simple way to automatically tune a good γ_* , that is parameter specific. The entries Γ_{ii}^{n+1} are then defined by Eq. (27).

5.2 Simulated data set

The convergence of the corresponding algorithms is illustrated on simulated data. Data are generated with the model defined by Eq. (30) and $N = 40, J = 12, D = 300$. The design matrix X_k is defined by Eq. (20), with components (X_{k1}, \dots, X_{kD}) drawn from $\mathcal{N}(0, \Gamma)$ with $\Gamma_{ii'} = 0.5^{|i-i'|}$ ($i, i' = 1, \dots, 300$). Parameter values are

$$[\mu_1, \mu_{1+(D+1)}, \mu_{1+2(D+1)}, \mu_{1+3(D+1)}, \mu_{1+4(D+1)}] = [6.61, 6.96, 5.77, 5.42, -0.51];$$

the other components are set to zero, except μ_4 and μ_{912} that are set to 1. The matrix Ω is diagonal with diagonal elements equal to (0.16, 0.16, 0.16, 0.04, 0.04).

The penalty function is set to

$$g(\tilde{\theta}) := \lambda \sum_{\ell \neq \{1+r(D+1), r=0, \dots, 4\}} |\tilde{\mu}_\ell|, \tag{32}$$

only the parameters corresponding to a covariate effect being penalized. The optimization problem Eq. (1) with regularization parameter $\lambda = 190$ is solved on this dataset with SAEM-pen and SAPG; we run SAPG with the random sequence $\{\Gamma_n, n \geq 0\}$ as described above (see Eq. 27) with $n_0 = 9500$. For both algorithms, the stochastic approximation step size was set to:

$$\delta_{n+1} = \begin{cases} 0.5 & \text{if } n \leq n_0 \\ \frac{0.5}{(n-n_0)^\beta} & \text{if } n > n_0 \end{cases} \tag{33}$$

We set $\alpha = 0.75$ and $\beta = 0.499$. Figure 8 shows the convergence of SAEM-pen and three parameterizations of SAPG: (i) a version with $\gamma^* = 0.005$ for all the components of θ , (ii) a version with $\gamma^* = 0.005$ for $\tilde{\mu}$, $\gamma^* = 0.0005$ for Σ and $\gamma^* = 0.03$ for σ and (iii) a version with adaptive random step sizes. For the four algorithms, all the parameters corresponding to a covariate effect are estimated to zero except the two components μ_4 and μ_{912} . The version of SAPG with a same γ^* for all the component is the one that converge the most slowly. When the γ^* is tuned differently according the type of parameters, the convergence of SAPG is accelerated. Algorithms SAEM-pen and SAPG with adaptive random step sizes have a similar fast convergence profile.

Figure 9 presents the evolution of four entries of the matrix Γ^n along the iterations of SAPG, corresponding to the components $\tilde{\mu}_{904}, \tilde{\mu}_{912}, \Sigma_{44}$ and σ . We can notice that they are not on the same scale. They vary during the first iterations and converge to limiting values before iteration $n_0 = 9500$.

Then the step sizes decrease to 0, following the definition given in Eq. (27).

5.3 Application to real data

Algorithms SAEM-pen and SAPG with matrix-valued random sequence $\{\Gamma^n, n \geq 0\}$ are applied to real data of the pharmacokinetic of dabigatran (*DE*) from two crossover clinical trials (Delavenne et al. 2013; Ollier et al. 2015). These 2 trials studied the drug–drug interaction between *DE* and different Pgp-inhibitors. From these 2 trials, the pharmacokinetics of *DE* are extracted from 15 subjects with no concomitant treatment with Pgp-inhibitors. The concentration of dabigatran is measured at 9 sampling times for each patient. Each subject is genotyped using the DMET[®] microarray from Affymetrix. Single nucleotide polymorphisms (SNP) showing no variability between subjects are removed and 264 SNP are included in the analysis.

Function f of the nonlinear mixed model is defined as the two-compartment pharmacokinetic model with first-order absorption previously described (see Eq. 30) (Delavenne et al. 2013). The penalty function g is defined by Eq. (32).

Because of the limited number of subjects, the influence of genetic covariates is only studied on V_c and Cl parameters, that characterize the elimination process and are the most likely to be influenced by the genetic. Finally, random effect variances of Q and V_p are set to 0.01 in accordance with previously published population pharmacokinetic of dabigatran (Delavenne et al. 2013). The other variance parameters are estimated. The penalized likelihood problem (Eq. 31) is solved on the data with the SAEM-pen and SAPG algorithms, for 40 different values of parameter λ . SAPG algorithm is run using the random sequence $\{\Gamma^n, n \geq 0\}$ given in Eq. (27). The best regularization parameter λ is chosen with a data-driven approach based on the EBIC criteria (Chen and Chen 2008).

Figure 10 shows the results. The regularization paths of Cl and V_c parameters using both algorithms correspond to the evolution of covariate coefficient estimates as a function of the value of λ . They are reconstructed with low noise for both algorithms, are very similar for high values of λ but less for lower values of λ .

Finally, the selected model has all covariates parameters set to zero. This means that none of the genetic covariates influence the distribution of the individual parameters. This result is not surprising given the low number of subjects and the fact that a large part of the interindividual variability is due to the dissolution process of the drug (Ollier et al. 2015) and is therefore not influenced by genetic covariates. This lack of relationship between dabigatran’s pharmacokinetic parameters and genetic covariates has already been highlighted in an other study (Gouin-Thibault et al. 2017).

