



HAL
open science

Constant Step Size Least-Mean-Square: Bias-Variance Trade-offs and Optimal Sampling Distributions

Alexandre Défossez, Francis Bach

► **To cite this version:**

Alexandre Défossez, Francis Bach. Constant Step Size Least-Mean-Square: Bias-Variance Trade-offs and Optimal Sampling Distributions. International Conference on Artificial Intelligence and Statistics (AISTATS), 2015, San Diego, United States. hal-01081578

HAL Id: hal-01081578

<https://inria.hal.science/hal-01081578>

Submitted on 9 Nov 2014

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

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

Public Domain

Constant Step Size Least-Mean-Square: Bias-Variance Trade-offs and Optimal Sampling Distributions

Alexandre Défossez

Département de Mathématiques

École Normale Supérieure

Paris, France

`alexandre.defossez@ens.fr`

Francis Bach

Département d'Informatique

École Normale Supérieure

Paris, France

`francis.bach@inria.fr`

November 29, 2014

We consider the least-squares regression problem and provide a detailed asymptotic analysis of the performance of averaged constant-step-size stochastic gradient descent (a.k.a. least-mean-squares). In the strongly-convex case, we provide an asymptotic expansion up to explicit exponentially decaying terms. Our analysis leads to new insights into stochastic approximation algorithms: (a) it gives a tighter bound on the allowed step-size; (b) the generalization error may be divided into a variance term which is decaying as $O(1/n)$, independently of the step-size γ , and a bias term that decays as $O(1/\gamma^2 n^2)$; (c) when allowing non-uniform sampling, the choice of a good sampling density depends on whether the variance or bias terms dominate. In particular, when the variance term dominates, optimal sampling densities do not lead to much gain, while when the bias term dominates, we can choose larger step-sizes that leads to significant improvements.

1 Introduction

For large-scale supervised machine learning problems, optimization methods based on stochastic gradient descent (SGD) lead to efficient algorithms that make a single or few passes over the data (Bottou and Le Cun, 2005; Bousquet and Bottou, 2008).

In recent years, for smooth problems, large step-sizes together with some form of averaging, have emerged as having optimal scaling in terms of number of examples, both with asymptotic (Polyak and Juditsky, 1992) and non-asymptotic (Bach and Moulines, 2011) results. However, these convergence rates in $O(1/n)$ are only optimal in the limit of large samples, and in practice where the asymptotic regime may not be reached, notably because of the high-dimensionality of the data, other non-dominant terms may come into play, which is the main question we are tackling in this paper.

We consider least-squares regression with constant-step-size stochastic gradient descent—a.k.a. least-mean-squares—(Macchi, 1995; Bach and Moulines, 2013), where the generalization error may be explicitly split into a *bias* term that characterizes how fast initial conditions are forgotten, and a *variance* term that is only impacted by the noise present in the prediction problem. In this paper, we first show that while the variance term is asymptotically dominant, the bias term may play a strong role, both in theory and in practice, that explains convergence behaviors typically seen in applications.

Another question that has emerged as important to improve convergence is the use of special sampling distributions (Nesterov, 2012; Needell et al., 2013; Zhao and Zhang, 2014). With our theoretical result, we can optimize the first-order asymptotic terms rather than traditional upper-bounds, casting a new light on the potential gains (or lack thereof) of such different sampling distributions.

More precisely, we make the following contributions:

- We provide in Section 3 a detailed asymptotic analysis of the performance of averaged constant-step-size SGD, with all terms up to exponentially decaying ones. We also give in Section 3.1 a tighter bound on the allowed step-size γ .
- In Section 3.4, the generalization error may be divided into a variance term which is (up to first order) decaying as $O(1/n)$, independently of the step-size γ , and a bias term that decays as $O(1/\gamma^2 n^2)$.
- When allowing non-uniform sampling, the choice of a good sampling density depends on whether the variance or bias terms dominate. In particular, as shown in Section 4, when the variance term dominates, optimal sampling densities do not lead to much gain, while when the bias term dominates, we can choose larger step-sizes that leads to significant improvements.

1.1 Problem setup

Let X be a random variable with values in \mathbb{R}^d and Y another random variable with values in \mathbb{R} . Throughout this paper, $\|\cdot\|$ denotes the Euclidean norm on \mathbb{R}^d . We assume that $\mathbb{E}[\|X\|^2] = \mathbb{E}[X^T X]$ is finite and we denote by $H = \mathbb{E}[X X^T] \in \mathbb{R}^{d \times d}$ the second-order moment matrix of X . Throughout the paper, we assume that H is invertible, or equivalently, in optimization terms, that we are in the strongly convex case (see, e.g., Nesterov, 2004). We denote by μ the smallest eigenvalue of H , so that we have $\mu > 0$. Note that in our asymptotic results, the leading terms do not depend explicitly on μ .

We wish to solve the following optimization problem:

$$\min_{w \in \mathbb{R}^d} \mathbb{E} \left[\|X^T w - Y\|^2 \right], \quad (1.1)$$

from a stream of independent and identically distributed samples $(X_i, Y_i)_{i \geq 0}$.

For any given $w \in \mathbb{R}^d$, we denote by $f(w) = \mathbb{E}[\|X^T w - Y\|^2]$ the expected loss; we denote by $w^* \in \mathbb{R}^d$ the optimal solution (as H is invertible, it is unique), and by $f^* = f(w^*) \in \mathbb{R}$ the value at the minimum.

This set-up covers two common situations:

- Single pass through the data*, where each observation is seen once and considered as an i.i.d. sample, which is the context we explicitly study in this paper; note that then, our bounds are on the *testing error*, i.e., on the expected error on unseen data.
- Multiple passes through a finite dataset*, where each sample (X_i, Y_i) is selected uniformly at random from the dataset; in this situation, the *training error* is explicitly minimized, a regularizer is often added and our bound corresponds to training errors. Moreover, dedicated algorithms (Schmidt et al., 2013; Shalev-Shwartz and Zhang, 2013) have then better convergence rates than stochastic gradient.

Averaged SGD with constant step-size. In this paper, we study the convergence of the algorithm described by Bach and Moulines (2013), which is averaged stochastic gradient descent with constant step-size, also often referred to as the averaged least-mean-squares algorithm (LMS) (see, e.g., Macchi, 1995).

From a starting point $w_0 \in \mathbb{R}^d$, at each iteration $i \geq 1$, an i.i.d. sample of (X_i, Y_i) is obtained and the following recursion is used:

$$\begin{aligned} w_i &= w_{i-1} - \gamma X_i (X_i^T w_{i-1} - Y_i), \\ \bar{w}_i &= \frac{1}{i+1} \sum_{k=0}^i w_k = \frac{1}{i+1} w_i + \frac{i}{i+1} \bar{w}_{i-1}, \end{aligned}$$

where $\gamma > 0$ is a user-defined step-size. We denote by $\varepsilon_i = X_i^T w^* - Y_i$ the residual. Note that by definition of w^* , $\mathbb{E}[\varepsilon_i X_i] = 0$. If the vector X includes a constant component (which is common in practice), then this implies that ε_i and X_i are *uncorrelated*. Note however that in general they are not *independent*, unless the model with independent homoscedastic noise is well-specified (which implies in particular that $E(Y_i | X_i) = X_i^T w^*$).

We denote by $f_i = f(\bar{w}_i)$ the value at the averaged iterate. When studying the properties of this algorithm it is more convenient to work with the following centered estimates

$$\eta_i = w_i - w^* \quad \text{and} \quad \bar{\eta}_i = \bar{w}_i - w^*,$$

for which one immediately gets

$$\eta_i = (I - \gamma X_i X_i^T) \eta_{i-1} + \gamma \varepsilon_i X_i,$$

which is the recursion that we study in this paper.

1.2 Related work

Stochastic gradient methods have been heavily studied. We mention in this section some of the works which are relevant for the present paper.

Analysis of stochastic gradient algorithms. Since the work of Nemirovski and Yudin (1983), it is known that the optimal convergence rate depends in general on the presence or absence of *strong convexity*, with rates of $O(1/n\mu)$ for μ -strongly convex problems, and $O(1/\sqrt{n})$ for non-strongly convex problems. Recently, for specific smooth situations with the square or logistic loss, these rates can be improved to $O(1/n)$ in both situations (Bach and Moulines, 2013). For least-squares, this is achieved with constant-step-size SGD, hence our main focus on this algorithm.

Asymptotic analysis of stochastic gradient descent. In this paper, we focus on finding asymptotic equivalents of the generalization errors of SGD algorithms (with explicit remainder terms). For decaying step-sizes and general loss functions, this was partially considered by Fabian (1968), but only without averaging. Moreover, the traditional analysis of Polyak-Rupert averaging (Polyak and Juditsky, 1992; Ruppert, 1988) also leads to asymptotic equivalents, also for decaying step-sizes, but only the (asymptotically dominant) variance terms are considered.

Non-uniform sampling. Non-uniform sampling has been already tackled from several points of views; for example, in the active learning literature, Kanamori and Shimodaira (2003) provide the optimal sampling density to optimize the generalization error (for an estimator obtained as the minimum of the empirical least-squares risk), leading to distributions that are the same than the one obtained in Section 4.3 (where the variance term dominates), for which the actual gains are limited.

Moreover, in the context of stochastic gradient descent methods, Needell et al. (2013); Zhao and Zhang (2014) show that by optimizing the sampling density, bounds on the convergence rates could be improved, but the actual gains are hard to quantify. Our focus on limits of convergence rates allows us to precisely quantify the gains and obtain extra insights (at least asymptotically).

2 Linear algebra prerequisites

Throughout our results we will use the following notations and results. These are necessary to provide explicit expressions for the constants in the asymptotic expansions.

For any real vector space V of finite dimension d , let $\mathcal{M}(V)$ be the space of linear operators over V which is isomorphic to the space of d -by- d matrices, with the usual results that composition becomes matrix multiplication. As a consequence we will use the same notation for the space of matrices and the space of endomorphisms.

