

Stochastic Differential Equations models for Least-Squares Stochastic Gradient Descent

Adrien Schertzer

*ENS Lyon and CNRS,
Lyon, France*

ASCHERTZ@UNI-BONN.DE

Loucas Pillaud-Vivien

*Ecole des Ponts ParisTech - Cermics
Marne-la-Vallée, France*

LOUCAS.PILLAUD-VIVIEN@ENPC.FR

Editor: Maxim Raginsky

Abstract

We study the dynamics of a continuous-time model of stochastic gradient descent (SGD) for the least-square problem. Indeed, pursuing the work of Li et al. (2019), we analyze stochastic differential equations (SDEs) that model SGD either in the case of the *training loss* (finite samples) or the population one (online setting). A key qualitative feature of the dynamics is the existence of a perfect interpolator of the data, irrespective of the sample size. In both scenarios, we provide precise, non-asymptotic rates of convergence to the (possibly degenerate) stationary distribution. Additionally, we describe this asymptotic distribution, offering estimates of its mean, deviations from it, and a proof of the emergence of heavy-tails related to the step-size magnitude. Numerical simulations supporting our findings are also presented.

Keywords: SGD, Least-Squares, Lyapunov method, Heavy Tails, SDEs.

1. Introduction

Stochastic gradient descent (SGD) is the workhorse of any large-scale machine learning pipeline. Described more than seventy years ago as a versatile stochastic algorithm (Robbins and Monro, 1951), it has been studied thoroughly since then (e.g. Pagès (2001); Benaïm (2006)), and applied with success as an efficient computational and statistical device for large-scale machine learning (Bottou and Bousquet, 2008). Yet, since the emergence of deep neural networks (DNNs), some new light has been shed on the use of the algorithm: among others, the role of SGD’s noise in the good generalization performances of DNNs (Zhang et al., 2021) has been described empirically (Xing et al., 2018), while the particular shape of its covariance is expected to play a predominant role in its dynamics (Andriushchenko et al., 2023).

In this direction, Li et al. (2019) proposed to approximate the SGD dynamics thanks to a stochastic differential equation (SDE) whose noise’s covariance matches the one of SGD. Since then, many works have leveraged this continuous perspective, attempting to better describe some of SGD’s phenomenology: among others, the role of the step size (HaoChen et al., 2021), the escape time and direction from local minimizers (Xie et al., 2020), study of the invariant distribution (Chaudhari and Soatto, 2018) or the heavy-tail phenomenon (Gur-

buzbalaban et al., 2021). However, lured by the strong analytical tools that offer SDE models, some misconceptions about the basic nature of SGD’s noise *in the machine learning context* have also emerged. Among others, an important message of Wojtowytsch (2023) is to recall that SGD’s noise has a particular shape and intensity that make SGD far from a *Langevin* type of dynamics where isotropic noise is added to the gradient at each step. In particular, if the data model can be perfectly interpolated, the invariant measure of SGD could be largely degenerate without the need for step-size decay: this is the case in the overparametrized regime (Berthier et al., 2020; Varre et al., 2021).

In this article, we take a step back from general-purpose studies of SGD and focus on the specific case of least-squares, where the predictor is a linear function of fixed features. Needless to say that linear predictors do not lead to state-of-the-art performance in most modern tasks, yet, the abundant literature on the *neural tangent kernel* (Jacot et al., 2018) recalled us that we have still to understand better linear models. Moreover, even if the article is not written within the setting of *kernel methods* for the sake of simplicity, every assumption is made in order to easily adapt to this setting, replacing \mathbb{R}^d equipped with its canonical Euclidean structure by an abstract reproducing kernel Hilbert space.

Purpose and contributions. The aim of the present paper is to show that, with a proper model, SDEs offer nice analytical tools that help to clean the analysis while capturing the qualitative as well as the quantitative essence of SGD’s dynamics. We present a wide range of results on both online and empirical SGD, demonstrating through quantitative analysis that the key difference lies in the model’s ability to achieve perfect interpolation. We also present systemically non-asymptotic rates of convergence in all the settings we present, either in ℓ_2 -norm in the interpolation regime (Theorems 4 and 13), or in Wasserstein distance (towards the invariant distribution) in the case of a noisy system (Theorems 7 and 15). In the latter case, we further investigate the invariant distribution: we pinpoint its location (Proposition 8), and more importantly, we demonstrate in Proposition 10 that although there is no heavy tail phenomenon in finite time, it emerges asymptotically. We finally address convergence of *variance reduction techniques* like time-averaging (Proposition 11) and decay of step-size (Proposition 12). Throughout the article, we try to present some theoretical background on the study of SDEs: among others, the use of Lyapunov potentials and of coupling methods that enables us to quantitatively study the speed of convergence to equilibrium.

Further related work. The formal links between stochastic gradient descent and its continuous model are studied in Li et al. (2019) on the theoretical side, where a weak error analysis is provided, and in Li et al. (2021) on the experimental side. In a similar article in spirit, Fontaine et al. (2021) provides convergence results on general convex function, but *with a systematic polynomial step-size decay*. Some results implying the use of SDEs to study the influence of the noise on the *implicit bias* the algorithm are given in the least-squares case in Ali et al. (2020), and is an active topic in general (Ziyin et al., 2021; Wu et al., 2018; Pesme et al., 2021; Pillaud-Vivien et al., 2022). Quantitative studies of the invariant distribution that focus on the particular shape and intensity of the noise covariance include Mori et al. (2022); Wojtowytsch (2024). Note that when the step-size is properly re-scaled with respect to the dimension, SGD converges in the high-dimensional limit to an SDE that is similar to the one we study here (Ben Arous et al., 2022), which is called *homogenized*

SGD in the least-square context (Paquette et al., 2022). Power laws of convergence toward stationarity related to the eigenvalue decay of the covariance matrix (capacity condition) are ubiquitous in statistics (Caponnetto and De Vito, 2007) and the study of SGD for least-squares (Dieuleveut and Bach, 2016; Berthier et al., 2020; Pillaud-Vivien et al., 2018; Cui et al., 2021; Bordelon and Pehlevan, 2021). Finally, the *heavy-tail phenomenon*, has been re-discovered lately as an interesting feature of SGD with multiplicative noise (Hodgkinson and Mahoney, 2021; Gurbuzbalaban et al., 2021) and more recently in the context of SDE in the work of Jiao and Keller-Ressel (2024).

Organisation of the paper. In Section 2, we present the general setup of SGD both in the population and empirical cases as well as the possibility of the data *in both cases* to be fully interpolated. Section 3 explains the relevance of building a consistent SDE model of SGD and recalls technical details related to SDEs. In Section 4, the results regarding the dynamics of SGD in the training case are given, both in the interpolation regime (Section 4.1) and in the noisy regime (Section 4.2). Section 5 is built similarly to the previous one and is devoted to online SGD. The proofs are postponed to the Appendix, for which precise references can be found in the main text.

2. Setup: Stochastic Gradient Descent on Least Squares Problems

In this section, we introduce the least-square problem that we consider throughout the article, putting emphasis on the difference between empirical distributions (training loss) and true ones (population loss). However, the central argument of the article is that, aside from the known differences between empirical and test cases, the primary qualitative distinction hinges on whether the possibility of loss can be zero (interpolation), irrespective of the discrete nature of the distribution. Stochastic gradient descent is introduced at the end of the section.