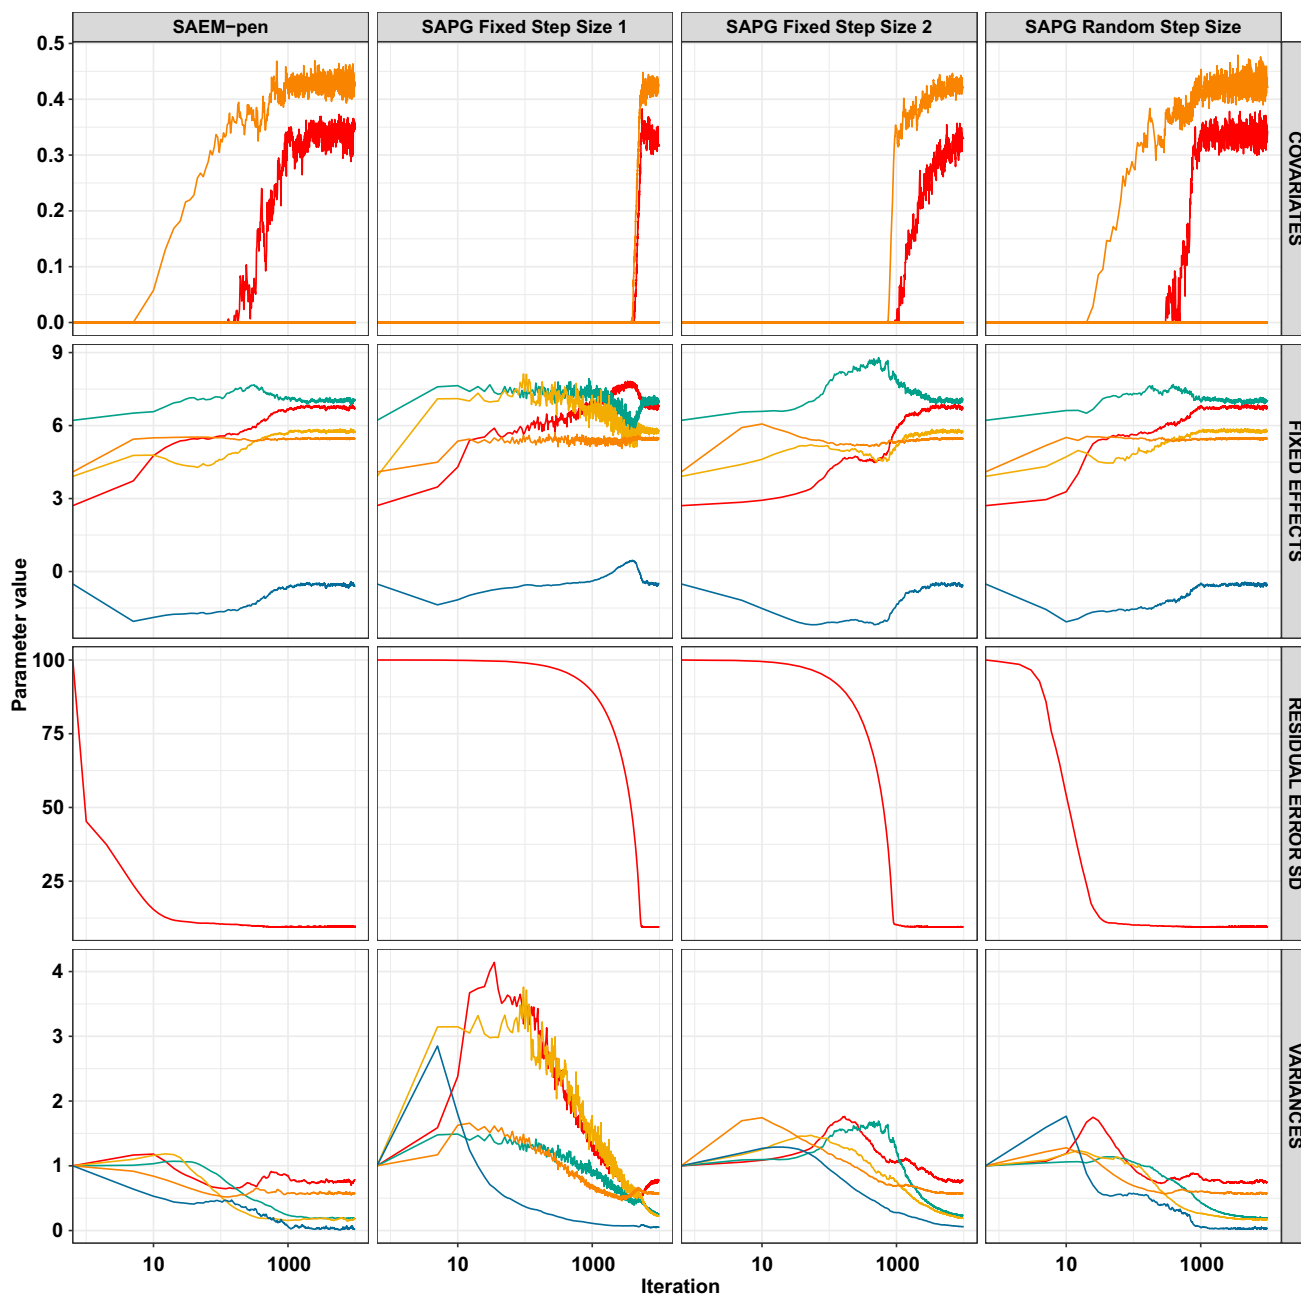


Fig. 8 Path of a run of SAEM-pen [left column] and three different parameterizations of SAPG: (i) with $\gamma^* = 0.005$ for all the components of θ [middle left column], (ii) with $\gamma^* = 0.005$ for $\tilde{\mu}$, $\gamma^* = 0.0005$ for Σ and $\gamma^* = 0.03$ for σ [middle right column] and (iii) with a random sequence $\{\Gamma^n, n \geq 0\}$ [right column]. For each algorithm, estimation of the standard deviation of the residual error σ [third row]; the variances of the $Z^{(k)}$'s, $\Omega_{11}, \dots, \Omega_{RR}$ [fourth row]; the path of the covariate param-

eters μ_i for $i \notin \{1, 1 + (D + 1), \dots, 1 + 4(D + 1)\}$ [first row]; the path of the intercept parameters $\mu_i, i \in \{1, 1 + (D + 1), \dots, 1 + 4(D + 1)\}$ [second row]. Each color corresponds to a specific parameter: orange line for Cl , red for V_c , blue line for k_a , yellow line for Q and green line for V_p . Note that the path of all the covariate parameters is zero except for two components. x -axis is in \log_{10} scale. (Color figure online)

6 Conclusion

In this work, we propose a new stochastic proximal-gradient algorithm to solve penalized maximum-likelihood problems when the likelihood is intractable: the gradient is approx-

imated through a stochastic approximation scheme. We provide a theoretical convergence analysis of this new algorithm and illustrate these results numerically on a simulated toy example in the case of a concave likelihood function. The robustness to the non-concave case is explored through a




Fig. 9 Evolution of Γ_{ii}^n with iterations n of SAPG, for four different values of i , corresponding to the components $\tilde{\mu}_{904}$ [left]; $\tilde{\mu}_{912}$ [middle left]; Σ_{44} [middle right]; σ [right]. Both x -axis and y -axis are in \log_{10} scale




Fig. 10 Regularization path of covariate parameters (Cl parameter on top, V_c parameter on bottom) obtained on dabigatran pharmacokinetic data for both SAEM-pen and SAPG algorithms. Black vertical dashed line corresponds to the λ value selected by $EBIC$. (Color figure online)

more challenging application to population pharmacokinetic analysis relying on penalized inference in nonlinear mixed effects models.