We denote by $\mathcal{I} = \mathcal{M}(\mathcal{M}(\mathbb{R}^d))$ the space of endomorphisms on the space of matrices over \mathbb{R}^d . One can index the rows and columns of a matrix $M \in \mathcal{I}$ by a pair (i, j) where $1 \leq i \leq d$ and $1 \leq j \leq d$. We will often denote by $M_{(i,j),(k,l)}$ an element of this matrix on matrices. In the following we will drop the domain of i, j, k, l, i', j' which is implicitly $\{1, 2, \dots, d\}$. Explicitly, if $A \in \mathcal{M}(\mathbb{R}^d)$ and $M \in \mathcal{I}$, then MA is defined through:

$$\forall (i, j)(MA)_{i,j} = \sum_{i'=1, j'=1}^d M_{(i,j),(i',j')} A_{i',j'}$$

We will mostly make no distinction between A as a vector in $\mathcal{M}(\mathbb{R}^d)$ on which elements in \mathcal{I} can operate and A as a matrix in $\mathcal{M}(\mathbb{R}^d)$. Then MA can be either usual matrix multiplication if $M, A \in \mathcal{M}(\mathbb{R}^d)$ or $M, A \in \mathcal{I}$ or application of M to A if $M \in \mathcal{I}$ and $A \in \mathcal{M}(\mathbb{R}^d)$. However, if $M \in \mathcal{I}$ and $A \in \mathcal{M}(\mathbb{R}^d)$, then AM does not make sense. For $P \in \mathcal{I}$, and any (i, j) , we will define $P_{i,j}$ the matrix in $\mathcal{M}(\mathbb{R}^d)$ with coefficient (i', j') given by $P_{(i,j),(i',j')}$.

For any $V \in \mathcal{S}(\mathbb{R}^d)$ (the set of symmetric matrices of size d), we will denote $\|V\|_{\text{op}}$ the operator norm of V or equivalently its eigenvalue with the largest absolute value. For any $M \in \mathcal{I}$ so that $\mathcal{S}(\mathbb{R}^d)$ is stable under M , we will take $\|M\|_{\text{op}}$ the operator norm of M restricted to $\mathcal{S}(\mathbb{R}^d)$, defined with respect to the Frobenius norm on $\mathcal{S}(\mathbb{R}^d)$, that is

$$\|M\|_{\text{op}} = \sup_{V \in \mathcal{S}(\mathbb{R}^d), \|V\|_F=1} \|MV\|_F.$$

Equivalently, it is given by the largest absolute value of the eigenvalues of M .

Finally, we will look more precisely at three elements of \mathcal{I} . For any given $A \in \mathcal{M}(\mathbb{R}^d)$, one can define A_L (resp. A_R) so that A_L is the matrix in \mathcal{I} representing left multiplication (resp. right

multiplication) by A . The coefficient of A_L and A_R are given by

$$\begin{aligned} \forall (i, j), (k, l), \quad (A_L)_{(i,j),(k,l)} &= \delta_{j,l} A_{i,k} \\ \forall (i, j), (k, l), \quad (A_R)_{(i,j),(k,l)} &= \delta_{i,k} A_{j,l}. \end{aligned}$$

If A is symmetric, then A_L and A_H are both symmetric operators, and $A_L + A_H$ is stable on the subspace of symmetric matrices, that we have denoted $\mathcal{S}(\mathbb{R}^d)$.

Let X be a random variable in \mathbb{R}^d , we consider the linear operator M on $\mathcal{M}(\mathbb{R}^d)$ defined by,

$$\forall A \in \mathcal{M}(\mathbb{R}^d), \quad MA = \mathbb{E} [(X^T A X) X X^T],$$

then, the coefficients of the associated matrix are given by

$$\forall (i, j, k, l), \quad M_{(i,j),(k,l)} = \mathbb{E} [X^{(i)} X^{(j)} X^{(k)} X^{(l)}],$$

where $X^{(i)}$ denote the i -th component of the vector X . The matrix M is clearly symmetric. One can also prove that it is stable on $\mathcal{S}(\mathbb{R}^d)$.

We then define $T = H_L + H_R - \gamma M$ with H_L , H_R and M as defined above for the random variable X defined in our setup. It is immediately stable over $\mathcal{S}(\mathbb{R}^d)$. We will denote μ_T the smallest eigenvalue of T .

3 Main results

We will present results about the convergence of the algorithm which are derived from the exact computation of the second-order moment matrix, which we refer to as the *covariance matrix*, $\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T]$. Since we consider a least-squares problem, we have

$$f_n - f^* = \text{Tr}(H \mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T]).$$

We will distinguish two terms, which can be assimilated to a variance/bias decomposition. The *variance* term Δ^{variance} can be defined as the covariance matrix we would get starting from the solution (that is, $\eta_0 = 0$). On the other hand, the *bias* term Δ^{bias} is defined as the covariance matrix we would get if the model was noiseless, that is $Y = X^T w^*$ and $\varepsilon = 0$.

Each of these two terms leads to contribution to $f_n - f^*$, that is $\text{Tr}(H \Delta^{\text{variance}})$ and $\text{Tr}(H \Delta^{\text{bias}})$. Under extra assumption that are discussed in the supplementary material, such that when X and ε are independent (i.e., well-specified model), the actual covariance matrix is exactly the sum of the bias and variance matrices, and thus

$$f_n - f^* = \text{Tr}(H \Delta^{\text{variance}}) + \text{Tr}(H \Delta^{\text{bias}}).$$

Moreover, even when this is not true, it has been noted by Bach and Moulines (2013) that

$$f_n - f^* \leq 2\text{Tr}(H \Delta^{\text{variance}}) + 2\text{Tr}(H \Delta^{\text{bias}}),$$

that is, the sum of the two terms is a factor of two away from the exact generalization error.

3.1 Improved step-size

Let us take T as defined in Section 2. We will also define two contraction factors,

$$\rho_T = \|I - \gamma T\|_{\text{op}} \quad \text{and} \quad \rho_H = \|I - \gamma H\|_{\text{op}}, \quad (3.1)$$

as well as $\rho = \max(\rho_T, \rho_H)$, where $\|\cdot\|_{\text{op}}$ is defined as the largest eigenvalue in absolute value.

Let us define γ_{\max} as the supremum of the set of $\gamma > 0$ verifying, $\forall A \in \mathcal{S}(\mathbb{R}^d)$ (the set of symmetric matrices of size $d \times d$):

$$2\text{Tr}(A^T H A) - \gamma \mathbb{E} [(X^T A X)^2] > 0, \quad (3.2)$$

or equivalently as the supremum of $\gamma > 0$ such that T is definite positive. One can actually show that we necessarily have that

$$\gamma_{\max} \leq 2/\text{Tr}(H), \quad (3.3)$$

and the following lemma (see proof in the Appendix):

Lemma 1. *Using the notations and assumptions of Section 1.1, define γ_{\max} as the supremum of $\gamma > 0$ such that*

$$\forall A \in \mathcal{S}(\mathbb{R}^d), \quad 2\text{Tr}(A^T H A) - \gamma \mathbb{E}[(X^T A X)^2] > 0. \quad (3.4)$$

If $0 < \gamma < \gamma_{\max}$ then T is positive definite and $\rho < 1$. More precisely, in dimension $d \geq 2$, we have

$$\begin{cases} \rho \leq 1 - 2\gamma \left(1 - \frac{\gamma}{\gamma_{\max}}\right) \mu & \text{if } 1 > \frac{\gamma}{\gamma_{\max}} \geq \frac{1}{2} \\ \rho \leq 1 - \gamma \mu & \text{otherwise.} \end{cases} \quad (3.5)$$

In dimension $d = 1$, we have

$$\rho \leq \max \left(|1 - \gamma \mu|, 1 - 2\gamma \left(1 - \frac{\gamma}{\gamma_{\max}}\right) \mu \right).$$

Otherwise, if $\gamma > \gamma_{\max}$, then $\rho > 1$.

Note that we may rewrite γ_{\max} as

$$\frac{2}{\gamma_{\max}} = \sup_{A \in \mathcal{S}(\mathbb{R}^d)} \frac{\mathbb{E}[(X^T A X)^2]}{\text{Tr}(A^T H A)},$$

which can be computed explicitly by a generalized eigenvalue problem once all second- and fourth-order moments of X are known. This is to be contrasted with the largest step-size $\gamma_{\max}^{\text{det}}$ for deterministic gradient descent, which is such that

$$\frac{2}{\gamma_{\max}^{\text{det}}} = \sup_{a \in \mathbb{R}^d} \frac{a^T H a}{a^T a}.$$

One can observe that for any distribution on X , we necessarily have $\gamma_{\max} \leq 2/\text{Tr}(H) \leq \gamma_{\max}^{\text{det}}$ so that the maximum stochastic step-size will always be smaller than the deterministic, one as one would expect.

Note also that the step-size provided by γ_{\max} is a strict improvement (see supplementary material) on the one proposed by Bach and Moulines (2013), which is equal to the supremum of the set of $\gamma > 0$ such that $\mathbb{E}[X X^T] - \gamma \mathbb{E}[(X^T X) X X^T]$ is positive definite.

We conjecture that the bound given by γ_{\max} is tight, namely that if γ is larger than γ_{\max} then there exists an initial condition η_0 such that the algorithm diverges.

3.2 Bias term

In this section, we provide an asymptotic expansion of the bias term Δ^{bias} .

Theorem 1 (Asymptotic covariance of the bias term). *Let $E_0 = \eta_0 \eta_0^T$. If $0 < \gamma < \gamma_{\max}$ and $\forall i \geq 1, \varepsilon_i = 0$, then*

$$\Delta^{\text{bias}} = \mathbb{E}[\bar{\eta}_n \bar{\eta}_n^T] = \frac{1}{n^2 \gamma^2} (H_L^{-1} + H_R^{-1} - \gamma I) (T^{-1} E_0) + O\left(\frac{\rho^n}{n}\right). \quad (3.6)$$

Explicit bounds are given in the proof in the Appendix.