2.1 The least-square problem: population and empirical losses.

We consider a regression problem with input/output pair $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ distributed according to a joint probability law $\rho \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R})$. To learn the rule linking inputs to outputs, we take a linear family of predictors $\{f_\theta : x \mapsto \langle \theta, x \rangle, \theta \in \mathbb{R}^d\}$ and aim to minimize the average *square loss* on the penalty $\ell(\theta, (X, Y)) := \frac{1}{2}(\langle \theta, X \rangle - Y)^2$

$$L(\theta) := \frac{1}{2} \mathbb{E}_{(X, Y) \sim \rho} \left[(\langle \theta, X \rangle - Y)^2 \right]. \quad (1)$$

Remark 1 (Link with RKHS) *Here, the predictor family consists of linear functions of the input data $x \in \mathbb{R}^d$. While for the sake of clarity, we will keep this linear family throughout the article, note that the same results apply for any family of linear predictors in an abstract reproducing kernel Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$, with feature $\mathbb{R}^d \ni x \mapsto \varphi(x) \in \mathcal{H}$, changing the dot product and family of predictors with the natural structure: $\{f_\theta : x \mapsto \langle \theta, \varphi(x) \rangle_{\mathcal{H}}, \theta \in \mathcal{H}\}$.*

In this article, in order to put emphasis on the possible difference between the two settings, we make a clear distinction between the *population* and the *empirical* cases. To keep the

notation simple, we refer to the population case whenever $\Omega := \text{supp}(\rho)$ is an open set of \mathbb{R}^d . On the contrary, we refer to the empirical case when ρ is a finite sum of atomic measures, i.e. there exist $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^d \times \mathbb{R}$ such that $\rho = \frac{1}{n} \sum_{i=1}^n \delta_{(x_i, y_i)}$. Obviously, the term empirical refers to the fact that, in this case, Eq.(1) can be seen as the training loss

$$L(\theta) = \frac{1}{2n} \sum_{i=1}^n (\langle \theta, x_i \rangle - y_i)^2. \quad (2)$$

In this case, the number $n \in \mathbb{N}^*$ will always denote the number of *observed inputs/outputs pairs* $(x_i, y_i)_{i=1\dots n}$, and we will use the notation $\mathbf{X} = [x_1, \dots, x_n]^\top \in \mathbb{R}^{n \times d}$ to denote the design matrix, and $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ to denote the output vector. With these notations, the training loss can be rewritten as $L(\theta) = \frac{1}{2n} \|\mathbf{X}\theta - \mathbf{y}\|_2^2$.

Note that even though for the sake of clarity we will sometimes treat them separately, these cases fit into the same framework and notation.

2.2 Noisy and noiseless settings.

For this, let us introduce the *set of interpolators*:

$$\mathcal{I} = \left\{ \theta \in \mathbb{R}^d, \text{ such that } L(\theta) = 0 \right\}. \quad (3)$$

We distinguish between the two following settings.

Noisy setting $\mathcal{I} = \emptyset$. In this case, the loss L is strictly positive, i.e. there exists $\sigma > 0$ such that $L(\theta) \geq \sigma^2$ for all $\theta \in \mathbb{R}^d$. The two following examples show that this situation typically arises for different reasons in the population and in the empirical cases. In the former, this is an effect of the noise model on the output, whereas in the latter, this is only a consequence of the underparametrization.

Example 1 (Noisy model) Assume that there exists $\theta_* \in \mathbb{R}^d$ and $\xi \in \mathbb{R}$ a random variable independent of X of mean zero and variance $2\sigma^2$, such that $Y = \langle \theta_*, X \rangle + \xi$. Then

$$L(\theta) = \frac{1}{2} \left\| \Sigma^{1/2}(\theta - \theta_*) \right\|^2 + \frac{1}{2} \mathbb{E}(\xi^2) \geq \sigma^2,$$

where we use the notation for the input covariance matrix: $\Sigma := \mathbb{E}_\rho[XX^\top] \in \mathbb{R}^{d \times d}$.

Example 2 (Underparametrized setting) Consider the underparametrized regime for which $n > d$, and assume that each *i.i.d.* input/output couple $(x_i, y_i)_{i=1\dots n}$ come from independent distributions that have densities (i.e. for all $i \in \llbracket 1, n \rrbracket$, x_i and y_i are independent and are distributed according to laws that are absolutely continuous with respect to Lebesgue). Then, we have that almost surely $\mathcal{I} = \emptyset$.

Noiseless setting $\mathcal{I} \neq \emptyset$. This means that there exists at least one perfect linear interpolator of the model. In the population setting, this corresponds to the strong assumption that the model is well-specified and noiseless (formally $\xi = 0$, if we refer to Example 1). Notably, this regime has received considerable attention recently to model the large expressive power of neural network (Ali et al., 2020; Berthier et al., 2020; Vaswani et al., 2019). Yet, in the empirical case, this typically and simply corresponds to the overparametrized regime $d \geq n$.

Example 3 (Overparametrized setting) Consider the overparametrized regime for which $d \geq n$, and assume that $((x_1, y_1), \dots, (x_n, y_n))$ are i.i.d. samples drawn from a distribution that has a density. Then, we have that the zero-loss set is the affine set $\mathcal{I} = \theta_* + \ker(\mathbf{X})$, where θ_* is any element of \mathcal{I} . Furthermore, $\dim \mathcal{I} \geq d - n \geq 0$ almost surely.

2.3 Stochastic gradient descent

Stochastic gradient descent (SGD) aims at minimizing a function through unbiased estimates of its gradient. While this method has been developed for a different purpose in the early 50's (Robbins and Monro, 1951), it is remarkable how SGD fits perfectly the modern large-scale machine learning framework (Bottou et al., 2018). Indeed, SGD iterative procedure consists in sampling at each time $t \in \mathbb{N}^*$ an independent draw $(x_t, y_t) \sim \rho$ and updating the predictor θ with respect to the local gradient calculated on this sample:

$$\theta_{t+1} = \theta_t - \gamma \nabla_{\theta} \ell(\theta_t, (x_t, y_t)), \quad (4)$$

where $\gamma > 0$ is the step size. Even if the population and empirical cases fall under the same general framework of stochastic approximation, let us detail the setups for these.

The population case: online SGD. The population case corresponds to what is referred to as *online SGD* in the literature (Bottou and Bousquet, 2008). In this case, for each time $t \in \mathbb{N}^*$, we have an independent draw $(x_t, y_t) \sim \rho$, and if we denote $\mathcal{F}_t = ((x_k, y_k), k \leq t)$ the natural adapted filtration, then $\mathbb{E}[\nabla_{\theta} \ell(\theta_t, (x_t, y_t)) | \mathcal{F}_{t-1}] = \nabla L(\theta_t)$. That is, $\nabla_{\theta} \ell(\theta_t, (x_t, y_t))$ is an unbiased estimate of the true gradient of the risk, and the recursion can be written in explicit form

$$\theta_{t+1} = \theta_t - \gamma (\langle \theta_t, x_t \rangle - y_t) x_t. \quad (5)$$

Remark that even though at iteration t , it has seen only t samples, the SGD recursion directly optimizes the population loss L . This is due to the fact that as long as the recursion progresses, we have access to *fresh samples* from the distribution ρ . This can be seen informally as the case $n = +\infty$ and is in contrast to the finite n empirical case. A convenient way to present this dynamics is to force the apparition of the true gradient and rewrite the rest as a martingale increment. Indeed, Eq. (5) reads:

$$\begin{aligned} \theta_{t+1} &= \theta_t - \gamma \nabla_{\theta} L(\theta_t) + \gamma (\nabla_{\theta} L(\theta_t) - (\langle \theta_t, x_t \rangle - y_t) x_t) \\ &= \theta_t - \gamma \mathbb{E}_{\rho} [(\langle \theta_t, X \rangle - Y) X] + \gamma m(\theta_t, (x_t, y_t)), \end{aligned} \quad (6)$$

where $m(\theta_t, (x_t, y_t)) := \mathbb{E}_{\rho} [(\langle \theta_t, X \rangle - Y) X] - (\langle \theta_t, x_t \rangle - y_t) x_t$.