Appendix

A Proof of Proposition 1

Lemma 1 Under the assumptions of Proposition 1, for any $\gamma \in (0, 1/L]$, $s \in \bar{S}(\Theta)$ and any $\theta, \theta' \in \Theta$,

$$\mathcal{Q}(\theta|\theta') \geq \mathcal{Q}(\theta'|\theta')$$

$$\begin{aligned}
 & - \frac{1}{2\gamma} \|\theta - \theta' - \gamma\{\nabla\phi(\theta') + J\psi(\theta')\bar{S}(\theta')\}\|^2 \\
 & + \frac{\gamma}{2} \|\nabla\phi(\theta') + J\psi(\theta')\bar{S}(\theta')\|^2. \tag{34}
 \end{aligned}$$

Proof Fix $\theta' \in \Theta$ and $s \in \bar{S}(\Theta)$. The derivative of the function $\theta \mapsto \mathcal{L}(\theta) := \phi(\theta) + \langle s, \psi(\theta) \rangle$ is $\nabla\phi(\theta) + J\psi(\theta)s$, and this gradient is L -Lipschitz. From a Taylor expansion to order 1 at θ' and since the gradient is Lipschitz, we have

$$\mathcal{L}(\theta) \geq \mathcal{L}(\theta') + \langle \nabla\phi(\theta') + J\psi(\theta')s, \theta - \theta' \rangle - \frac{L}{2} \|\theta - \theta'\|^2.$$

We then choose $s = \bar{S}(\theta')$, use $L \leq 1/\gamma$ and conclude by the equality $2 \langle a, b \rangle - \|a\|^2 = \|b\|^2 - \|a - b\|^2$. \square

Proof of Proposition 1 We prove that $\mathcal{Q}(\theta_{n+1}|\theta_n) - g(\theta_{n+1}) \geq \mathcal{Q}(\theta_n|\theta_n) - g(\theta_n)$ so that the sequence $\{\theta_n, n \geq 0\}$ defined by Eq. (10) is a sequence satisfying Eq. (15).

By Lemma 1, it holds for any $\theta \in \Theta$ and any $\gamma \in (0, 1/L]$

$$\begin{aligned}
 & \mathcal{Q}(\theta|\theta_n) - g(\theta) \\
 & \geq \mathcal{Q}(\theta_n|\theta_n) + \frac{\gamma}{2} \|\nabla\phi(\theta_n) + J\psi(\theta_n)\bar{S}(\theta_n)\|^2 \\
 & \quad - \frac{1}{2\gamma} \|\theta - \theta_n - \gamma\{\nabla\phi(\theta_n) + J\psi(\theta_n)\bar{S}(\theta_n)\}\|^2 - g(\theta).
 \end{aligned}$$

Note that the RHS and the LHS are equal when $\theta = \theta_n$ so that for any point τ which maximizes the RHS, it holds $\mathcal{Q}(\tau|\theta_n) - g(\tau) \geq \mathcal{Q}(\theta_n|\theta_n) - g(\theta_n)$. This concludes the proof upon noting that such a point τ is unique and equal to θ_{n+1} given by Eq. (10).

B Technical lemmas

Define

$$\Delta_{k:n} := \prod_{j=k}^n (1 - \delta_j), \quad 0 \leq k \leq n, \quad \Delta_{n+1:n} = 1,$$

$$D_k := \sum_{n \geq k} \Delta_{k:n}.$$

Lemma 2 For any $n \geq 2$, $\sum_{j=2}^n \Delta_{j+1:n} \delta_j = 1 - \Delta_{2:n}$.

Proof For any $j \leq n$, we have $\Delta_{j+1:n} - \Delta_{j:n} = \delta_j \Delta_{j+1:n}$ from which the result follows.

Lemma 3 Let $\beta \in (0, 1)$ and $\delta_\star > 0$. Set $\delta_n = \delta_\star n^{-\beta}$ for any $n \geq 1$. Then for any k large enough,

$$\delta_k D_k \leq 1 + O(k^{\beta-1}).$$

Furthermore, $|\delta_{n+1} D_{n+2} - \delta_n D_{n+1}| = O(1/n^{1+(1-\beta)\wedge\beta})$.

The proof of Lemma 3 relies on standard Taylor’s expansions with explicit formulas for the remainder. The proof is omitted.

Lemma 4 Let $\beta \in [0, 1)$ and $\delta_\star > 0$. For any r , when $n \rightarrow \infty$,

$$\sum_{j=2}^n j^{-r} \prod_{k=j}^n \left(1 - \frac{\delta_\star}{k^\beta}\right) = O(n^{\beta-r}).$$

Proof We have

$$\begin{aligned} \prod_{k=j}^n \left(1 - \frac{\delta_\star}{k^\beta}\right) &\leq \exp\left(-\delta_\star \sum_{k=j}^n k^{-\beta}\right) \\ &\leq \exp\left(-\frac{\delta_\star}{1-\beta} \left\{n^{1-\beta} - j^{1-\beta}\right\}\right). \end{aligned}$$

Let $q_\star \geq 0$ such that for any $q \geq q_\star$, $q(1-\beta) + 1 - r > 0$. For any constant $D > 0$, there exist constants C, C' (whose value can change upon each appearance) such that

$$\begin{aligned} &\sum_{j=2}^n j^{-r} \exp(Dj^{1-\beta}) \\ &= \sum_{q \geq 0} \frac{D^q}{q!} \sum_{j=2}^n j^{-r+q(1-\beta)} \\ &\leq Cn^{1-r} \sum_{q \geq q_\star} \frac{D^q}{q!} \frac{n^{q(1-\beta)}}{(q+1)(1-\beta)} \frac{(q+1)(1-\beta)}{q(1-\beta)+1-r} + C'n \\ &\leq Cn^{\beta-r} \sum_{q \geq q_\star} \frac{D^q}{(q+1)!} n^{(q+1)(1-\beta)} + C'n \\ &\leq C \exp(Dn^{1-\beta}) n^{\beta-r}. \end{aligned}$$

This concludes the proof. □

Lemma 5 Let $\{A_n, n \geq 0\}$ be a sequence of $d' \times q$ matrices and $\{\sigma_n, n \geq 0\}$ be a sequence of $q \times 1$ vectors. Let $\{S_n^{\text{sa}}, n \geq 0\}$ be given by Eq. (12). For any $n \geq 2$

$$\begin{aligned} A_n (S_n^{\text{sa}} - \sigma_{n-1}) &= \Delta_{2:n} A_1 (S_1^{\text{sa}} - \sigma_0) \\ &+ \sum_{j=2}^n \Delta_{j:n} (A_j - A_{j-1}) (S_{j-1}^{\text{sa}} - \sigma_{j-2}) \\ &+ \sum_{j=2}^n \Delta_{j:n} A_j (\sigma_{j-2} - \sigma_{j-1}) \\ &+ \sum_{j=2}^n \Delta_{j+1:n} \delta_j A_j \left(m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) - \sigma_{j-1}\right). \end{aligned}$$