A detailed proof is given in the Appendix. Using Lemma 1, we know that $\rho < 1$ so that (3.6) converges as n^{-2} . From that we can derive that the rate of convergence for $\text{Tr}(H \Delta^{\text{bias}})$, that will be of order n^{-2} as well. Although the dependency of $A(\gamma)$ is complex, one can easily derive an equivalent when γ tends to zero, and we have

$$\lim_{n \rightarrow \infty} n^2 \text{Tr}(H \Delta^{\text{bias}}) \underset{\gamma \rightarrow 0}{\sim} \gamma^{-2} \eta_0^T H^{-1} \eta_0. \quad (3.7)$$

3.3 Variance term

In this section, we provide an asymptotic expansion of the variance term Δ^{variance} .

Theorem 2 (Asymptotic covariance of the variance term). *Let $\Sigma_0 = \mathbb{E}[\varepsilon^2 X X^T]$ and let assume that $\eta_0 = 0$. If $0 < \gamma < \gamma_{\max}$ then Δ^{variance} is equal to:*

$$\mathbb{E}[\bar{\eta}_n \bar{\eta}_n^T] = \frac{1}{n} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} \Sigma_0 - \frac{1}{\gamma n^2} (H_L^{-1} + H_R^{-1} - \gamma I) (I - \gamma T) T^{-2} \Sigma_0 + O\left(\frac{\rho^n}{n}\right). \quad (3.8)$$

Explicit bounds are given in the proof in the Appendix.

A detailed proof is given in the Appendix.

Unsurprisingly, the asymptotic behavior of the variance term is dominant over the bias one as it decreases only as n^{-1} , which is the overall convergence rate of this algorithm of least-mean-squares as noted by Bach and Moulines (2013).

It is also possible to get a simpler equivalent when γ goes to 0:

$$\lim_{n \rightarrow \infty} n \text{Tr}(H \Delta^{\text{variance}}) \underset{\gamma \rightarrow 0}{\sim} \mathbb{E}[\varepsilon^2 X^T H^{-1} X].$$

If we further assume that the noise ε is independent of X , then we recover the usual

$$\lim_{n \rightarrow \infty} n \text{Tr}(H \Delta^{\text{variance}}) \underset{\gamma \sim 0}{\sim} d \sigma^2,$$

where $\sigma = \mathbb{E}[\varepsilon^2]$, which is the Cramer-Rao bound for such a problem. It is also interesting to notice that this is the exact same result as the one obtained by Polyak and Juditsky (1992) with a decreasing step-size.

Finally, note that if γ is small, the term in $\frac{1}{\gamma n^2} C(\gamma) \Sigma_0$ is always positive so that there is no risk of it exploding for small values of γ (unlike for the bias term).

3.4 Comparing both terms

As seen above with an asymptotic expansion around $\gamma = 0$, for n sufficiently large, the bias and variance terms are of order:

$$\begin{aligned} \text{Tr}(H \Delta^{\text{bias}}) &\sim \frac{1}{\gamma^2 n^2} \eta_0^T H^{-1} \eta_0 \\ \text{Tr}(H \Delta^{\text{variance}}) &\sim \frac{1}{n} \mathbb{E}[\varepsilon^2 X^T H^{-1} X]. \end{aligned}$$

The different behaviors of the bias and variance terms lead to two regimes, one in $(\gamma n)^{-2}$ and one in n^{-1} that can clearly be observed on synthetic data. On real world data, one will often observe a mixture of the two, depending on the step-size and the difficulty of the problem. Experimental results on both synthetic and real world data will be presented in Section 5.

4 Optimal sampling

Changing the sampling density may be interesting in several situations, in particular (a) in presence of outliers (i.e., points with large norms) and (b) classification problems with asymmetric costs (see Section 4.2).

4.1 Impact of sampling

Using the two previous theorems, we can now try to optimize the sampling distribution to increase performance. We will sample from a distribution q instead of the given distribution p . Since we wish to keep the same objective function, we will use importance weights $c(X, Y)$, so that if we denote by $\mathbb{E}_p[A]$ the expectation of a random variable A under the probability distribution given by p over A we have

$$\mathbb{E}_p[|X^T w - Y|^2] = \mathbb{E}_q[c(X, Y) |X^T w - Y|^2].$$

First, one can notice that, from a practical point of view, we must restrict ourselves to q that are absolutely continuous with respect to p as one cannot invent samples. In order to be able to define c we also need p to be absolutely continuous with respect to q , so that

$$c = \frac{dp}{dq}.$$

Besides, c^{-1} is defined as $c^{-1} = \frac{dq}{dp}$.

A key consequence of using least-squares is that for a given (q, c) pair, we only have to sample using q and scale X, Y and $\varepsilon = X^T w^* - Y$ by $\sqrt{c(X, Y)}$. Thus we can use the two previous theorems for $X' = \sqrt{c(X, Y)} X$ and $Y' = \sqrt{c(X, Y)} Y$ and sampling X, Y according to q .

As of now, we will assume that almost surely $X \neq 0$. Indeed, when $X_i = 0$, we perform no update so that we can just ignore such points.

One can notice that for any $A, B \in \{X, Y, \varepsilon\}$,

$$\mathbb{E}_q[A'B'] = \mathbb{E}_q[c(X, Y)AB] = \mathbb{E}_p[AB],$$

and thus all second-order moments are unchanged under resampling. This is the case for the matrix $H = \mathbb{E}_p [XX^T] = \mathbb{E}_q [X'X'^T]$ for instance.

However, for terms of order 4 like T , $\mathbb{E}[(X^T X)XX^T]$ or Σ_0 , an extra c appears and we have for instance

$$\mathbb{E}_q [(X'^T X')X'X'^T] = \mathbb{E}_p \left[\frac{p(X, Y)}{q(X, Y)} (X^T X)XX^T \right].$$

It means that while H will not be changed, T is impacted in non trivial ways, as is T^{-1} . This makes it tricky to truly optimize sampling for any γ . However, when assuming γ small, it is possible to optimize the limit we obtained for $\gamma \rightarrow 0$ (see Sections 4.3 and 4.4). The experiments we ran (see Section 5) seem to confirm that this is a valid assumption for values of γ as high as $\gamma_{\max}/2$.

4.2 Asymmetric binary classification

As a motivation for this work, we will present one practical application of resampling which is binary classification with highly asymmetric classes. Assume we have $Y \in \{-1, 1\}$ and that $\mathbb{P}\{Y = 1\}$ and $\mathbb{P}\{Y = -1\}$ are highly unbalanced, as it can be the case in various domains, such as ad click prediction or object detection, etc. Then, it can be useful in practice to give more weight the less frequent class (see, e.g., Perronnin et al., 2012, and references therein). This is equivalent to multiplying both X and Y by some constant $\sqrt{c_Y}$. A common choice is for instance to take $c_y = 1/\mathbb{P}\{Y = y\}$ which will give the same importance in the loss to both classes.

However, these weights will make the gradients from the less frequent class huge compared to the usual updates. This is likely to impact the convergence of the algorithm. In that case it is easy to notice that taking $c(x, y) = 1/c_y$ will leave the gradients unchanged but will favor sampling examples from the less frequent class.

4.3 Optimal sampling for the variance term

Let assume that we are only interested in the long term performance for our algorithm. Ultimately the variance term will be driving the performance and we need to optimize it.

Exactly optimizing the sampling for this case is uneasy as it impacts both Σ_0 and the terms $B(\gamma)$ and $C(\gamma)$ in Theorem 2, in a non trivial way. However, if we assume a small step-size γ , then we just have to minimize

$$\mathbb{E}_q [\varepsilon'^2 X'^T H^{-1} X'] = \mathbb{E}_p [c(X, Y) \varepsilon^2 X^T H^{-1} X],$$

under the constraint that $\mathbb{E}_p [c^{-1}(X, Y)] = 1$ so that q is a distribution. Using the Cauchy-Schwarz inequality, we have that

$$\begin{aligned} \mathbb{E}_p [c(X, Y) \varepsilon^2 X^T H^{-1} X] &= \mathbb{E}_p [c(X, Y) \varepsilon^2 X^T H^{-1} X] \mathbb{E}_p [c^{-1}(X, Y)] \\ &\geq (\mathbb{E}_p [|\varepsilon| \sqrt{X^T H^{-1} X}])^2. \end{aligned}$$

When $X \neq 0$ almost surely, then this lower-bound is achieved for

$$c^{-1}(X, Y) = \frac{|\varepsilon| \sqrt{X^T H^{-1} X}}{\mathbb{E}_p [|\varepsilon| \sqrt{X^T H^{-1} X}]},$$

which requires prior knowledge of H and ε .

In that case, we obtain

$$\lim_{n \rightarrow \infty} n \text{Tr}(H \Delta^{\text{variance}}) = \left(\mathbb{E} [|\varepsilon| \sqrt{X^T H^{-1} X}] \right)^2.$$

One can notice that this is the exact same optimal sampling as the one obtained in the active learning set-up by Kanamori and Shimodaira (2003).

Again, it is possible to slightly simplify this expression when ε and X are independent, as we obtain

$$\lim_{n \rightarrow \infty} n \text{Tr}(H \Delta^{\text{variance}}) = \sigma^2 \left(\mathbb{E} [\sqrt{X^T H^{-1} X}] \right)^2,$$

with $\sigma^2 = \mathbb{E}[\varepsilon^2]$. At this point it is important to realize that the gain we have here is of the order of

$$\mathbb{E} \left[\sqrt{X^T X} \right]^2 / \mathbb{E} [X^T X]. \quad (4.1)$$

During our experimentations on usual datasets, we have observed that this factor was always between 1/2 and 1 and thus there is little to be gained when optimizing the variance term.