The empirical case: training SGD. In this case, the main difference is that we consider a finite number of samples, so the SGD recursion iterates over them uniformly at random. Mathematically, this corresponds to taking, at each time $t \in \mathbb{N}^*$ an independent sample of the uniform distribution $i_t \sim \text{Unif}(\{1, \dots, n\})$ and consider the unbiased estimate $\nabla_{\theta} \ell(\theta_t, (x_{i_t}, y_{i_t}))$ of the true gradient of the training loss Eq.(2). For this, the adapted filtration writes $\mathcal{F}_t = ((x_k, y_k)_{k \leq n}, (i_k)_{k \leq t})$ and the recursion reads:

$$\theta_{t+1} = \theta_t - \gamma (\langle \theta_t, x_{i_t} \rangle - y_{i_t}) x_{i_t}. \quad (7)$$

Despite the fact that this dynamics looks very similar, the important difference is that, the batch of samples being fixed, the dynamics can select several times the same pair (x_k, y_k) . This is the reason why, informally, after $t \geq \Theta(n)$ iterations, the training dynamics Eq.(7) will deviate from the online SGD presented in Eq.(5). To end this section, let us also rewrite this dynamics with respect to the gradient descent plus the martingale increments:

$$\begin{aligned}\theta_{t+1} &= \theta_t - \gamma \nabla_{\theta} L(\theta_t) + \gamma (\nabla_{\theta} L(\theta_t) - (\langle \theta_t, x_{i_t} \rangle - y_{i_t}) x_{i_t}) \\ &= \theta_t - \gamma \frac{1}{n} \sum_{i=1}^n (\langle \theta_t, x_i \rangle - y_i) x_i + \gamma m(\theta_t, (x_{i_t}, y_{i_t})) ,\end{aligned}\tag{8}$$

where $m(\theta_t, (x_t, y_t)) := \frac{1}{n} \sum_{i=1}^n (\langle \theta_t, x_i \rangle - y_i) x_i - (\langle \theta_t, x_{i_t} \rangle - y_{i_t}) x_{i_t}$. Note that these equations are, in fact, exactly the same as those presented in the population case, simply considering that the distribution is an empirical one, i.e. $\rho = \frac{1}{n} \sum_{i=1}^n \delta_{(x_i, y_i)}$. Nonetheless, we decided to present them explicitly for the sake of clarity.

3. Continuous model of SGD

In this section, we present stochastic differential equation (SDE) models for SGD. First, we provide general necessary conditions for accurately modeling SGD. We then instantiate these conditions more precisely.

3.1 The requirement of a SDE model

As mentioned above, the decomposition of the SGD recursion as in equations (6),(8) is a generic feature for stochastic descent (Benaïm, 2006), as it has led to the celebrated *ODE method* (Harold et al., 1997) to study stochastic approximation of this type. Going further, this decomposition between a *drift* ∇L and *local martingale term* $m(\cdot, (x, y))$ is reminiscent of the decomposition that occurs for Itô processes, i.e. solution to a SDE of the type

$$d\theta_t = b(t, \theta_t)dt + \sqrt{\gamma} \sigma(t, \theta_t)dB_t,\tag{9}$$

where $(B_t)_{t \geq 0}$ is a Brownian motion of \mathbb{R}^d . These types of models have been extensively studied in the literature (Khasminskii, 2011), as they not only possess strong modeling capabilities but also provide useful tools such as Itô calculus for mathematical analysis. One of the aims of this article is to link the SGD dynamics to processes that are exemplary in the SDE literature. In order to establish this link, the first question we need to address is the following

What SDE model fits well with the SGD dynamics?

This question has received great attention in the last decade, and a good principle to answer this is to turn to *stochastic modified equations* (Li et al., 2019). This is a natural way to build SDE models since they are consistent in the infinitesimal step-size limit with SGD. More precisely, this is a weak approximation result for the marginal laws at discrete times in the small-step limit. This does not imply pathwise convergence, but only convergence of the marginal distributions at each discrete time. In order to build such models, there are two requirements:

- (i) The drift term $b(t, \theta_t)$ should match $-\nabla L(\theta_t)$.
- (ii) The noise factor σ should have the same covariance as the local martingale m , i.e.

$$\sigma(t, \theta_t)\sigma(t, \theta_t)^\top = \mathbb{E}_\rho \left[m(\theta_t, (x_t, y_t))m(\theta_t, (x_t, y_t))^\top \mid \mathcal{F}_{t-1} \right]. \quad (10)$$

Besides technical assumptions, these are the two requirements presented in Li et al. (2019, Theorem 3) to show that the SDE model is *consistent* in the small step-size limit with the SGD recursion. Going beyond the approximation concerns tackled by (i) and (ii), it has recently been observed that the SGD noise carries a specific shape that the SDE model should carry as well (Wojtowytsch, 2023; Pesme et al., 2021). This requirement has a more *qualitative nature* but it is important to fully capture the essence of the SGD dynamics

- (iii) The noise term $\sigma(t, \theta_t)dB_t$ should span the same space as $m(\theta_t, (x_t, y_t))$.

We will see below, that, in order to build a SDE model of SGD, this third requirement is particularly important in the empirical case, where the noise a strong degeneracy.

3.2 Explicit form of the SDE models

Let us first explicitly write the multiplicative noise factor $\sigma(t, \theta_t)$ in the population and empirical cases. Then, we derive the expression of the SDE models that we analyze later.

The population case. In the population case, the calculation has already been made in Ali et al. (2020) and, defining $r_X(\theta) = \langle \theta_t, X \rangle - Y \in \mathbb{R}$, the residual random variable, we have

$$\sigma(t, \theta_t)\sigma(t, \theta_t)^\top = \left(\mathbb{E}_\rho \left[r_X(\theta_t)^2 X X^\top \right] - \mathbb{E}_\rho [r_X(\theta_t)X] \mathbb{E}_\rho [r_X(\theta_t)X]^\top \right).$$

In this case, there is no geometric specificity of the noise and we choose $\sigma(t, \theta_t)$ as the PSD square root of the right hand side, that is, we define, for all $\theta \in \mathbb{R}^d$,

$$\sigma(\theta) := \left(\mathbb{E}_\rho \left[r_X(\theta)^2 X X^\top \right] - \mathbb{E}_\rho [r_X(\theta)X] \mathbb{E}_\rho [r_X(\theta)X]^\top \right)^{1/2} \in \mathbb{R}^{d \times d}, \quad (11)$$

and we study the following SDE model:

$$d\theta_t = -\mathbb{E}_\rho [(\langle \theta_t, X \rangle - Y) X] dt + \sqrt{\gamma} \sigma(\theta_t) dB_t, \quad (12)$$

where $(B_t)_{t \geq 0}$ is a Brownian motion of \mathbb{R}^d and $\sigma \in \mathbb{R}^{d \times d}$ is given by (11).