Proof By definition of S_n^{sa} , it holds $A_n (S_n^{\text{sa}} - \sigma_{n-1}) = (1 - \delta_n) A_{n-1} (S_{n-1}^{\text{sa}} - \sigma_{n-2}) + B_n$ where

$$\begin{aligned} B_n &:= (1 - \delta_n)(A_n - A_{n-1}) (S_{n-1}^{\text{sa}} - \sigma_{n-2}) \\ &+ (1 - \delta_n) A_n (\sigma_{n-2} - \sigma_{n-1}) \\ &+ \delta_n A_n \left(m_n^{-1} \sum_{k=1}^{m_n} S(Z_{k,n-1}) - \sigma_{n-1}\right). \end{aligned}$$

By iterating, we have

$$A_n (S_n^{\text{sa}} - \sigma_{n-1}) = \Delta_{2:n} A_1 (S_1^{\text{sa}} - \sigma_0) + \sum_{j=2}^n \Delta_{j+1:n} B_j,$$

from which the lemma follows. □

Lemma 6 Assume (HH4a). Let $\{S_n^{\text{sa}}, n \geq 0\}$ be given by Eq. (12). Then

$$\begin{aligned} \sup_{n \geq 0} \mathbb{E} \left[\left\| m_n^{-1} \sum_{j=1}^{m_n} S(Z_{j,n-1}) \right\|^2 \right] &< \infty, \\ \sup_{n \geq 0} \mathbb{E} \left[m_n^{-1} \sum_{j=1}^{m_n} W(Z_{j,n-1}) \right] &< \infty, \\ \sup_{n \geq 0} \mathbb{E} \left[\|S_n^{\text{sa}}\|^2 \right] &< \infty. \end{aligned}$$

Proof By (H4a), there exists a constant $C < \infty$ such that for any $n \geq 1$ and $1 \leq j \leq m_n$, $\|S(Z_{j,n-1})\|^2 \leq C W(Z_{j,n-1})$. In addition, by the drift assumption on the kernels P_θ , we have

$$\begin{aligned} \mathbb{E} [W(Z_{j,n-1})] &= \mathbb{E} [P_{\theta_{n-1}} W(Z_{j-1,n-1})] \\ &\leq \lambda \mathbb{E} [W(Z_{j-1,n-1})] + b \\ &\leq \lambda^j \mathbb{E} [W(Z_{0,n-1})] + \frac{b}{1-\lambda}. \end{aligned}$$

Similarly, by using $Z_{0,n-1} = Z_{m_{n-1},n-2}$, we have

$$\mathbb{E} [W(Z_{0,n-1})] \leq \lambda^{m_{n-1}} \mathbb{E} [W(Z_{0,n-2})] + \frac{b}{1-\lambda}.$$

A trivial induction shows that

$$\sup_n \sup_{j \leq m_n} \mathbb{E} [W(Z_{j,n-1})] < \infty,$$

from which the first two results follow. For the third one, by Lemma 5 applied with $A_n = I$ (the identity matrix) and $\sigma_n = 0$, we have for any $n \geq 1$,

$$S_n^{sa} = \Delta_{2:n} S_1 + \sum_{j=2}^n \Delta_{j+1:n} \delta_j m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}).$$

By the Minkowsky inequality and the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we have

$$\begin{aligned} \mathbb{E} [\|S_n^{sa}\|^2] &\leq 2 (\Delta_{2:n})^2 \mathbb{E} [\|S_1^{sa}\|^2] \\ &\quad + 2 \sup_j \mathbb{E} \left[\left\| m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) \right\|^2 \right] \\ &\quad \times \left(\sum_{j=2}^n \Delta_{j+1:n} \delta_j \right)^2. \end{aligned}$$

By definition, $\Delta_{2:n} \in [0, 1]$ and by Lemma 2,

$$\sup_n \sum_{j=2}^n \Delta_{j+1:n} \delta_j < \infty.$$

Hence, $\sup_{n \geq 0} \mathbb{E} [\|S_n^{sa}\|^2] < \infty$. □

Define the proximal-gradient operator

$$T_\gamma(\theta) := \text{Prox}_{\gamma,g}(\theta + \gamma \nabla \ell(\theta)).$$

Lemma 7 Assume H1, H2 and H3. Let $\{S_n^{sa}, n \geq 0\}$ be given by (12). Then, for the sequence $\{\theta_n, n \geq 0\}$ given by Algorithm 2,

(i) There exists a constant C such that almost-surely, for any $n \geq 0$,

$$\|\theta_{n+1} - \theta_n\| \leq C \gamma_{n+1} (1 + \|S_{n+1}^{sa} - \bar{S}(\theta_n)\|).$$

(ii) There exists a constant C' such that almost-surely, for any $n \geq 0$,

$$\begin{aligned} &\|\gamma_{n+1} \Psi(\theta_n) - \gamma_n \Psi(\theta_{n-1})\| \\ &\leq C' (|\gamma_{n+1} - \gamma_n| + \gamma_n^2 (1 + \|S_n^{sa} - \bar{S}(\theta_{n-1})\|)). \end{aligned}$$

(iii) There exists a constant C'' such that almost-surely, for any $n \geq 0$,

$$\begin{aligned} &\|\gamma_{n+1} T_{\gamma_{n+1},g}(\theta_n) - \gamma_n T_{\gamma_n,g}(\theta_{n-1})\| \\ &\leq C'' (|\gamma_{n+1} - \gamma_n| + \gamma_n \gamma_{n+1} \\ &\quad + \gamma_n^2 (1 + \|S_n^{sa} - \bar{S}(\theta_{n-1})\|)). \end{aligned}$$

Proof The proof of (i) is on the same lines as the proof of Atchadé et al. (2017, Lemma 15), and is omitted. For (ii), we write by using (H3b) and (H3c),

$$\begin{aligned} &\|\gamma_{n+1} \Psi(\theta_n) - \gamma_n \Psi(\theta_{n-1})\| \\ &\leq |\gamma_{n+1} - \gamma_n| \|\Psi(\theta_n)\| + \gamma_n \|\Psi(\theta_n) - \Psi(\theta_{n-1})\| \\ &\leq |\gamma_{n+1} - \gamma_n| \sup_{\Theta} \|\Psi\| + L \gamma_n \|\theta_n - \theta_{n-1}\|. \end{aligned}$$