4.4 Optimal sampling for the bias term

Although asymptotically the variance term will be the largest one, it is possible that initially the bias one is non negligible and it can be interesting to optimize for it. This is all the more possible as it depends much more on the step-size γ and if γ is too small, the bias term can stay larger than the variance term for many iterations.

If we assume γ small, then we can approximate the bias term by the expression given by (3.7), that is, proportional to $1/(\gamma^2 n^2)$. In this case, it is clear that we want to increase γ_{\max} and, because second-order moments are not impacted by resampling, it has no effect other than changing γ_{\max} . Numerical experiments tends to show that increasing γ is beneficial even for γ close to $\frac{\gamma_{\max}}{2}$. Beyond this limit, the approximation (3.7) is no longer sustainable and besides, exponentially decreasing terms can start to grow quite large.

The maximum step-size we can take is given by the tighter condition from Section 3.1, that is, $\forall A \in \mathcal{S}(\mathbb{R}^d)$,

$$2\text{Tr}(A^T H A) - \gamma \mathbb{E}[(X^T A X)^2] > 0, \quad (4.2)$$

which implies that

$$\frac{2}{\mathbb{E}[X^T X]} \geq \gamma_{\max}, \quad (4.3)$$

using (3.3). As this upper bound on γ_{\max} only depends on moments of order 2, and that those moments are not changed by resampling, (4.3) is an upper bound on any γ_{\max} for a given optimization problem, no matter how we resample. It turns out it can be achieved by the resampling given by

$$c_*^{-1}(X, Y) = \frac{X^T X}{\mathbb{E}_p[X^T X]} \text{ a.s.},$$

which, unlike the variance term, does not require the knowledge of H . We have

$$\begin{aligned} H - \gamma \mathbb{E}_q [c_*(X, Y)^2 X^T X X X^T] &= 2H - \gamma \mathbb{E}_p \left[\frac{\mathbb{E}_p[X^T X]}{X^T X} (X^T X) X X^T \right] \\ &= H(2 - \gamma \mathbb{E}_p[X^T X]), \end{aligned}$$

which is positive definite as soon as $\gamma < \frac{2}{\text{Tr}(H)}$. Besides, one can prove that $H - \gamma \mathbb{E}[(X^T X) X X^T] \succ 0$ is a stronger form of (4.2) and implies it, as we already noted in Section 3.1. This means that using the resampling defined by c_* , we have $\gamma_{\max} = \frac{2}{\mathbb{E}[X^T X]}$ which is thus not improvable.

If $\gamma_{\max}^{(0)}$ is the maximum step-size before resampling and $\gamma_{\max}^{(1)}$ is the maximum step-size after resampling, then the gain for $f_n - f^*$ is a factor $(\gamma_{\max}^{(0)}/\gamma_{\max}^{(1)})^2$. It can be hard to evaluate, but from our experiments (see Section 5) it was common to observe gain factor of 1/100 or 1/400 while the gain for the variance term was limited to 1/2.

Unlike the variance term, the resampling in itself here has an impact only through a larger step-size. Resampling while keeping the same step-size will often lead to almost identical performances for the bias term. It is interesting to note that when $H = I$, this sampling will exactly have no impact at all on the variance term, and when $H \neq I$, it will only impact it marginally.

Link with other algorithms. When using this resampling and $\gamma = 1/\mathbb{E}[X^T X]$, the update step becomes

$$w_i = w_{i-1} - \frac{1}{X_i^T X_i} (X_i^T w_{i-1} - Y_i), \quad (4.4)$$

where X_i is sampled from q_* . This is very similar to normalized least mean squares (NLMS) by Bershada (1986), i.e., we first normalize X (and Y by the same factor), and then we run the usual

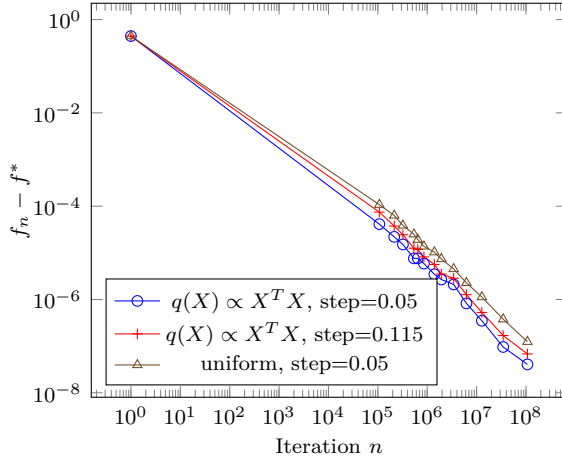


Figure 1: Convergence on *Yahoo* dataset without weights.

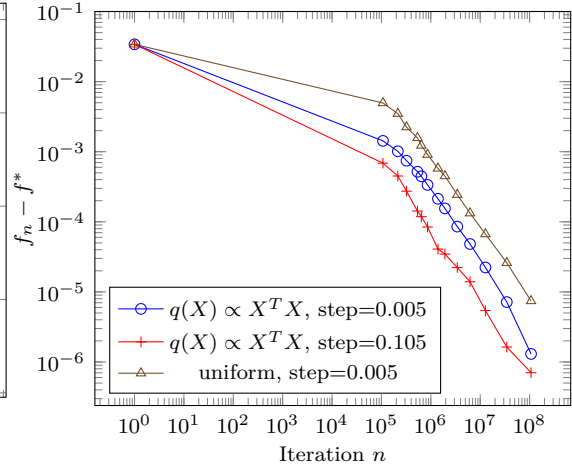


Figure 2: Convergence on *Yahoo* dataset with weights.

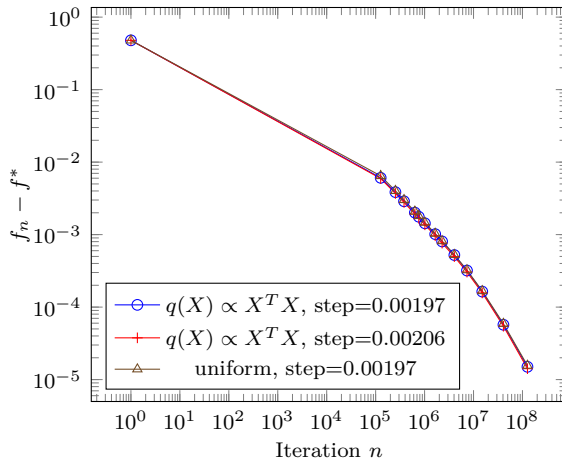


Figure 3: Convergence on *Sido* dataset without weights.

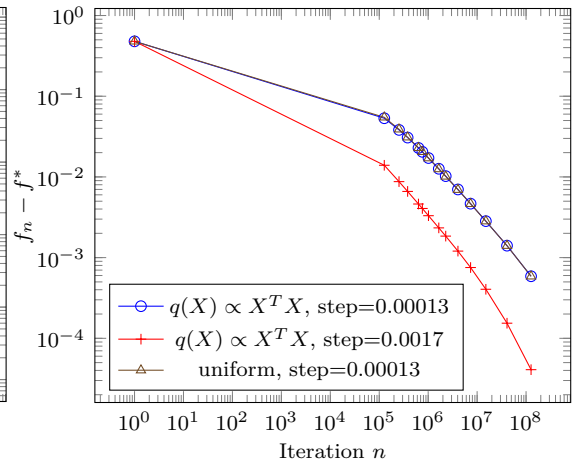


Figure 4: Convergence on *Sido* dataset with weights.

stochastic gradient descent with a step-size of 1. However, while NLMS does not optimize the same overall objective function, we remember the norm of X in c_* and sample large ones more often and keep the same overall objective function. One can also notice some links with implicit stochastic gradient descent (ISGD) by Toulis et al. (2014), where the update rule is

$$w_i = w_{i-1} - \frac{\gamma_i}{1 + \gamma_i X_i^T X_i} (X_i^T w_{i-1} - Y_i), \quad (4.5)$$

which is similar to NLMS and (4.4) when γ_i is large. As γ_i is a decreasing step-size it means that during the early iterations, ISGD will behave like NLMS before switching to a regular stochastic gradient descent as γ_i goes to 0. This comforts us in the idea that a large step-size is crucial during the early stages, as highlighted by our analysis.

5 Experiments

Yahoo and Sido datasets. We have tried to observe evidence of our predictions into two unbalanced datasets, “yahoo” and “sido”. The “yahoo” dataset¹ is composed of millions of triple (ads, context, click) from the Yahoo front page where click is 1 if the user clicked on the given ad in the given context (composed of 136 Boolean features) and 0 otherwise. It is very unbalanced, as the click rate is low. For this experiment we only looked at the rows corresponding to a specific

¹webscope.sandbox.yahoo.com/

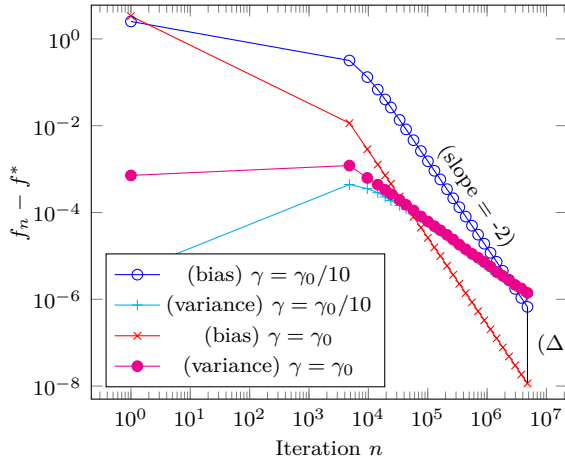


Figure 5: Convergence per term on synthetic data.