The empirical case. The empirical case is, from a geometric perspective, a bit more subtle. Indeed, as can be seen directly from Eq.(7), iterates SGD remain in the low-dimensional space of dimension at most n : $\theta_0 + \text{span}(x_1, \dots, x_n) \subset \mathbb{R}^d$. Hence, it is primordial for a good SDE model that the noise reflects this degenerate structure. Let us apply for the SDE noise factor calculation. Let us define $r_x(\theta) := (\langle \theta_t, x_1 \rangle - y_1, \dots, \langle \theta_t, x_n \rangle -$

$y_n)^\top \in \mathbb{R}^n$, the vector of the residuals, then

$$\begin{aligned} \sigma(t, \theta_t) \sigma(t, \theta_t)^\top &= \left(\frac{1}{n} \mathbf{X}^\top \text{diag}(r_x(\theta_t))^2 \mathbf{X} - \frac{1}{n^2} \mathbf{X}^\top r_x(\theta_t) r_x(\theta_t)^\top \mathbf{X} \right) \\ &= \frac{1}{n} \mathbf{X}^\top \left(\text{diag}(r_x(\theta_t))^2 - \frac{1}{n} r_x(\theta_t) r_x(\theta_t)^\top \right) \mathbf{X} \\ &= \frac{1}{n} \mathbf{X}^\top \left(\text{diag}(r_x(\theta_t)) - \frac{1}{n} r_x(\theta_t) \mathbf{1}^\top \right) \left(\text{diag}(r_x(\theta_t)) - \frac{1}{n} r_x(\theta_t) \mathbf{1}^\top \right)^\top \mathbf{X}, \end{aligned}$$

where for any $\theta \in \mathbb{R}^d$, we define $\text{diag}(\theta) \in \mathbb{R}^{d \times d}$ as the diagonal matrix with the coordinates of θ as diagonal entries and $\mathbf{1} \in \mathbb{R}^n$ the vector of ones. In the case where $n < d$, the overall covariance matrix is degenerate as its rank is at most n : this is because the “noise of SGD” naturally lies in $\text{span}(x_1, \dots, x_n)$. Hence, it is important to keep this degeneracy and this is the reason why we choose *the* square root that preserves it:

$$\sigma(\theta_t) := \frac{1}{\sqrt{n}} \mathbf{X}^\top R_x(\theta_t) \in \mathbb{R}^{d \times n},$$

where we defined $R_x(\theta) := \text{diag}(r_x(\theta)) - \frac{1}{n} r_x(\theta) \mathbf{1}^\top \in \mathbb{R}^{n \times n}$. With these notations, we study the following SDE model:

$$d\theta_t = -\frac{1}{n} \mathbf{X}^\top (\mathbf{X} \theta_t - y) dt + \sqrt{\frac{\gamma}{n}} \mathbf{X}^\top R_x(\theta_t) dB_t, \quad (13)$$

where $(B_t)_{t \geq 0}$ is a Brownian motion of \mathbb{R}^n .

Initial condition and moments. Initialization will be taken at some $\theta_0 \in \mathbb{R}^d$ which can be considered as a random variable. Standard choices for the law of θ_0 include the standard Gaussian of \mathbb{R}^d , or a dirac measure on some vector θ_0 , e.g. $\theta_0 = 0$. In both cases, its law ρ_0 has moments of all orders, and since drift and multiplicative noise of the SDEs (12)-(13) have at most linear growth, Theorem 3.5 of Khasminskii (2011) shows that the marginal laws $(\rho_t)_{t \geq 0}$ all have moments at all orders at any time $t \geq 0$.

Now that we motivated the SDE models in the previous part, and stated them in Eqs.(12)-(13), we study their convergence property. This is the purpose of the main results that are stated in the two following sections.

4. SGD on the training loss

Recall that this corresponds to the finite data set case with n input/output pairs $(x_i, y_i)_{i=1 \dots n}$, stacked into data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and data output vector $y \in \mathbb{R}^d$. We study in this section the SDE given in equation (13). We assume that all data are bounded, i.e. there is some $K > 0$ such that $\|x_i\|^2 \leq K$ for all $i \in \llbracket 1, n \rrbracket$. Let us introduce first an important element of the model. In this section, we define

Definition 2 Let \mathbf{X}^\dagger denote the pseudo-inverse of the design matrix \mathbf{X} . We define

$$\theta_* = \mathbf{X}^\dagger y + (I - \mathbf{X}^\dagger \mathbf{X}) \theta_0. \quad (14)$$

Note that, in generic cases, in the underparametrized case for which $d \leq n$, we have $\mathbf{X}^\dagger = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$. Then, $\mathbf{X}^\dagger \mathbf{X} = I$, and $\theta_* = \mathbf{X}^\dagger \mathbf{y}$ is the *ordinary least squares estimator* and does not depend on θ_0 . In the generic overparameterized case for which $n \leq d$, we have $\mathbf{X}^\dagger = \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1}$, and hence $\mathbf{X} \theta_* = \mathbf{X} \mathbf{X}^\dagger \mathbf{y} + (\mathbf{X} - \mathbf{X} \mathbf{X}^\dagger \mathbf{X}) \theta_0 = \mathbf{y}$, which shows that $\theta_* \in \mathcal{I}$.

Finally, for both cases, we define $\Sigma := \frac{1}{n} \mathbf{X}^\top \mathbf{X} \in \mathbb{R}^{d \times d}$ the design covariance matrix. With this notation, the SDE becomes

$$d\theta_t = -\Sigma(\theta_t - \theta_*)dt + \sqrt{\frac{\gamma}{n}} \mathbf{X}^\top R_x(\theta_t) dB_t,$$

where $(B_t)_{t \geq 0}$ is a Brownian motion of \mathbb{R}^n .

Comparison with standard processes. As already said, the dynamics in the noiseless and noisy cases are of different natures because of the possibility to cancel the multiplicative noise term. Indeed, for the sake of clarity, assuming $n = d$ and $\mathbf{X}/\sqrt{n} = \Sigma = I_d$. If the noise can cancel, $R_x(\theta_t)$ has the shape of a linear term like $\text{diag}(\theta - \theta_*)$, and each the coordinate of the difference $\eta = \theta - \theta_*$ follows a one-dimensional *Geometric Brownian motion* $d\eta_t = -\eta_t dt + \sqrt{\gamma} \eta_t dB_t$, for which it is known that $\eta_t \rightarrow 0$ almost surely, which corresponds to $\theta_t \rightarrow \theta_*$. This comparison is the governing principle of the analysis of this setup. Otherwise, if the noise cannot cancel and is strictly lower bounded, then, under the same proxy, the movement of each coordinate of η resembles the SDE $d\eta_t = -\eta_t dt + \sqrt{\gamma} \sqrt{\eta_t^2 + \sigma^2} dB_t$ which looks like a Ornstein-Uhlenbeck process if $\eta \ll \sigma$, but with a noise that has a multiplicative part $\eta_t dB_t$ when $\eta \gg \sigma$. These are known in one dimension under the name of *Pearsons diffusions* (Forman and Sørensen, 2008) and exhibit stationary distribution with heavy tails.

The difference between these two settings is the reason why we divide the results into two different subsections.