We then conclude by (ii). The LHS in (iii) is upper bounded by

$$\begin{aligned} &|\gamma_{n+1} - \gamma_n| \sup_{\gamma \in (0,1/L]} \sup_{\theta \in \Theta} \|T_{\gamma,g}(\theta)\| \\ &\quad + \gamma_n \|T_{\gamma_{n+1},g}(\theta_n) - T_{\gamma_n,g}(\theta_{n-1})\|. \end{aligned}$$

Under H4a, there exists a constant C such that for all $\gamma, \gamma' \in (0, 1/L]$ and $\theta, \theta' \in \Theta$ (see Atchadé et al. 2017, Proposition 12)

$$\begin{aligned} &\sup_{\gamma \in (0,1/L]} \sup_{\theta \in \Theta} \|T_{\gamma,g}(\theta)\| < \infty, \\ &\|T_{\gamma,g}(\theta) - T_{\gamma',g}(\theta')\| \leq C (\gamma + \gamma' + \|\theta - \theta'\|). \end{aligned}$$

We then conclude by (i). □

Lemma 8 Assume H4. For any $\theta \in \Theta$, there exists a function $\hat{S}_\theta : Z \rightarrow \mathbb{R}^q$ such that $S - \bar{S}(\theta) = \hat{S}_\theta - P_\theta \hat{S}_\theta$ and $\sup_{\theta \in \Theta} \|\hat{S}_\theta\|_{\sqrt{W}} < \infty$. In addition, there exists a constant C such that for any $\theta, \theta' \in \Theta$,

$$\|P_\theta \hat{S}_\theta - P_{\theta'} \hat{S}_{\theta'}\|_{\sqrt{W}} \leq C \|\theta - \theta'\|.$$

Proof Set $\hat{S}_\theta(z) := \sum_{n \geq 0} (P_\theta^n S(z) - \bar{S}(\theta))$. Observe that, when exists, this function satisfies $S - \bar{S}(\theta) = \hat{S}_\theta - P_\theta \hat{S}_\theta$. Note that under (H4a)–(H4b), there exist C and $\rho \in (0, 1)$ such that for any $\theta \in \Theta$,

$$\sum_{n \geq 0} \|P_\theta^n S(z) - \bar{S}(\theta)\| \leq C |S|_{\sqrt{W}} \left(\sum_{n \geq 0} \rho^n \right) \sqrt{W}(z);$$

the RHS is finite, thus showing that \hat{S}_θ exists. This inequality also proves that $\sup_{\theta \in \Theta} \|\hat{S}_\theta\|_{\sqrt{W}} < \infty$. The Lipschitz property

is established in Fort et al. (2011b, Lemma 4.2) and its proof uses (H4c). \square

C Proof of Proposition 2

Throughout this section, set $\|U\|_{L_2} := \mathbb{E}[\|U\|^2]^{1/2}$. By Lemma 5, $\|S_n^{\text{sa}} - \bar{S}(\theta_{n-1})\|_{L_2} \leq \sum_{i=1}^3 \mathcal{T}_{i,n}$ with

$$\begin{aligned} \mathcal{T}_{1,n} &:= \Delta_{2:n} \|S_1^{\text{sa}} - \bar{S}(\theta_0)\|_{L_2}, \\ \mathcal{T}_{2,n} &:= \sum_{j=2}^n \Delta_{j:n} \|\bar{S}(\theta_{j-1}) - \bar{S}(\theta_{j-2})\|_{L_2}, \\ \mathcal{T}_{3,n} &:= \left\| \sum_{j=2}^n \Delta_{j+1:n} \delta_j \left(m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) \right) \right\|_{L_2}. \end{aligned}$$

Since $\Delta_{2:n} \leq \exp(-\delta_\star \sum_{j=2}^n j^{-\beta})$, then

$$\mathcal{T}_{1,n} = O\left(\exp(-\delta_\star (1 - \beta)^{-1} n^{1-\beta})\right).$$

By (H3c), Lemma 6 and Lemma 7, there exists a constant C such that $\mathcal{T}_{2,n} \leq C \sum_{j=2}^n \Delta_{j:n} \gamma_{j-1}$. By Lemma 4, this yields $\mathcal{T}_{2,n} = O(n^{\beta-\alpha})$. For the last term, we use a martingale decomposition.

By Lemma 8, there exists a function \hat{S}_θ such that

$$S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) = \hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k,j-1}),$$

and $\sup_{\theta \in \Theta} \left| \hat{S}_\theta \right|_{\sqrt{W}} < \infty$. Hence, we write

$$m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) = \partial M_j + R_{j,1} + R_{j,2}$$

with

$$\begin{aligned} \partial M_j &:= m_j^{-1} \sum_{k=1}^{m_j} \left\{ \hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k-1,j-1}) \right\}, \\ R_{j,1} &:= m_j^{-1} \left\{ P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{0,j-1}) - P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}) \right\}, \\ R_{j,2} &:= m_j^{-1} \left\{ P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{0,j}) \right\}; \end{aligned}$$

we used that $Z_{0,j} = Z_{m_j,j-1}$. Upon noting that ∂M_j is a martingale increment, and

$$\hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k-1,j-1})$$

is a martingale increment, we have by two successive applications of Hall and Heyde (1980, Theorem 2.10):

$$\left\| \sum_{j=2}^n \Delta_{j+1:n} \delta_j \partial M_j \right\|_{L_2} \leq C \left(\sum_{j=2}^n \Delta_{j+1:n}^2 \frac{\delta_j^2}{m_j} \right)^{1/2}.$$

By Lemma 4, this term is $O(n^{-(\beta+c)/2})$. For the second term, we write

$$\begin{aligned} &\sum_{j=2}^n \Delta_{j+1:n} \delta_j R_{j,1} \\ &= \Delta_{3:n} \frac{\delta_2}{m_2} P_{\theta_1} \hat{S}_{\theta_1}(Z_{0,1}) - \frac{\delta_n}{m_n} P_{\theta_n} \hat{S}_{\theta_n}(Z_{0,n}) \\ &\quad + \sum_{j=2}^{n-1} \left(\Delta_{j+2:n} \frac{\delta_{j+1}}{m_{j+1}} - \Delta_{j+1:n} \frac{\delta_j}{m_j} \right) P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}). \end{aligned}$$

By Lemmas 6 and 8, the RHS is $O(n^{-(\beta+c)} + n^{-(1+c)})$ so that this second term is $O(n^{-(\beta+c)})$. Finally, for the third term, by using Lemmas 6, 7 and 8, we write

$$\left\| \sum_{j=2}^n \Delta_{j+1:n} \delta_j R_{j,2} \right\|_{L_2} \leq \sum_{j=2}^n \Delta_{j+1:n} \frac{\delta_j}{m_j} \gamma_j.$$

Again by Lemma 4, this last term is $O(n^{-(\alpha+c)})$. Therefore, $\mathcal{T}_{3,n} = O(n^{-(\beta+c)/2})$.