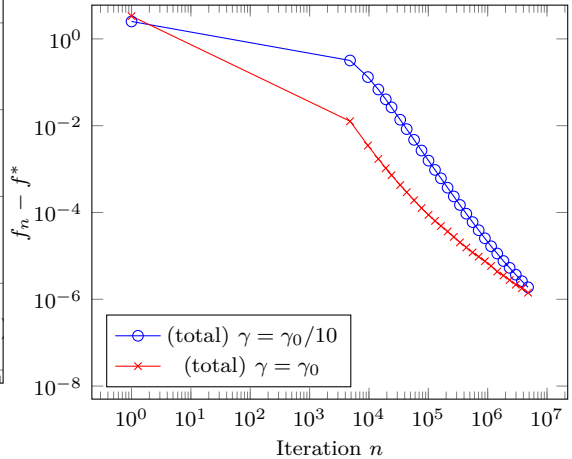


Figure 6: Convergence on synthetic data.

ad (107355 rows), with a click rate of 0.03%. We used it both with and without weights in order to give the same importance to both clicks and non clicks as explained in Section 4.2.

On Figures 1 and 2, we compared the incidence of the step-size and the resampling of performance. When using weights, the maximum step-size happens to be divided by ten, but when resampling proportionally to $X^T X$, we recover almost the same step-size as without weights. Moreover, without weights (Figure 1), the lowest step-size performs best, which tends to indicate that the variance term is dominant. However, with weights (Figure 2), one can observe the more pronounced dependency in γ , which show that the bias term became non negligible. Comparable results, but obtained on the “sido” dataset² are given in Figures 3 and 4. This dataset is composed of 12678 points with 4932 features. The less frequent class represents 3.6% of the points.

We did not plot the graphs obtained when sampling proportionally to $\sqrt{X^T H^{-1} X}$ as they are mostly the same as without resampling. If we were to achieve the regime where the variance is completely dominant, the error would be at best be divided by two for the “yahoo” dataset and would be almost the same for “sido” (divided by 1.02). These potential gains were computed directly using the expression (4.1).

Synthetic data and bias-variance decomposition. We also observed exactly the bias and variance terms on synthetic data. The data consist in an infinite stream of points X sampled from a normal distribution with covariance matrix H , so that the eigenvalues of H are $(\frac{1}{i})_{1 \leq i \leq 25}$. Y is given by $Y = X^T w_* + \varepsilon$ for a fixed $w_* \in \mathbb{R}^{25}$ and $\varepsilon \sim \mathcal{N}(0, 1)$. On Figures 5 and 6 one can see the decomposition of the error between variance and bias terms (Figure 5) as well as the sum of both the variance and bias error (Figure 6). We see that both the variance and bias curves quickly reach their asymptotic regime with a slope of -1 for the variance in log/log space and -2 for the bias as expected. We ran the algorithm with two step-sizes, one being ten times larger than the other. We observe the ratio of 100 as expected between the two bias curves and almost no difference at all for the variance ones, except for the first iterations.

One can also see the effect of the step-size γ at a fixed number of iterations on Figure 7. Due to the symmetry between n and γ in the expression of the bias term, one can notice the resemblance at first between this curve and the one obtained when plotting the $f_n - f^*$ against n . However, when the step-size is large, we get sooner into the regime where the variance dominates. At this point we observe almost no influence of the step-size on $f_n - f^*$. When getting closer to the maximum step-size, convergence becomes very slow as ρ becomes close to 1.

6 Conclusion

In this paper, we have provided a tighter analysis of constant-step-size LMS, leading to a better understanding of the convergence of the algorithm at different stages, in particular regarding how the initial condition is forgotten.

We were able to deduce different sampling schemes depending on what regime we are in. Sampling proportionally to $\sqrt{X^T H^{-1} X}$ is always asymptotically the best method. The potential gain is however limited most of the time. Besides, for datasets that are more “difficult”, that is with

²www.causality.inf.ethz.ch/data/SIDO.html

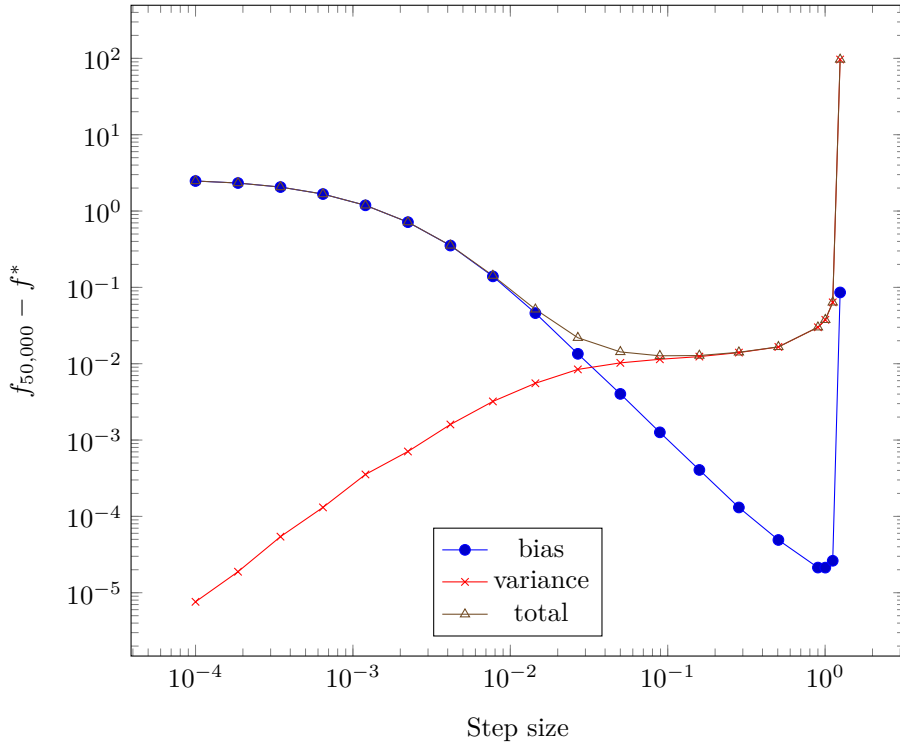


Figure 7: Impact of step-size on error, with its bias/variance decomposition.

moments that increases quickly, forgetting the initial condition can happen arbitrary slow due to the strong dependency in the step-size. If this is the case, then sampling proportionally to $X^T X$ will allow us to take a much larger step-size which will then lead to a smaller error.

Our work can be extended in several ways: for simplicity we have focused on least-squares problems where the bias/variance decomposition is explicit. It would be interesting to see how these results can be extended to other smooth losses such as logistic regression, where constant-step size SGD does not converge to the global optimum (Nedic and Bertsekas, 2000; Bach and Moulines, 2013). Moreover, we have only provided results in expectations and a precise study of higher-order moments would give a better understanding of additional potential effects of resampling.

Acknowledgements

This work was partially supported by a grant from the European Research Council (SIERRA project 239993). We thank Aymeric Dieuleveut and Nicolas Flammarion for interesting discussions related to this work.

Appendix

We give hereafter the proofs for the different results in the main paper. Unless otherwise specified, references are to the present Appendix. We first give a more thorough definition of the space in which our operators live. We then proceed to a proof of Lemma 1. Finally we detail the computation that allowed us to derive both theorems in this paper.

Proof of Lemma 1

We will first need some preliminary results in order to provide a proof of Lemma 1.

Some Lemmas

Lemma 2. *Let $A \in \mathcal{S}(\mathbb{R}^d)$ be any symmetric matrix, then*

$$\forall x \in \mathbb{R}^d, (x^T Ax)^2 \leq \text{Tr}((x^T x)Axx^T A).$$

Proof. Using Cauchy-Schwarz inequality, one has

$$\begin{aligned} (x^T Ax)^2 &= [x^T(Ax)]^2 \leq (Ax)^T(Ax)(x^T x) \\ &= x^T AAx(x^T x) = \text{Tr}((x^T x)Axx^T A). \end{aligned}$$

□

The following lemma is the proof of equation (2.4) in the original paper.

Lemma 3. *Let $H \in \mathcal{S}(\mathbb{R}^d)$ be a positive semi-definite matrix. If $\gamma > 0$ is so that*

$$\forall A \in \mathcal{S}(\mathbb{R}^d), 2\text{Tr}(A^T HA) - \gamma \mathbb{E}[(X^T AX)^2] > 0$$

then

$$\gamma < \frac{2}{\text{Tr}(H)}.$$

Proof. Let $A \in \mathcal{S}(\mathbb{R}^d)$, $2\text{Tr}(A^T HA) - \gamma \mathbb{E}[(X^T AX)^2] > 0$ implies that with Jensen's inequality,

$$\begin{aligned} 2\text{Tr}(A^T HA) - \gamma \text{Tr}(AH)^2 &= 2\text{Tr}(A^T HA) - \gamma \text{Tr}(A \mathbb{E}[XX^T])^2 \\ &= 2\text{Tr}(A^T HA) - \gamma \mathbb{E}[X^T AX]^2 \\ &> 0. \end{aligned}$$

Then, let $(u_i)_i \in \mathbb{R}^{d \times d}$ an orthogonal basis that diagonalizes H and λ_i the eigenvalues associated with each eigenvector. Then, taking $A = \sum_i u_i u_i^T$, we get

$$\begin{aligned} 2\text{Tr}(A^T HA) - \gamma \text{Tr}(AH)^2 &= 2\text{Tr}\left(\sum_{i,j} u_i u_i^T H u_j u_j^T\right) - \gamma \left(\sum_i u_i^T H u_i\right)^2 \\ &= 2\left(\sum_i \lambda_i\right) - \gamma \left(\sum_i \lambda_i\right)^2 \geq 0, \end{aligned}$$

so that

$$\gamma < \frac{2}{\sum_i \lambda_i} = \frac{2}{\text{Tr}(H)}.$$

□

Lemma 4. *Let $\gamma > 0$, we can define $T = H_L + H_R - \gamma M$ as in Section 2. If $\gamma < \frac{2}{\text{Tr}(H)}$, then*

$$I - \gamma T \succ -I.$$

and if we are in dimension 1,

$$I - \gamma T \succeq 0$$

Proof. The Lemma is equivalent to $\forall A \in \mathcal{S}(\mathbb{R}^d), A \neq 0 \Rightarrow \langle A, (2I - \gamma T)A \rangle > 0$.