4.1 The noiseless case

In this section, we let $(\theta_t)_{t \geq 0}$ follow the dynamics given by Eq.(13) initialized at $\theta_0 \in \mathbb{R}^d$ and assume that $\ker(\mathbf{X}) \neq \{0\}$ and that $\mathcal{I} \neq \emptyset$. Note that this generically occurs in the overparameterized case for which $n \leq d$. Recall that we have defined $\theta_* = \mathbf{X}^\dagger \mathbf{y} + (I - \mathbf{X}^\dagger \mathbf{X}) \theta_0$ in Eq.(14), we now present a geometric interpretation of θ_* in this case,

Lemma 3 *The vector θ_* is the orthogonal projection of θ_0 into \mathcal{I} , that is*

$$\theta_* = \underset{\mathbf{X}\theta = \mathbf{y}}{\operatorname{argmin}} \|\theta - \theta_0\|^2. \tag{15}$$

The proof is deferred to Section A.1 in the Appendix. Note that this fact is often known as the implicit bias of least-squares methods and has statistical consequences studied under the name of “benign overfitting” (Bartlett et al., 2020, 2021).

Remarkably, despite the randomness of the SDE (13), the convergence is almost sure towards θ_* , with explicit rates that we show below. In the following, to simplify the presentation and avoid burdening the reader with slightly different technical conditions in each result, we enforce throughout the more restrictive but unified step-size bound $\gamma < \frac{1}{3K}$. In the following, for any $\alpha > 0$, we define $\Sigma^{-\alpha}$ by

$$\Sigma^{-\alpha} := Q \Lambda^{-\alpha} Q^\top,$$

where $\Sigma = Q\Lambda Q^\top$ is the spectral decomposition of Σ , and $\Lambda^{-\alpha}$ is the diagonal matrix with entries

$$(\Lambda^{-\alpha})_{ii} = \begin{cases} \lambda_i^{-\alpha} & \text{if } \lambda_i > 0, \\ 0 & \text{if } \lambda_i = 0. \end{cases}$$

Theorem 4 *Let $(\theta_t)_{t \geq 0}$ follows the dynamics given by Eq.(13) initialized at $\theta_0 \in \mathbb{R}^d$, then for $\gamma < \frac{1}{3K}$, we have that $(\theta_t)_{t \geq 0}$ converges almost surely to θ_* with the following rates*

(i) **Parametric rate.** *For all $t \geq 0$, we have that*

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \|\theta_0 - \theta_*\|^2 e^{-\mu(2-K\gamma)t}, \quad (16)$$

where $\mu > 0$ is the smallest non zero eigenvalue of Σ .

(ii) **Non-parametric rate.** *For all $t \geq 0$, we have that for all $\alpha > 0$,*

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \left[\frac{1}{\|\theta_0 - \theta_*\|^{-2/\alpha} + C_\alpha t} \right]^\alpha, \quad (17)$$

where we defined

$$C_\alpha = \frac{1}{2\alpha} (\langle \theta_0 - \theta_*, \Sigma^{-\alpha}(\theta_0 - \theta_*) \rangle) + \frac{\gamma K_\alpha}{2 - K\gamma} \|\theta_0 - \theta_*\|^2)^{-1/\alpha} \text{ and } K_\alpha = \max_{i \leq n} \langle x_i, \Sigma^{-\alpha} x_i \rangle.$$

Hence, for all $t \geq 0$,

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \min \left\{ \|\theta_0 - \theta_*\|^2 e^{-\mu t}, \inf_{\alpha \in \mathbb{R}^+} \left[\frac{1}{\|\theta_0 - \theta_*\|^{-2/\alpha} + C_\alpha t} \right]^\alpha \right\},$$

and, if μ is non-zero but very small (e.g. 10^{-10}), this inequality describes well the difference between the transient regime of convergence that is polynomial and the asymptotic regime that is exponential but occurs after time-scale $1/\mu$. These polynomial convergence bounds, associated with the covariance upperbound, K_α , are common in the study of SGD in RKHS (Dieuleveut and Bach, 2016; Lin and Rosasco, 2016; Berthier et al., 2020). Second, by Markov inequality, remark that the estimates (i) and (ii) give convergence rates in probability. Furthermore, one steadily checks that one can derive the same kind of inequality as in (i) for the expected iterates $(\mathbb{E}[\theta_t])_{t \geq 0}$. In this case, one obtains the overestimation $\|\theta_0 - \theta_*\|^2 e^{-2\mu t}$ and from the perspective of the rate given by the Theorem, the stochastic nature of the dynamic seems to slow down the convergence. Finally, note that considering θ_0 as a random variable, all expectations can be thought of as conditional expectations $\mathbb{E}[\cdot | \theta_0]$. The results are illustrated in Figure 1 and 2.

4.2 The noisy case

In this section too, $(\theta_t)_{t \geq 0}$ follows the dynamics given by Eq.(13) initialized at $\theta_0 \in \mathbb{R}^d$. However, we assume in this subsection that $\ker \mathbf{X} = \{0\}$ and that $\mathcal{I} = \emptyset$. Note that this is generically the case if $n \geq d$. The important consequence of this is that the loss is uniformly lower bounded by a positive number, indeed, we have

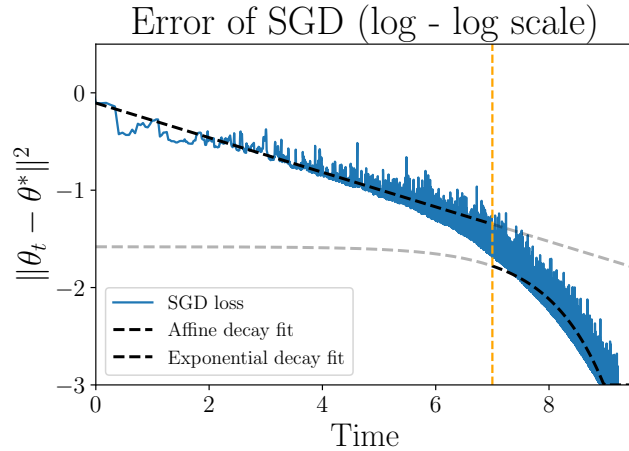


Figure 1: Plot showing the error of SGD along time for an overparametrized regime where $n = 100$ and $d = 200$. The samples $(x_i)_{i \leq n}$ come from a Gaussian distribution with a covariance whose eigenvalues decay as a power law ($\alpha = 0.2$) and $(y_i)_{i \leq n} \sim \mathcal{N}(0, I_n)$. The vertical dotted orange line illustrates the separation between the two regimes depicted by Theorem 4, the polynomial one before (a straight line in a log-log plot, shown with an affine fit with slope $\alpha \sim 0.2$) and the exponential line, after typical time scale $1/\mu$. This illustrates perfectly the rates of convergence shown in Theorem 4.




Figure 2: Display of a two-dimensional projection of 10 trajectories of SGD for $n = 100$, $d = 200$, in the case that there is an perfect interpolator θ_* . The ellipses represent the level curves of the training loss.

Lemma 5 *Let $\sigma^2 = L(\theta_*)$, we have, for all $\theta \in \mathbb{R}^d$ that $L(\theta) \geq \sigma^2$.*

Proof The function $\theta \mapsto L(\theta)$ is convex, hence any minimizer of L satisfies the equation $\nabla L(\theta) = 0$. This is the case of θ_* , as $\nabla L(\theta) = \Sigma(\theta - \theta_*)$. ■

Hence, the noise covariance matrix is (most of the time) uniformly bounded away from zero, and this implies that there exists a unique invariant measure to the SDE equation. This, in contrast to the noiseless case ($\mathcal{I} \neq \emptyset$) where any measure that is supported in \mathcal{I} is naturally invariant.