D Proof of Theorem 4

Throughout the proof, we will write S_{n+1} instead of S_{n+1}^{sa} .

Proof of Theorem 4 We prove the almost-sure convergence of the three random sums given in Theorem 1. The third one is finite almost-surely since its expectation is finite (see Proposition 3). The first two ones are of the form $\sum_n A_{n+1} (S_{n+1} - \bar{S}(\theta_n))$ where A_{n+1} is, respectively,

$$A_{n+1} = \gamma_{n+1} (T_{\gamma_{n+1},g}(\theta_n))', \quad A_{n+1} = \gamma_{n+1} \Psi(\theta_n).$$

Note that $A_{n+1} \in \mathcal{F}_n$ (the filtration is defined by Eq. (17)). By Lemma 7 and (H3b–c), for both cases, there exists a constant C such that almost-surely, for any $n \geq 0$,

$$\begin{aligned} \|A_{n+1} - A_n\| &\leq C \left(|\gamma_{n+1} - \gamma_n| + \gamma_n^2 + \gamma_n \gamma_{n+1} \right) \dots \\ &\quad \times (1 + \|S_n - \bar{S}(\theta_{n-1})\|), \\ \|A_{n+1}\| &\leq C \gamma_{n+1}. \end{aligned}$$

We then conclude by Proposition 4.

Proposition 3 Assume (H4a) and

$$\sup_{\theta \in \Theta} (\|\Psi(\theta)\| + \|\bar{S}(\theta)\|) < \infty.$$

Then there exists a constant C such that

$$\sum_n \gamma_{n+1}^2 \mathbb{E} \left[\|\Psi(\theta_n) (S_{n+1} - \bar{S}(\theta_n))\|^2 \right] \leq C \sum_n \gamma_{n+1}^2.$$

Proof We write

$$\begin{aligned} & \mathbb{E} \left[\|\Psi(\theta_n) (S_{n+1} - \bar{S}(\theta_n))\|^2 \right] \\ & \leq 2 \sup_{\Theta} \|\Psi(\theta)\|^2 \left(\sup_n \mathbb{E} \left[\|S_n\|^2 \right] + \sup_{\Theta} \|\bar{S}\|^2 \right), \end{aligned}$$

and conclude by Lemma 6. □

Proposition 4 Let $\{\theta_n, n \geq 0\}$ be given by Algorithm 2. Assume H1, H3, (H4a–b) and (H5a) and (H4a–b). In the biased case, assume also (H4c) and (H5b). Let $\{A_n, n \geq 0\}$ be a sequence of $d' \times q$ random matrices such that for any $n \geq 0$, $A_{n+1} \in \mathcal{F}_n$, and there exists a constant C_* such that almost-surely

$$\|A_{n+1}\| \leq C_* \gamma_{n+1}, \tag{35}$$

$$\|A_{n+1} - A_n\| \leq C_* a_{n+1} (1 + \|S_n - \bar{S}(\theta_{n-1})\|); \tag{36}$$

here $a_{n+1} := \gamma_n \gamma_{n+1} + \gamma_n^2 + |\gamma_{n+1} - \gamma_n|$. Then, almost-surely, the series $\sum_n A_{n+1} (S_{n+1} - \bar{S}(\theta_n))$ converges.

By Lemma 5 applied with $\sigma_n = \bar{S}(\theta_n)$, we decompose this sum into four terms:

$$\mathcal{T}_1 := \sum_{n \geq 2} \Delta_{2:n} A_1 (S_1 - \bar{S}(\theta_0))$$

$$= D_2 A_1 (S_1 - \bar{S}(\theta_0)),$$

$$\mathcal{T}_2 := \sum_{n \geq 2} \sum_{j=2}^n \Delta_{j:n} (A_j - A_{j-1}) (S_{j-1} - \bar{S}(\theta_{j-2}))$$

$$= \sum_{j \geq 2} D_j (A_j - A_{j-1}) (S_{j-1} - \bar{S}(\theta_{j-2})),$$

$$\mathcal{T}_3 := \sum_{n \geq 2} \sum_{j=2}^n \Delta_{j:n} A_j (\bar{S}(\theta_{j-2}) - \bar{S}(\theta_{j-1}))$$

$$= \sum_{j \geq 2} D_j A_j (\bar{S}(\theta_{j-2}) - \bar{S}(\theta_{j-1})),$$

$$\mathcal{T}_4 := \sum_{n \geq 2} \sum_{j=2}^n \delta_j \Delta_{j+1:n} A_j \left(m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) \right)$$

$$= \sum_{j \geq 2} \delta_j (1 + D_{j+1}) A_j \left(m_j^{-1} \sum_{k=1}^{m_j} S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) \right).$$

We have by using Eq. (35),

$$D_2 \|A_1 (S_1 - \bar{S}(\theta_0))\| \leq D_2 C_* \gamma_1 \left(\|S_1\| + \sup_{\Theta} \|\bar{S}\| \right).$$

By (H5a), $D_2 < \infty$ so the RHS is finite thus implying that \mathcal{T}_1 is finite almost-surely.

Using Eq. (36), there exists a constant C such that

$$\begin{aligned} & \mathbb{E} \left[\sum_{j \geq 2} D_j \| (A_j - A_{j-1}) (S_{j-1} - \bar{S}(\theta_{j-2})) \|^2 \right] \\ & \leq C \left(1 + \sup_n \mathbb{E} \left[\|S_n\|^2 \right] + \sup_{\Theta} \|\bar{S}\|^2 \right) \sum_{j \geq 2} a_j D_j. \end{aligned}$$

By (H3b)–(H3c), (H5a), and Lemma 6, the RHS is finite thus implying that \mathcal{T}_2 is finite almost-surely.

Similarly, there exists a constant C such that

$$\begin{aligned} & \mathbb{E} \left[\sum_{j \geq 2} D_j \| A_j (\bar{S}(\theta_{j-2}) - \bar{S}(\theta_{j-1})) \|^2 \right] \\ & \leq C \sum_{j \geq 2} \gamma_j D_j \mathbb{E} \left[\|\bar{S}(\theta_{j-1}) - \bar{S}(\theta_{j-2})\|^2 \right]. \end{aligned}$$

By (H3c), the RHS is bounded (up to a multiplicative constant) by $\sum_{j \geq 2} \gamma_j D_j \mathbb{E} [\|\theta_{j-1} - \theta_{j-2}\|]$; and by (H5a) and Lemmas 6 and 7, this sum is finite. Hence, \mathcal{T}_3 is finite almost-surely.