If we are in dimension $d = 1$, then we have $I - \gamma T = 1 - 2\gamma h + \gamma m^2$ where $h = \mathbb{E}[X^2]$ and $m = \mathbb{E}[X^4] \geq h^2$ so that $I - \gamma T \geq (1 - \gamma h)^2 \geq 0$.

Let now assume we are in dimension two or more. Let $A \in \mathcal{S}(\mathbb{R}^d)$ with $A \neq 0$. Let $P \in \mathbb{R}^{d \times d}$ be an orthogonal matrix such that $PHP^{-1} = D$ where D is diagonal with eigenvalues ordered in decreasing order, with $\lambda_i = D_{i,i}$ and $\lambda_1 = L$. We will denote $U = PAP^{-1} = PAP^T$.

$$\begin{aligned}
\langle A, (2I - \gamma T)A \rangle &= \text{Tr}(A^T(2I - \gamma T)A) \\
&= 2\text{Tr}(A^T A) - 2\gamma \text{Tr}(A^T H A) + \gamma^2 \mathbb{E}[(X^T A X)^2] \\
&\geq 2\text{Tr}(A^T A) - 2\gamma \text{Tr}(A^T H A) + \gamma^2 \mathbb{E}[(X^T A X)]^2 \\
&= 2\text{Tr}(A^T A) - 2\gamma \text{Tr}(A^T H A) + \gamma^2 \text{Tr}(A H)^2 \\
&= 2\text{Tr}(U^T U) - 2\gamma \text{Tr}(U^T D U) + \gamma^2 \text{Tr}(U D)^2 \\
&= \sum_{i,j=1}^d 2U_{i,j}^2 - 2\gamma U_{i,j}^2 \lambda_i + \gamma^2 U_{i,i} U_{j,j} \lambda_i \lambda_j \\
&= \left(\sum_{i \neq j} 2U_{i,j}^2 (2 - \gamma(\lambda_i + \lambda_j)) \right) + \sum_{i=1}^d 2U_{i,i}^2 - 2\gamma U_{i,i}^2 \lambda_i + \gamma^2 \left(\sum_{i=1}^d U_{i,i} \lambda_i \right)^2.
\end{aligned}$$

The first sum immediately defines a definite positive form over the subspace generated by $(U_{i,j})_{i \neq j}$ as $\gamma < \frac{2}{\lambda_i + \lambda_j}$ for all $i \neq j$. The second part also defines a bilinear form over the orthogonal subspace generated by $(U_{i,i})_{1 \leq i \leq d}$. $2I - \gamma T$ is definite positive if and only if those two forms are definite positive. We will introduce $x_i = U_{i,i}$ so that the second form is given by $x^T G x$ where $G = 2I - 2\gamma \text{Diag}(\Lambda) + \gamma^2 \Lambda \Lambda^T$, with $\Lambda = (\lambda_i)_{1 \leq i \leq d}$ and $\text{Diag}(\Lambda)$ the diagonal matrix with values from Λ on the diagonal.

We can decompose G as

$$G = \begin{pmatrix} B & \gamma^2 \lambda_1 C^T \\ \gamma^2 \lambda_1 C & D \end{pmatrix},$$

with $B = 2 - 2\gamma \lambda_1 + \gamma^2 \lambda_1$, $C = (\lambda_i)_{2 \leq i \leq d}$ and $D = 2I - 2\gamma \text{Diag}(C) + \gamma^2 C C^T$. Using the Schur complement condition for positive definiteness, we have that $G \succ 0$ if and only if $D \succ 0$ and $B - \gamma^4 \lambda_1^2 C^T D^{-1} C > 0$. We immediately have that $D \succ 0$ as $I - \gamma \text{Diag}(C) \succ 0$, indeed, for all $d \geq i \geq 2$, we have that $\gamma \lambda_i < 1$.

Let us introduce $E = 2I - 2\gamma \text{Diag}(C)$, then we have

$$D^{-1} = E^{-1} - \frac{\gamma^2}{1 + \gamma^2 C^T E^{-1} C} E^{-1} C C^T E^{-1}.$$

We will assume that $\sum_{i=2}^d \lambda_i < \lambda_1$, otherwise one trivially has that $\gamma \lambda_1 < 1$ and $G \succ 0$. Let us denote

$$\begin{aligned}
q &= C^T E^{-1} C \\
&= \sum_{i=2}^d \frac{\lambda_i^2}{2(1 - \lambda_i \gamma)} \\
&\leq \frac{(\sum_{i=2}^d \lambda_i)^2}{2(1 - \gamma \sum_{i=2}^d \lambda_i)} \\
&= \frac{l^2}{2(1 - \gamma l)},
\end{aligned}$$

where $l = \sum_{i=2}^d \lambda_i$. We will take $l = \lambda_1 \alpha$ so that $0 < \alpha < 1$. We have

$$\begin{aligned}
B - \gamma^4 \lambda_1^2 C^T D^{-1} C &= \gamma^2 \lambda_1^2 + 2 - 2\lambda_1 \gamma - \gamma^4 \lambda_1^2 \left(q - \frac{\gamma^2 q^2}{1 + \gamma^2 q} \right) \\
&= \frac{\gamma^2 \lambda_1^2}{1 + \gamma^2 q} - 2\lambda_1 \gamma + 2 \\
&\geq \frac{\gamma^2 \lambda_1^2}{1 + \gamma^2 \frac{l^2}{2(1 - \gamma l)}} - 2\lambda_1 \gamma + 2.
\end{aligned}$$

Denoting $y = \gamma\lambda_1$, we get

$$B - \gamma^4 \lambda_1^2 C^T D^{-1} C = \frac{2y^2(1 - y\alpha)}{2 - 2y\alpha + \alpha^2 y^2} - 2y + 2.$$

Using standard analysis tools, one can show that the last quantity is positive for $0 < y < \frac{2}{1+\alpha}$ and $0 < \alpha < 1$. As a conclusion, G is definite positive and so is $2I - \gamma T$. \square

Lemma 5. *Let $\gamma > 0$, we can define $T = H_L + H_R - \gamma M$ which is symmetric and is stable over $\mathcal{S}(\mathbb{R}^d)$.*

If

$$\forall A \in \mathcal{S}(\mathbb{R}^d), 2\text{Tr}(A^T H A) - \gamma \mathbb{E}[(X^T A X)^2] > 0,$$

or (this second assumption implies the first one)

$$\mathbb{E}[X X^T] - \gamma \mathbb{E}[X^T X X X^T] \succ 0,$$

then

- $\|I - \gamma H\|_{op} < 1$,
- $T \succ 0$,
- $\|I - \gamma T\|_{op} < 1$.

Proof. We should first notice that using Lemma 3, we necessarily have

$$\gamma < \frac{2}{\text{Tr}(H)}. \quad (6.1)$$

We first need, $I - \gamma H \prec I$ which is always true as long as H is invertible (i.e. H is positive). Then we need $I - \gamma H \succ -I$, or $\gamma H \prec 2I$, which means $\gamma < \frac{2}{L}$ where L is H largest eigenvalue. However this is implied by (6.1).

Now, we need $I - \gamma T \prec I$, i.e., $T \succ 0$ (this will also prove T invertible). This is equivalent to

$$\forall A \in \mathcal{S}(\mathbb{R}^d), A \neq 0 \Rightarrow \langle A, T A \rangle > 0$$

Let us compute this term for $A \in \mathcal{S}(\mathbb{R}^d)$ with $A \neq 0$

$$\begin{aligned} \langle A, T A \rangle &= \text{Tr}(A^T (T A)) \\ &= \text{Tr}(A^T A H + A^T H A - \gamma A^T \mathbb{E}[X X^T A X X^T]) \\ &= 2\text{Tr}(A^T H A) - \mathbb{E}[(X^T A X)^2] \quad \text{and we can stop here if we have first assumption} \\ &\geq \text{Tr}(A^T (2H - \gamma \mathbb{E}[X^T X X X^T]) A) \quad \text{using Lemma 2} \end{aligned}$$

A sufficient condition here is that $K = 2H - \gamma \mathbb{E}[X^T X X X^T] \succ 0$. Indeed, let $I = \text{Ker}(A)^\perp$ be the orthogonal space of the kernel of A , which is stable under A as A is symmetric, so we can define A' the restriction of A to I which is invertible. It is of dimension greater than 1 as A is not 0. K defines on I a bilinear symmetric definite positive application K' . Then, $\text{Tr}(A^T K A) = \text{Tr}(A'^T K' A') > 0$ because $A'^T K' A'$ is also symmetric definite positive.

Finally, we want $I - \gamma T \succ -I$. Using Lemma 4, this is a direct consequence of (6.1). \square

Proof of Lemma 1

Let assume $0 < \gamma < \gamma_{\max}$. Lemma 5 already tells us that our operators have good properties as we have $\rho < 1$ and $T \succ 0$. We will now get a finer result in order to have an explicit bound on ρ depending on γ .

As we will be using different values for γ we will explicitly mark the dependency in γ for T by writing $T(\gamma)$. We will only consider $0 < \gamma < \gamma_{\max}$ so that $T(\gamma)$ is positive. We will denote by $L_{T(\gamma)}$ the largest eigenvalue of $T(\gamma)$ and by $\mu_{T(\gamma)}$ its smallest. We then have

$$\rho_T(\gamma) = \max(1 - \gamma \mu_{T(\gamma)}, \gamma L_{T(\gamma)} - 1).$$

One should also notice that the smallest eigenvalue of $H_L + H_R$ is 2μ and the largest $2L$.