Semi-group. Due to the non-degeneracy of the noise term in this case, it is more convenient to track the dynamics of the probability measure $\rho_t := \text{Law}(\theta_t)$ for any $t \geq 0$, which is the time marginal of the process initialized at θ_0 and defined such that $\mathbb{E}[f(\theta_t)] = \int f d\rho_t$, for any $f : \mathbb{R}^d \rightarrow \mathbb{R}$ smooth enough. In fact the SDE Eq.(13) has an associated semi-group \mathcal{P}_t , defined so that for all $\eta \in \mathbb{R}^d$, $(\mathcal{P}_t f)(\eta) := \mathbb{E}[f(\theta_t) | \theta_{t=0} = \eta]$. Under a local strong ellipticity assumption (i.e., the diffusion matrix $\sigma(\cdot)\sigma(\cdot)^T$ is uniformly positive definite on compacts), classical regularization results for SDEs imply that, even if the initial condition is singular, such as $\rho_0 = \delta_{\theta_0}$, the law ρ_t of θ_t admits a density with respect to the Lebesgue measure for any $t > 0$, and has a finite second moment (Friedman, 2008, Section 4 of Chapter 9). We place ourselves in this setup where $\rho_0 \in \mathbb{P}_2(\mathbb{R}^d)$, the set of probability measures μ such that $\int \|\theta\|^2 d\mu(\theta) < +\infty$.

Infinitesimal generator. It is classical that the evolution of expectations of *observables* $(f(\theta_t))_{t \geq 0}$, for f smooth and compactly supported, is given by the Dynkin formula (Khasminskii, 2011, Lemma 3.2):

$$\frac{d}{dt} \mathbb{E}[f(\theta_t)] = \mathcal{L} \mathbb{E}[f(\theta_t)] , \tag{18}$$

where the operator \mathcal{L} , often referred to as the *infinitesimal generator* of the dynamics, acting on test functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ sufficiently smooth, is defined as

$$(\mathcal{L}f)(\theta) = -\langle \Sigma(\theta - \theta_*), \nabla f(\theta) \rangle + \frac{\gamma}{2n} \sum_{i,j=1}^d [\mathbf{X}^\top R_{\mathbf{x}}^2(\theta) \mathbf{X}]_{ij} \partial_{ij} f(\theta) . \tag{19}$$

Integrating by parts the Dynkin formula eq.(18) gives that the the time-marginal law of $(\theta_t)_{t \geq 0}$ satisfies the (parabolic) Fokker-Planck equation (at least in the weak sense):

$$\partial_t \rho_t = \mathcal{L}^* \rho_t ,$$

with the operator \mathcal{L}^* is the adjoint \mathcal{L} for the standard dot product of $L^2(\mathbb{R}^d)$, and writes

$$(\mathcal{L}^* f)(\theta) = \text{div} [\Sigma(\theta - \theta_*) f(\theta)] + \frac{\gamma}{2n} \sum_{i,j=1}^d \partial_{ij} \left(\left[\mathbf{X}^\top R_{\mathbf{x}}^2(\theta) \mathbf{X} \right]_{ij} f(\theta) \right) .$$

Dynkin formula enables, as in the deterministic case for gradient flow, the use of Lyapunov functions. This is a very useful tool to study the asymptotic behavior of stochastic processes. This is the objective of the following lemma:

Lemma 6 *Let $V(\theta) = \frac{1}{2}\|\theta - \theta_*\|^2$, we have the inequality, for all $\theta \in \mathbb{R}^d$,*

$$\mathcal{L}V(\theta) \leq -2(1 - \gamma\mathsf{K}/2)L(\theta) + 2\sigma^2 . \quad (20)$$

Consequently, for $\gamma < \frac{1}{3\mathsf{K}}$, there exists a stationary process for the SDE (13).

The proof is postponed to the Appendix A.2.

4.2.1 INVARIANT MEASURE AND CONVERGENCE.

Any invariant measure ρ_∞ satisfies the (elliptic) Partial Differential Equation (PDE): $\mathcal{L}^*\rho_\infty = 0$. Moreover, by ergodicity, the law of the iterates eventually converges towards this unique invariant measure. In the following result we also provide a quantitative statement, in Wasserstein distance, on the speed of convergence of the dynamics towards such a measure.

Theorem 7 *Let $(\theta_t)_{t \geq 0}$ follows the dynamics given by Eq.(13) initialized at θ_0 distributed according to $\rho_0 \in \mathbb{P}_2(\mathbb{R}^d)$, then for $\gamma < \frac{1}{3\mathsf{K}}$, there exists a unique stationary distribution $\rho^* \in \mathbb{P}_2(\mathbb{R}^d)$, and quantitatively,*

(i) **Parametric rate.** *For all $t \geq 0$, we have that*

$$W_2^2(\rho_t, \rho^*) \leq W_2^2(\rho_0, \rho^*)e^{-2\mu(1-2\gamma\mathsf{K})t}, \quad (21)$$

where $\mu > 0$ is the smallest non zero eigenvalue of Σ .

(ii) **Non-parametric rate.** *For all $t \geq 0$, we have that $\forall \alpha > 0$,*

$$W_2^2(\rho_t, \rho^*) \leq \left[\frac{1}{W_2^2(\rho_0, \rho^*)^{-2/\alpha} + C_\alpha t} \right]^\alpha, \quad (22)$$

with $C_\alpha = \frac{2(1-2\gamma\mathsf{K})}{\alpha} \left(\mathbb{E} [\langle \theta_0 - \Theta_, \Sigma^{-\alpha}(\theta_0 - \Theta_*) \rangle] + \frac{2\gamma\mathsf{K}_\alpha}{1-2\gamma\mathsf{K}} W_2^2(\rho_0, \rho^*) \right)^{-1/\alpha}$, and where the expectation in the first term is taken w.r.t. the optimal Wasserstein coupling between $\theta_0 \sim \rho_0$ and $\Theta_* \sim \rho_*$.*

The proof of these results is deferred in Appendix, Section A.2.1 and rests on coupling techniques standard in the asymptotic behavior of SDEs (von Renesse and Sturm, 2005) for the Wasserstein distance. One could adapt the argument to derive rates of convergence for other *probabilistic distances* that have a coupling representations like the total variation (Hairer, 2010); however, due to the multiplicative noise, convergence rates in the “natural metric” given by $L^2(\rho^*)$ are more difficult to obtain. This could be the occasion for a future investigation (see the recent work of Cattiaux and Guillin (2023) for a possibility to overcome this). Similarly as for the noiseless cases, there is a difference between time-scale of convergence between the early regime $t \ll 1/\mu$ that is polynomial and the asymptotic time $t \geq 1/\mu$ that is exponentially fast.