We give the proof of the convergence of the last term in the biased case: $\mathbb{E} [S(Z_{k,n}) | \mathcal{F}_n] \neq \bar{S}(\theta_n)$. The proof in the unbiased case corresponds to the following lines with $R_{j,1} = R_{j,2} = 0$ and $\hat{S}_\theta = S$. Set $\bar{D}_j := \delta_j (1 + D_{j+1})$. By Lemma 8, there exists \hat{S}_θ such that

$$\begin{aligned} & S(Z_{k,j-1}) - \bar{S}(\theta_{j-1}) \\ & = \hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k,j-1}), \end{aligned}$$

and $\sup_{\theta \in \Theta} \left| \hat{S}_\theta \right|_{\sqrt{W}} < \infty$. Hence, we have

$$\mathcal{T}_4 = \sum_{j \geq 2} \bar{D}_j A_j (\partial M_j + R_{j,1} + R_{j,2}),$$

where

$$\partial M_j := m_j^{-1} \sum_{k=1}^{m_j} \left(\hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k-1,j-1}) \right),$$

$$R_{j,1} := m_j^{-1} \left(P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{0,j-1}) - P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}) \right),$$

$$R_{j,2} := m_j^{-1} \left(P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{0,j}) \right).$$

Upon noting that $\mathbb{E}[A_j \partial M_j | \mathcal{F}_{j-1}] = 0$, the almost-sure convergence of the series $\sum_j \bar{D}_j A_j \partial M_j$ is proved by checking criteria for the almost-sure convergence of a martingale. By (35), there exists a constant C such that

$$\sum_j \bar{D}_j^2 \mathbb{E} \left[\|A_j \partial M_j\|^2 \right] \leq C \sum_j \frac{\gamma_j^2}{m_j^2} \bar{D}_j^2 \dots \times \mathbb{E} \left[\left\| \sum_{k=1}^{m_j} \left(\hat{S}_{\theta_{j-1}}(Z_{k,j-1}) - P_{\theta_{j-1}} \hat{S}_{\theta_{j-1}}(Z_{k-1,j-1}) \right) \right\|^2 \right].$$

By (H5a), Lemma 6 and (Hall and Heyde 1980, Theorem 2.10), the RHS is finite. Hall and Heyde (1980, Theorem 2.17) implies that $\sum_j \bar{D}_j A_j \partial M_j$ is finite almost-surely. For the second term, we write

$$\sum_{j \geq 2} \bar{D}_j A_j R_{j,1} = m_2^{-1} \bar{D}_2 A_2 P_{\theta_1} \hat{S}_{\theta_1}(Z_{0,1}) + \sum_{j \geq 2} \left(m_{j+1}^{-1} \bar{D}_{j+1} A_{j+1} - m_j^{-1} \bar{D}_j A_j \right) P_{\theta_j} \hat{S}_{\theta_j}(Z_{0,j}),$$

so that, by Lemmas 6 and 8, this series is finite almost-surely if $\sum_j \mathbb{E} \left[\|m_{j+1}^{-1} \bar{D}_{j+1} A_{j+1} - m_j^{-1} \bar{D}_j A_j\| \right] < \infty$. From Eqs. (35) and (36), there exists a constant C such that

$$\|m_{j+1}^{-1} \bar{D}_{j+1} A_{j+1} - m_j^{-1} \bar{D}_j A_j\| \leq C \gamma_{j+1} \left| m_{j+1}^{-1} \bar{D}_{j+1} - m_j^{-1} \bar{D}_j \right| + m_j^{-1} \bar{D}_j a_{j+1} (1 + \|S_j - \bar{S}(\theta_{j-1})\|).$$

H5 and Lemma 6 imply that

$$\sum_j \mathbb{E} \left[\|m_{j+1}^{-1} \bar{D}_{j+1} A_{j+1} - m_j^{-1} \bar{D}_j A_j\| \right] < \infty.$$

Finally, by (35), Lemmas 6 to 8, there exists a constant C such that

$$\sum_{j \geq 2} \mathbb{E} \left[\bar{D}_j \|A_j R_{j,2}\| \right] \leq C \sum_{j \geq 2} \gamma_j^2 m_j^{-1} \bar{D}_j.$$

The RHS is finite by H5 thus implying that $\sum_j \bar{D}_j A_j R_{j,2}$ is finite almost-surely.

References

Andrieu, C., Moulines, E.: On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Prob.* **16**(3), 1462–1505 (2006)
 Andrieu, C., Thoms, J.: A tutorial on adaptive MCMC. *Stat. Comput.* **18**(4), 343–373 (2008)
 Atchadé, Y., Fort, G., Moulines, E.: On perturbed proximal gradient algorithms. *J. Mach. Learn. Res.* **18**(10), 1–33 (2017)