We have $T(\gamma_{\max}) \geq 0$ using Lemma 5. For any $0 < \gamma < \gamma_{\max}$ we can define $\alpha = \frac{\gamma}{\gamma_{\max}}$. Then we have

$$\begin{aligned} T(\gamma) &= (1 - \alpha)(H_L + H_R) + \alpha(H_L + H_R) - \alpha\gamma_{\max}M \\ &= (1 - \alpha)(H_L + H_R) + \alpha T(\gamma_{\max}) \\ &\geq (1 - \alpha)(H_L + H_R) \\ &\geq 2(1 - \alpha)\mu, \end{aligned}$$

so that $\mu_{T(\gamma)} \geq 2(1 - \alpha)\mu$.

Using Lemma 4 we have that $T(\gamma_{\max}) \leq \frac{2L}{\gamma_{\max}}$ so that we obtain

$$\begin{aligned} T(\gamma) &= (1 - \alpha)(H_L + H_R) + \alpha T(\gamma_{\max}) \\ &\leq 2(1 - \alpha)L + \frac{2\alpha}{\gamma_{\max}}. \end{aligned}$$

As a consequence if we take

$$\begin{aligned} a(\gamma) &= 1 - 2\alpha\gamma_{\max}(1 - \alpha)\mu \\ b(\gamma) &= 2(1 - \alpha)\alpha\gamma_{\max}L + 2\alpha^2 - 1 \end{aligned}$$

we have $\rho_T(\gamma) = \max(a(\gamma), b(\gamma))$. Besides, if we are in dimension $d = 2$ or more,

$$\begin{aligned} a(\gamma) - b(\gamma) &= 2 - 2\alpha\gamma_{\max}(L + \mu)(1 - \alpha) - 2\alpha^2 \\ &\geq 2 - 4\alpha(1 - \alpha) - 2\alpha^2 \quad \text{as } \gamma_{\max}(L + \mu) \leq 2 \\ &= 2 + 2\alpha^2 - 4\alpha \\ &\geq 0, \end{aligned}$$

so that,

$$\rho_T(\gamma) \leq 1 - 2\gamma\mu\left(1 - \frac{\gamma}{\gamma_{\max}}\right).$$

In dimension $d = 1$, the same result holds as we have $1 - \gamma T \geq 1 - 2\gamma H + \gamma^2 H^2 \geq 0$ so that $\gamma L_{T(\gamma)} - 1 \leq 0$.

We can now look at ρ_H which is given by $\rho_H = \max(1 - \gamma\mu, \gamma L - 1)$.

Let assume we are in dimension 2 or more, then we have $1 - \gamma\mu \geq \gamma L - 1$ so that $\rho_H = 1 - \gamma\mu$. In dimension 1, we have $\rho_H = |1 - \gamma\mu|$. Comparing ρ_H and ρ_T we obtain the result of this Lemma.

Finally, if $\gamma > \gamma_{\max}$, then T has a negative eigenvalue and so $\rho_T > 1$ and $\rho > 1$.

Proof of the theorems

Complete expression of the covariance matrix

Let us recall that we have the update rule

$$\eta_i = (I - \gamma X_i X_i^T)\eta_i + \gamma \varepsilon_i X_i. \quad (6.2)$$

We can then introduce the following matrices

$$M_{k,j} = \left(\prod_{i=k+1}^j (I - \gamma X_i X_i^T) \right)^T \in \mathbb{R}^{d \times d},$$

and by iterating over (6.2) we obtain,

$$\eta_n = \gamma \sum_{k=1}^n M_{k,n} X_k \varepsilon_k + M_{0,n} \eta_0.$$

We have

$$\begin{aligned} \bar{\eta}_n &= \frac{\gamma}{n} \sum_{j=0}^{n-1} \sum_{k=1}^j M_{k,j} X_k \varepsilon_k + \frac{1}{n} \sum_{j=0}^{n-1} M_{0,j} \eta_0 \\ &= \frac{\gamma}{n} \sum_{k=1}^{n-1} \left(\sum_{j=k}^{n-1} M_{k,j} \right) X_k \varepsilon_k + \frac{1}{n} \sum_{j=0}^{n-1} M_{0,j} \eta_0. \end{aligned}$$

One can already see the decomposition between the variance and bias term, one depending only on η_0 and the other on ε .

If we assume that ε_k is independent of X_k , then we can immediately see that when computing $\mathbb{E}[\bar{\eta}_n \bar{\eta}_n^T]$, cross-terms between bias and variance will be zero as they will contain only one ε_k . If that is not true, then extra cross-terms will appear and there is no longer a simple bias/variance decomposition. Let us look at one of the cross terms,

$$\frac{\gamma}{n^2} \mathbb{E} [M_{k,j} X_k \varepsilon_k \eta_0^T M_{0,p}].$$

If $p < k$, then one can immediately notice that $X_k \varepsilon_k$ will be independent from the rest so that the term will be 0, as it is always true that $\mathbb{E}[\varepsilon X] = 0$. If not, X_k will also appear in $M_{0,p}$ as a factor $I - \gamma X_k X_k^T$ so that the term can be expressed as $G(\mathbb{E}[X_k \varepsilon_k \eta_0^T X_k X_k^T])$ where G is a linear operator obtained using the independance of the other X_i and ε_i for $i \neq k$. As a consequence, we can recover a simple decomposition as soon as

$$\forall 1 \leq i, j, k \leq d, \mathbb{E} [X^{(i)} X^{(j)} X^{(k)} \varepsilon] = 0,$$

where $X^{(i)}$ is the i -th component of X .

In any case, because of Minkowski's inequality as noted in Bach and Moulines (2013), we always have that

$$f_n^{\text{total}} - f^* \leq 2(f_n^{\text{bias}} - f^*) + 2(f_n^{\text{variance}} - f^*),$$

so that we are never too far from the true error when assuming X and ε independent.

Proof for the bias term

First, let us assume that $\varepsilon_k = 0$ a.s. Then we have

$$\bar{\eta}_n = \frac{1}{n} \sum_{j=0}^{n-1} M_{0,j} \eta_0,$$

and

$$\begin{aligned}
\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{1}{n^2} \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(\mathbb{E} \left[M_{0,i} \eta_0 \eta_0^T M_{0,i}^T + \sum_{j=i+1}^{n-1} M_{0,i} \eta_0 \eta_0^T M_{0,i}^T M_{0,i,j}^T + \sum_{j=0}^{i-1} M_{j,i} M_{0,j} \eta_0 \eta_0^T M_{0,j}^T \right] \right) \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] + \sum_{j=i+1}^{n-1} \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] (I - \gamma H)^{j-i} \right. \\
&\quad \left. + \sum_{j=0}^{i-1} (I - \gamma H)^{i-j} \mathbb{E} [M_{0,j}^T \eta_0 \eta_0^T M_{0,j}] \right) \text{ because of independence assumptions,} \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] + \sum_{j=i+1}^{n-1} \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] (I - \gamma H)^{j-i} \right) \\
&\quad + \frac{1}{n^2} \sum_{j=0}^{n-1} \left(\sum_{i=j+1}^{n-1} (I - \gamma H)^{i-j} \mathbb{E} [M_{0,j}^T \eta_0 \eta_0^T M_{0,j}] \right) \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] \right. \\
&\quad \left. + \sum_{j=i+1}^{n-1} (\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] (I - \gamma H)^{j-i} + (I - \gamma H)^{j-i} \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T]) \right) \\
&\quad \text{by exchanging the role of } i \text{ and } j \text{ in the last equation,} \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] \right. \\
&\quad + \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] ((I - \gamma H) - (I - \gamma H)^{n-i}) (\gamma H)^{-1} \\
&\quad \left. + (\gamma H)^{-1} ((I - \gamma H) - (I - \gamma H)^{n-i}) \mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] \right).
\end{aligned}$$

We only used the fact that X_i and X_j are independent as soon as $i \neq j$, so that we can condition on X_1, \dots, X_i to obtain $M_{1,i} (I - \gamma H)^{j-i}$. Now we need to express $\mathbb{E} [(I - \gamma X_i X_i^T) A (I - \gamma X_i X_i^T)]$ for A some matrix that is independent of X_i . Using the notation we introduced, we have immediately that

$$\begin{aligned}
\mathbb{E} [(I - \gamma X_i X_i^T) A (I - \gamma X_i X_i^T)] &= A - \gamma A H - \gamma H A + \gamma^2 \mathbb{E} [X^T A X X X^T] \\
&= (I - \gamma H_R - \gamma H_L + \gamma^2 M) A \\
&= (I - \gamma T) A.
\end{aligned}$$

Then we have, with \mathcal{F}_{i-1} the σ field generated by X_1, \dots, X_{i-1} ,

$$\begin{aligned}
\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T] &= \mathbb{E} [\mathbb{E} [M_{0,i} \eta_0 \eta_0^T M_{0,i}^T | \mathcal{F}_{i-1}]] \\
&= \mathbb{E} [\mathbb{E} [(I - \gamma X_i X_i^T) M_{0,i-1} \eta_0 \eta_0^T M_{0,i-1}^T (I - \gamma X_i X_i^T) | \mathcal{F}_{i-1}]] \\
&= \mathbb{E} [(I - \gamma T) M_{0,i-1} \eta_0 \eta_0^T M_{0,i-1}^T] \\
&= (I - \gamma T) \mathbb{E} [M_{0,i-1} \eta_0 \eta_0^T M_{0,i-1}^T].
\end{aligned}$$

and by iterating this process, we obtain

$$\begin{aligned}
\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{1}{n^2} \sum_{i=0}^{n-1} (I - \gamma T)^i E_0 \\
&\quad + ((I - \gamma T)^i E_0) ((I - \gamma H) - (I - \gamma H)^{n-i}) (\gamma H)^{-1} \\
&\quad + (\gamma H)^{-1} ((I - \gamma H) - (I - \gamma H)^{n-i}) ((I - \gamma T)^i E_0) \\
&= \frac{1}{n^2} \sum_{i=0}^{n-1} \left(I + [(I - \gamma H)_L - (I - \gamma H)_L^{n-i}] (\gamma H_L)^{-1} \right. \\
&\quad \left. + [(I - \gamma H)_R - (I - \gamma H)_R^{n-i}] (\gamma H_R)^{-1} \right) (I - \gamma T)^i E_0.
\end{aligned}$$