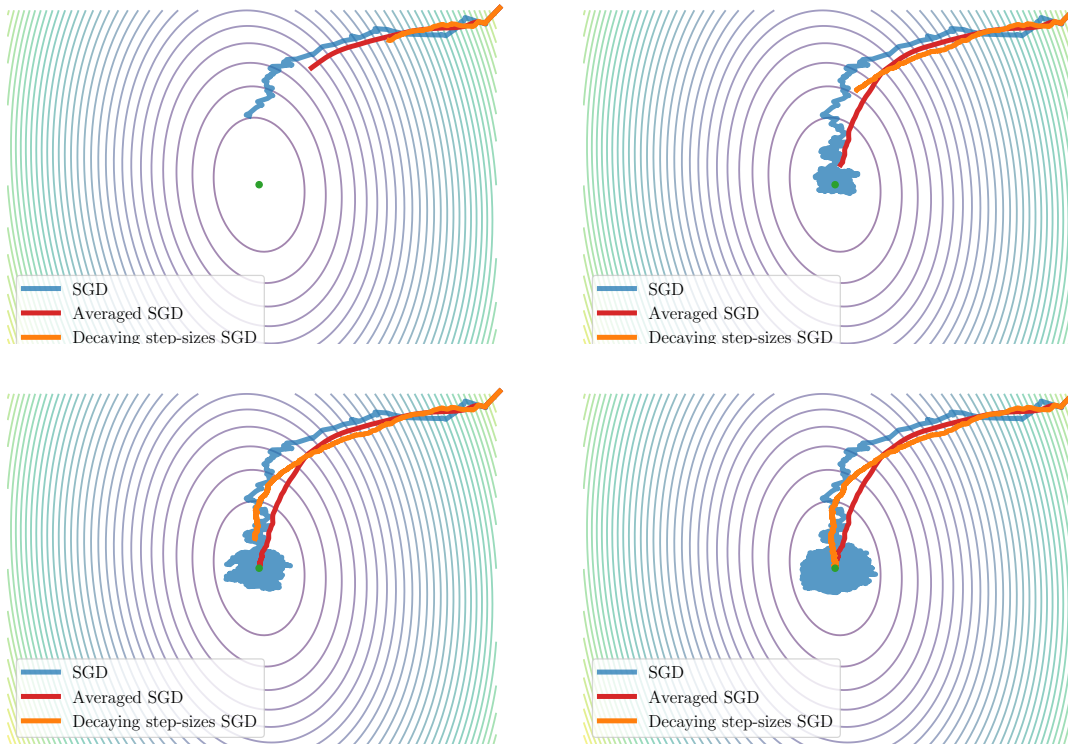


Figure 3: Four plots showing the trajectory of SGD in the noisy setting. The arrow of time goes from top left to bottom right. We see that the two variance reduction methods (time average and decaying step-sizes) converge towards θ_* (confirming Propositions 11 and 12), while plain SGD has a stationary distribution with certain fluctuations around its mean θ_* as explained in Theorem 7 and Proposition 8. Plain SGD is faster to its invariant distribution as the variance reduction methods as shown in the convergence rates provided in the results.

4.2.2 LOCALIZATION OF THE INVARIANT MEASURE.

Above, we have shown that the law of the iterates converges towards a measure ρ^* at a certain speed. We know that the invariant measure satisfies the equation $\mathcal{L}^*\rho^* = 0$; yet, this does not really give any practical description or insight on it. Notably, we want to understand how its localization depends on θ_* and on the parameters σ^2, γ or \mathbf{X} . This is the purpose of the following proposition.

Proposition 8 *Let $\Theta_* \sim \rho^*$, then for $\gamma < \frac{1}{3\mathbf{K}}$ we have*

$$\mathbb{E}[\Theta_*] = \theta_* \text{ , and } \mathbb{E}[\|\Theta_* - \theta_*\|^2] \leq \frac{\gamma\mathbf{K}\sigma^2}{\mu(1 - \gamma\mathbf{K})} . \quad (23)$$

The proof can be found in the Appendix, Section A.2.2. Note that, as expected, the deviation of Θ_* to its mean goes to 0 as σ or γ goes to 0. Moreover, in the limit $\gamma \ll 1$,

we have $\mathbb{E} \|\Theta_* - \theta_*\|^2 \sim \gamma \mathsf{K} \sigma^2 / \mu$, which reflects, up to constants, the scale given by the Gaussian calculation provided in Remark 9.

Remark 9 *Let $\sigma > 0$ a positive constant, we now consider the simplified model where $R_x = \sigma I_n$, so that the SDE dynamics becomes*

$$d\theta_t = -\frac{1}{n} \mathsf{X}^\top (\mathsf{X} \theta_t - \mathsf{y}) dt + \sqrt{\frac{\gamma \sigma^2}{n}} \mathsf{X}^\top dB_t.$$

This model is particularly convenient as it allows for an explicit computation of the invariant distribution. Solving $\mathcal{L}^ \rho^* = 0$ in this case gives for all $\theta \in \mathbb{R}^d$,*

$$\rho^*(\theta) \propto e^{-\frac{\|\theta - \theta_*\|^2}{\gamma \sigma^2}}.$$

That is to say that ρ^ is the Gaussian of mean θ_* and covariance $\frac{\gamma \sigma^2}{2} I_d$.*

However, recall that $R_x(\cdot)$ is in fact linear, this implies that for large $\|\theta\|$, the noise in the SDE amplifies: this has consequences on the tails of the distribution, see next paragraph.

Heavy-tails or not? Recall that at any time $t > 0$ all the moments of the law of θ_t exist (provided the initial distribution had such moments). Yet, depending on the step-size, moments of the invariant distribution might not exist. More precisely, we show that if one fixes a step-size $\gamma > 0$, all moments up to a certain value $\alpha(\gamma)$ of the invariant distribution exist and all higher moments do not. In other words, the step size is a direct control of the tail of the asymptotic distribution.

Proposition 10 *For $n \geq 2d$ and a fixed $\gamma < \frac{1}{3\mathsf{K}}$, there exists $\alpha > 0$ large enough (specifically, $\alpha > \frac{C}{\gamma}$ for C a constant depending only on X, y) such that*

$$\mathbb{E}(\|\Theta_*\|^\alpha) = +\infty.$$

The proof of this result can be found in the Appendix, Section A.2.2. This result is in accordance with the fact that multiplicative noise in SGD induces heavy-tails of its asymptotic distribution (Hodgkinson and Mahoney, 2021; Gurbuzbalaban et al., 2021; Jiao and Keller-Ressel, 2024).

4.2.3 CONVERGENCE OF VARIANCE REDUCED DYNAMICS

If the aim is to have an estimator that converges at long times to the optimal θ_* , it is necessary to use *variance reduction techniques*: here we show how time averages and step-size decay can be used in order to achieve this.

Ergodicity: convergence of the time average. Classical ergodic theorems tell that the time average along a trajectory converges towards spatial mean taken according to the invariant measure, i.e. $\frac{1}{t} \int_0^t \varphi(\theta_s) ds \rightarrow \mathbb{E}_{\rho^*} [\varphi(\Theta_*)]$, as t goes to infinity, for every smooth function φ . Taking $\varphi = I$, we thus expect that $\bar{\theta}_t := \frac{1}{t} \int_0^t \theta_s ds \rightarrow \mathbb{E}_{\rho^*} [\Theta_*] = \theta_*$, as t goes to infinity, as explained in Proposition 8. We quantify the convergence speed in the following.

Proposition 11 *We have for all $t \geq 0$ that*

$$\mathbb{E} \|\Sigma(\bar{\theta}_t - \theta_*)\|^2 \leq \frac{8\gamma \mathsf{K} \sigma^2}{t} + \frac{10\|\theta_0 - \theta_*\|^2}{t^2}.$$

Step-size decay. In this section, the step-size $\gamma = \gamma_t$ depends on the time t and tends to 0 as t tends to ∞ . In this case, $(\theta_t)_{t \geq 0}$ eventually converges.