Bartholomew, D., Knott, M., Moustaki, I.: *Latent Variable Models and Factor Analysis*. Wiley Series in Probability and Statistics, 3rd edn. Wiley, Chichester (2011)
 Bauschke, H., Combettes, P.: *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, New York (2011)
 Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* **2**(1), 183–202 (2009)
 Benveniste, A., Métivier, M., Priouret, P.: *Adaptive Algorithms and Stochastic Approximations, Applications of Mathematics*, vol. 22. Springer, Berlin (1990)
 Bertrand, J., Balding, D.: Multiple single nucleotide polymorphism analysis using penalized regression in nonlinear mixed-effect pharmacokinetic models. *Pharmacogenet. Genomics* **23**(3), 167–174 (2013)
 Bickel, P.J., Doksum, K.A.: *Mathematical Statistics—Basic Ideas and Selected Topics*, vol. 1, 2nd edn. Texts in Statistical Science Series. CRC Press, Boca Raton (2015)
 Chen, J., Chen, Z.: Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* **95**(3), 759–771 (2008)
 Chen, H., Zeng, D., Wang, Y.: Penalized nonlinear mixed effects model to identify biomarkers that predict disease progression. *Biometrics* **73**, 1343–1354 (2017)
 Combettes, P., Pesquet, J.: Proximal splitting methods in signal processing. In: Bauschke, H., Burachik, R., Combettes, P., Elser, V., Luke, D., Wolkowicz, H. (eds.) *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer Optimization and Its Applications, vol. 49. Springer, New York (2011)
 Combettes, P.L., Pesquet, J.C.: Stochastic quasi-fejér block-coordinate fixed point iterations with random sweeping. *SIAM J. Optim.* **25**(2), 1221–1248 (2015)
 Combettes, P., Pesquet, J.: Stochastic approximations and perturbations in forward–backward splitting for monotone operators. *Online J. Pure Appl. Funct. Anal.* **1**(1), 1–37 (2016)
 Combettes, P.L., Wajs, V.R.: Signal recovery by proximal forward–backward splitting. *Multiscale Model. Simul.* **4**(4), 1168–1200 (2005)
 Delavenne, X., Ollier, E., Basset, T., Bertolotti, L., Accassat, S., Garcin, A., Laporte, S., Zufferey, P., Mismetti, P.: A semi-mechanistic absorption model to evaluate drug–drug interaction with dabigatran: application with clarithromycin. *Br. J. Clin. Pharmacol.* **76**(1), 107–113 (2013)
 Delyon, B., Lavielle, M., Moulines, E.: Convergence of a stochastic approximation version of the EM algorithm. *Ann. Stat.* **27**(1), 94–128 (1999)
 Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B* **39**, 1–38 (1977)
 Fort, G., Moulines, E., Priouret, P.: Convergence of adaptive and interacting Markov chain Monte Carlo algorithms. *Ann. Stat.* **39**(6), 3262–3289 (2011a)
 Fort, G., Moulines, E., Priouret, P.: Convergence of adaptive and interacting Markov chain Monte Carlo algorithms. *Ann. Stat.* **39**(6), 3262–3289 (2011b)
 Fort, G., Jourdain, B., Kuhn, E., Lelièvre, T., Stoltz, G.: Convergence of the Wang–Landau algorithm. *Math. Comput.* **84**(295), 2297–2327 (2015)
 Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**(1), 1–22 (2010)
 Gouin-Thibault, I., Delavenne, X., Blanchard, A., Siguret, V., Salem, J., Narjoz, C., Gaussem, P., Beaune, P., Funck-Brentano, C., Azizi, M., et al.: Interindividual variability in dabigatran and rivaroxaban exposure: contribution of abcb1 genetic polymorphisms and inter-

- action with clarithromycin. *J. Thromb. Haemost.* **15**(2), 273–283 (2017)
- Hall, P., Heyde, C.C.: Probability and mathematical statistics. In: Hall, P., Heyde, C.C. (eds.) *Martingale Limit Theory and Its Application*. Academic Press, New York (1980)
- Kushner, H., Yin, G.: *Stochastic Approximation and Recursive Algorithms and Applications: Applications of Mathematics*, vol. 35, 2nd edn. Springer, New York (2003)
- Lehmann, E., Casella, G.: *Theory of Point Estimation*. Springer, New York (2006)
- Levine, R.A., Fan, J.: An automated (Markov chain) Monte Carlo EM algorithm. *J. Stat. Comput. Simul.* **74**(5), 349–359 (2004)
- Lin, J., Rosasco, L., Villa, S., Zhou, D.: Modified Fejer sequences and applications. Technical report. [arXiv:1510.04641v1](https://arxiv.org/abs/1510.04641v1) [math.OC] (2015)
- McLachlan, G., Krishnan, T.: *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics, 2nd edn. Wiley-Interscience, Hoboken (2008)
- Meyer, R.R.: Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *J. Comput. Syst. Sci.* **12**(1), 108–121 (1976)
- Meyn, S., Tweedie, R.L.: *Markov Chains and Stochastic Stability*, 2nd edn. Cambridge University Press, Cambridge (2009)
- Moreau, J.J.: Fonctions convexes duales et points proximaux dans un espace hilbertien. *Compt. Rendus Math. l'Acad. Sci.* **255**, 2897–2899 (1962)
- Ng, S., Krishnan, T., McLachlan, G.: The EM algorithm. In: Gentle, J., Härdle, W., Mori, Y. (eds.) *Handbook of Computational Statistics—Concepts and Methods*, vol. 1, 2nd edn, pp. 139–172. Springer, Heidelberg (2012)
- Ollier, E., Hodin, S., Basset, T., Accassat, S., Bertoletti, L., Mismetti, P., Delavenne, X.: In vitro and in vivo evaluation of drug–drug interaction between dabigatran and proton pump inhibitors. *Fundam. Clin. Pharmacol.* **29**(6), 604–614 (2015)
- Ollier, E., Samson, A., Delavenne, X., Viallon, V.: A saem algorithm for fused lasso penalized nonlinear mixed effect models: application to group comparison in pharmacokinetics. *Comput. Stat. Data Anal.* **95**, 207–221 (2016)
- Parikh, N., Boyd, S.: Proximal Algorithms. *Found. Trends Optim.* **1**(3), 123–231 (2013)
- Robert, C.P., Casella, G.: *Monte Carlo Statistical Methods*. Springer Texts in Statistics, 2nd edn. Springer, New York (2004)
- Roberts, G., Rosenthal, J.: Coupling and ergodicity of adaptive MCMC. *J. Appl. Prob.* **44**, 458–475 (2007)
- Rosasco, L., Villa, S., Vu, B.: Convergence of a stochastic proximal gradient algorithm. Technical report. [arXiv:1403.5075v3](https://arxiv.org/abs/1403.5075v3) (2014)
- Rosasco, L., Villa, S., Vu, B.: A stochastic inertial forward–backward splitting algorithm for multivariate monotone inclusions. *Optimization* **65**(6), 1293–1314 (2016)
- Saksman, E., Vihola, M.: On the ergodicity of the adaptive Metropolis algorithm on unbounded domains. *Ann. Appl. Prob.* **20**(6), 2178–2203 (2010)
- Samson, A., Lavielle, M., Mentré, F.: The SAEM algorithm for group comparison tests in longitudinal data analysis based on non-linear mixed-effects model. *Stat. Med.* **26**(27), 4860–4875 (2007)
- Schreck, A., Fort, G., Moulines, E.: Adaptive equi-energy sampler: convergence and illustration. *ACM Trans. Model. Comput. Simul.* **23**(1), 5 (2013)
- Städler, N., Bühlmann, P., van de Geer, S.: ℓ_1 -penalization for mixture regression models. *Test* **19**(2), 209–256 (2010)
- Wei, G., Tanner, M.: A Monte-Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *J. Am. Stat. Assoc.* **85**, 699–704 (1990)
- Wu, C.F.: On the convergence properties of the EM algorithm. *Ann. Stat.* **11**(1), 95–103 (1983)
- Zangwill, W.: *Nonlinear Programming: A Unified Approach*. Prentice-Hall International Series in Management. Prentice-Hall, Englewood Cliffs (1969)