Let us define

$$A_n = -\frac{1}{n^2} \sum_{i=0}^{n-1} ((\gamma H_R)^{-1} (I - \gamma H)^{n-i} + (\gamma H_L)^{-1} (I - \gamma H)^{n-i}) ((I - \gamma T)^i E_0)$$

$$\|A_n\|_F \leq \frac{2d}{n\gamma\mu} \rho^n \|E_0\|_F,$$

which is decaying exponentially. We now have

$$\begin{aligned} \mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{1}{n^2} \sum_{i=0}^{n-1} (I + (I - \gamma H_L)(\gamma H_L)^{-1} + (I - \gamma H_R)(\gamma H_R)^{-1}) (I - \gamma T)^i E_0 + A_n \\ &= \frac{1}{\gamma^2 n^2} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} (I - (I - \gamma T)^n) E_0 + A_n. \end{aligned}$$

Again, we have some exponential terms, that we will regroup in B_n with

$$B_n = -\frac{1}{\gamma^2 n^2} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} (I - \gamma T)^n E_0$$

$$\|B_n\|_F \leq \frac{d}{n^2 \gamma^2 \mu_T} \rho_T^n \left(\frac{2}{\mu} - \gamma \right) \|E_0\|_F,$$

and we have

$$\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] = \frac{1}{n^2 \gamma^2} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} E_0 + A_n + B_n.$$

We can bound $A_n + B_n$ by

$$\|A_n + B_n\|_F \leq \frac{d \rho^n \|E_0\|_F}{\gamma n} \left(\frac{2}{\mu} + \frac{1}{\mu_T n \gamma} \left(\frac{2}{\mu} - \gamma \right) \right),$$

which completes the first assertion of Theorem 1.

Proof for the variance term

Let assume now that $\eta_0 = 0$, then we have

$$\bar{\eta}_n = \frac{\gamma}{n} \sum_{k=1}^{n-1} \left(\sum_{j=k}^{n-1} M_{k,j} \right) X_k \varepsilon_k,$$

and

$$\begin{aligned} \mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{\gamma^2}{n^2} \mathbb{E} \left[\sum_{k,l=1}^{n-1} \left(\sum_{j=k}^{n-1} M_{k,j} \right) X_k \varepsilon_k \varepsilon_l X_l^T \left(\sum_{p=l}^{n-1} M_{l,p}^T \right) \right] \\ &= \frac{\gamma^2}{n^2} \mathbb{E} \left[\sum_{k=1}^{n-1} \left(\sum_{j=k}^{n-1} M_{k,j} \right) X_k \varepsilon_k \varepsilon_k X_k^T \left(\sum_{p=k}^{n-1} M_{k,p}^T \right) \right]. \end{aligned}$$

Indeed, we can remove terms where $k \neq l$: if we have for instance $l < k$, then $X_l \varepsilon_l$ will be independent from the rest of the terms and as $\mathbb{E}[X_l \varepsilon_l] = 0$, the term will be 0.

By using mostly the same method as for the bias term, we obtain that

$$\begin{aligned}
\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{\gamma^2}{n^2} \sum_{k=1}^{n-1} \sum_{j=k}^{n-1} (I - \gamma T)^{j-k} \Sigma_0 \\
&\quad + \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1} \left((I - \gamma T)^{j-k} \Sigma_0 \right) \\
&\quad + \left((I - \gamma T)^{j-k} \Sigma_0 \right) \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1} \\
&= \frac{\gamma^2}{n^2} \sum_{j=1}^{n-1} \sum_{k=1}^j (I - \gamma T)^{j-k} \Sigma_0 \\
&\quad + \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1} \left((I - \gamma T)^{j-k} \Sigma_0 \right) \\
&\quad + \left((I - \gamma T)^{j-k} \Sigma_0 \right) \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1} \\
&= \frac{\gamma^2}{n^2} \sum_{j=1}^{n-1} (I - (I - \gamma T)^j) (\gamma T)^{-1} \Sigma_0 \\
&\quad + \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1} (I - (I - \gamma T)^j) (\gamma T)^{-1} \Sigma_0 \\
&\quad + (I - (I - \gamma T)^j) (\gamma T)^{-1} \Sigma_0 \left((I - \gamma H) - (I - \gamma H)^{n-j} \right) (\gamma H)^{-1}.
\end{aligned}$$

As for the bias, we can bound some terms:

$$\begin{aligned}
C_n &= \frac{\gamma^2}{n^2} \sum_{j=1}^{n-1} \left((I - \gamma H)_L^{n-j} (\gamma H_L)^{-1} + (I - \gamma H)_R^{n-j} (\gamma H_R)^{-1} \right) (I - \gamma T)^j (\gamma T)^{-1} \Sigma_0 \\
\|C_n\|_F &\leq \frac{2d}{n\mu\mu_T} \rho^n \|\Sigma_0\|_F.
\end{aligned}$$

Now we have,

$$\begin{aligned}
\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] &= \frac{1}{n^2} \sum_{j=1}^{n-1} (H_L^{-1} + H_R^{-1} - \gamma I) (I - (I - \gamma T)^j) T^{-1} \Sigma_0 + C_n \\
&= \frac{1}{n} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} \Sigma_0 + D_n + C_n,
\end{aligned}$$

where D_n is defined by

$$\begin{aligned}
D_n &= -\frac{1}{n^2} \sum_{j=1}^{n-1} (H_L^{-1} + H_R^{-1} - \gamma I) (I - \gamma T)^j T^{-1} \Sigma_0 \\
&= -\frac{1}{\gamma n^2} (H_L^{-1} + H_R^{-1} - \gamma I) (I - \gamma T) T^{-2} \Sigma_0 + D'_n.
\end{aligned}$$

D'_n are again exponentially decreasing terms:

$$\begin{aligned}
D'_n &= \frac{1}{\gamma n^2} (H_L^{-1} + H_R^{-1} - \gamma I) (I - \gamma T)^n T^{-2} \Sigma_0 \\
\|D'_n\|_F &\leq \frac{d}{\gamma^2 \mu_T^2 n} \left(\frac{2}{\mu} - \gamma \right) \rho_T^n \|\Sigma_0\|_F,
\end{aligned}$$

so that we have

$$\mathbb{E} [\bar{\eta}_n \bar{\eta}_n^T] = \frac{1}{n} (H_L^{-1} + H_R^{-1} - \gamma I) T^{-1} \Sigma_0 - \frac{1}{\gamma n^2} (H_L^{-1} + H_R^{-1} - \gamma I) (I - \gamma T) T^{-2} \Sigma_0 + C_n + D'_n. \tag{6.3}$$

We can bound $C_n + D'_n$ by

$$\|C_n + D'_n\|_F \leq \frac{d\rho^n \|\Sigma_0\|_F}{n} \left(\frac{1}{n\gamma\mu_T^2} \left(\frac{2}{\mu} - \gamma \right) + \frac{2}{\mu\mu_T} \right).$$

This concludes the proof of Theorem 2.

References

- Bach, F. and E. Moulines (2011). Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. In *Adv. NIPS*.
- Bach, F. and E. Moulines (2013). Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$. In *Adv. NIPS*.
- Bershad, N. (1986). Analysis of the normalized LMS algorithm with gaussian inputs. *Speech and Signal Processing, IEEE Transactions on Acoustics* 34(4), 793–806.
- Bottou, L. and Y. Le Cun (2005). On-line learning for very large data sets. *Applied Stochastic Models in Business and Industry* 21(2), 137–151.
- Bousquet, O. and L. Bottou (2008). The tradeoffs of large scale learning. In *Adv. NIPS*.
- Fabian, V. (1968). On asymptotic normality in stochastic approximation. *The Annals of Mathematical Statistics* 39(4), 1327–1332.
- Kanamori, T. and H. Shimodaira (2003). Active learning algorithm using the maximum weighted log-likelihood estimator. *Journal of statistical planning and inference* 116(1), 149–162.
- Macchi, O. (1995). *Adaptive processing: The least mean squares approach with applications in transmission*. Wiley West Sussex.
- Nedic, A. and D. Bertsekas (2000). Convergence rate of incremental subgradient algorithms. *Stochastic Optimization: Algorithms and Applications*, 263–304.
- Needell, D., N. Srebro, and R. Ward (2013). Stochastic gradient descent and the randomized kaczmarz algorithm. Technical Report 1310.5715, arXiv.
- Nemirovski, A. S. and D. B. Yudin (1983). *Problem complexity and method efficiency in optimization*. Wiley & Sons.
- Nesterov, Y. (2004). *Introductory Lectures on Convex Optimization: a Basic Course*. Kluwer Academic Publishers.
- Nesterov, Y. (2012). Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization* 22(2), 341–362.
- Perronnin, F., Z. Akata, Z. Harchaoui, and C. Schmid (2012). Towards good practice in large-scale learning for image classification. In *Proc. CVPR*.
- Polyak, B. T. and A. B. Juditsky (1992, July). Acceleration of stochastic approximation by averaging. *SIAM J. Control Optim.* 30(4), 838–855.
- Ruppert, D. (1988). Efficient estimations from a slowly convergent Robbins-Monro process. Technical report, Cornell University Operations Research and Industrial Engineering.
- Schmidt, M., N. L. Roux, and F. Bach (2013). Minimizing finite sums with the stochastic average gradient. Technical Report 00860051, HAL.
- Shalev-Shwartz, S. and T. Zhang (2013). Stochastic dual coordinate ascent methods for regularized loss minimization. *JMLR* 14, 567–599.
- Toulis, P., J. Rennie, and A. M. Airolidi (2014). Statistical analysis of stochastic gradient methods for generalized linear models. In *Proc. ICML*.
- Zhao, P. and T. Zhang (2014). Stochastic optimization with importance sampling. Technical Report 1401.2753, arXiv.