Proposition 12 *Let $\gamma_t = 1/(2K + t^\alpha)$ and $\alpha > 1$, we have*

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \frac{C_\alpha}{t^{\alpha-1}},$$

with $C_\alpha = e^{-\alpha} (\mathbb{E}[\|\theta_0 - \theta_*\|^2])e(2(\alpha - 1)/\mu)^{\alpha-1} + (2\alpha/\mu)^\alpha \sigma^2 + \frac{2^{1+\alpha} K \sigma^2}{(\alpha-1)}$.

This result of polynomial decay when the step-size decay polynomially is similar to the one provided in the work of Fontaine et al. (2021), and is re-proven here for the sake of completeness. For both results, the proofs are postponed to the Appendix, Section A.2.3. All results on convergence are illustrated in Figure 3.

5. Online SGD (population loss)

In the previous section, on SGD in the *empirical setting*, we have shown that the core of the qualitative description relies on the fact that the noise can cancel or not. The situation is similar for online SGD. In fact most of the calculation that we have shown so far transfer almost immediately within this setting. This is the reason why we will only show the results concerning the convergence in the *noisy* and the *noiseless* settings and describe briefly what properties will be similar than the ones exhibited in the previous section.

First, we recall the main equation governing the dynamics (12):

$$d\theta_t = -\mathbb{E}_\rho [(\langle \theta_t, X \rangle - Y) X] dt + \sqrt{\gamma} \sigma(\theta_t) dB_t,$$

where $(B_t)_{t \geq 0}$ is a Brownian motion of \mathbb{R}^d and $\sigma \in \mathbb{R}^{d \times d}$ is given by

$$\sigma(\theta) := \left(\mathbb{E}_\rho \left[r_X(\theta)^2 X X^\top \right] - \mathbb{E}_\rho [r_X(\theta) X] \mathbb{E}_\rho [r_X(\theta) X]^\top \right)^{1/2} \in \mathbb{R}^{d \times d}.$$

Let us define $\Sigma = \mathbb{E}_\rho [X X^\top] \in \mathbb{R}^{d \times d}$ the covariance matrix that we assume invertible and $\theta_* = \Sigma^{-1} \mathbb{E}_\rho [Y X] \in \mathbb{R}^d$. We have that for all $\theta \in \mathbb{R}^d$, $\mathbb{E}_\rho [(\langle \theta, X \rangle - Y) X] = \Sigma(\theta - \theta_*)$, and hence the dynamics writes:

$$d\theta_t = -\Sigma(\theta_t - \theta_*) dt + \sqrt{\gamma} \sigma(\theta_t) dB_t.$$

Despite the difference between the empirical and population learning setups, the story here is similar to what is described in Section 4. Indeed, the core of the behavior of the algorithm relies only on the fact that θ_* cancels also the noise matrix σ , i.e $\sigma(\theta_*) = 0$. If this is the case, i.e. we are in the interpolation (or noiseless) regime, then the same analysis as in Subsection 4.1 applies and the movement resembles multivariate geometric Brownian motion. If uniformly $\sigma(\theta) \succ 0$, then there will be a non degenerate invariant measure and the same results as Subsection 4.2 go through.

5.1 The noiseless case

Here we assume that the distribution ρ is such that $Y = \langle X, \theta_* \rangle$ and that the features are almost surely bounded, i.e., $\|X\|^2 \leq K$ for some $K > 0$. This means that we are in the noiseless/interpolation regime where the input/output distribution admits a linear interpolator. As previously, in this case, despite the randomness of the SDE, the convergence is almost sure towards θ_* , with explicit rates that we show below.

Theorem 13 *Let $(\theta_t)_{t \geq 0}$ follows the dynamics given by Eq.(13) initialized at $\theta_0 \in \mathbb{R}^d$, then for $\gamma < \frac{1}{3K}$, we have that $\|\theta_t - \theta_*\|$ converges almost surely to 0 with the following rates*

(i) **Parametric rate.** *For all $t \geq 0$, we have that*

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \|\theta_0 - \theta_*\|^2 e^{-\mu(2-K\gamma)t}, \quad (24)$$

where $\mu > 0$ is the smallest non zero eigenvalue of Σ .

(ii) **Non-parametric rate.** *For all $t \geq 0$, we have that for all $\alpha > 0$,*

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \left[\frac{1}{\|\theta_0 - \theta_*\|^{-2/\alpha} + C_\alpha t} \right]^\alpha, \quad (25)$$

where we defined

$$C_\alpha = \frac{1}{2\alpha} (\langle \eta_0, \Sigma^{-\alpha} \eta_0 \rangle + \frac{\gamma K_\alpha}{2 - K\gamma} \|\theta_0 - \theta_*\|^2)^{-1/\alpha} \text{ and } K_\alpha = \max_{i \leq n} \langle x_i, \Sigma^{-\alpha} x_i \rangle.$$

Hence, as in the case of the training dynamics, we have that for all $t \geq 0$,

$$\mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \min \left\{ \|\theta_0 - \theta_*\|^2 e^{-\mu t}, \inf_{\alpha \in \mathbb{R}^+} \left[\frac{1}{\|\theta_0 - \theta_*\|^{-2/\alpha} + C_\alpha t} \right]^\alpha \right\},$$

and, if μ is non-zero but very small (e.g. 10^{-10}), this inequality describes well the difference between the transient regime of convergence that is polynomial and the asymptotic regime that is exponential but occurring after time-scale $1/\mu$. We recalled here the result as in the previous section for the sake of completeness, but the result and the proof of the theorem (which can be found in the Appendix, Section B.1) are rather similar to the training case. This shows that what really matters is the ability of the model to be interpolated or not and not the finite sample size.

5.2 The noisy case

In this section, we suppose that we are in the noisy regime. This corresponds to the fact that there exists a constant $a > 0$ such that for all $\theta \in \mathbb{R}^d$, we have $\sigma^2(\theta)/L(\theta) \gtrsim a^2 I_d$, as well as $L(\theta) \geq a^2$, that is to say that neither the loss nor the multiplicative noise can cancel. Let us give first a concrete example when this happens.

Example 4 (Gaussian model) *Assume that $X \sim \mathcal{N}(0, \Sigma)$, with $\Sigma \succeq \mu I_d$ and that there exists $\theta_* \in \mathbb{R}^d$ and $\xi \in \mathbb{R}$ a random variable independent of X of mean zero and variance $2\sigma^2$, such that $Y = \langle \theta_*, X \rangle + \xi$. Then, by some forth order Gaussian moment calculation given in (Berthier et al., 2020, Lemma 1), we have the exact calculation*

$$\sigma(\theta)^2 = (\Sigma(\theta - \theta_*)) (\Sigma(\theta - \theta_*))^\top + 2L(\theta)\Sigma \gtrsim 2\mu L(\theta)I_d,$$

as well as $L(\theta) \geq \sigma^2$.

Semi-group and Fokker-Planck equation. As for the training setup, due to the non-degeneracy of the noise term in this case, it is more convenient to track the dynamics of the probability measure $\rho_t := \text{Law}(\theta_t)$ for any $t \geq 0$. We place ourselves in the situation $\rho_0 \in \mathbb{P}_2(\mathbb{R}^d)$ and recall that $(\theta_t)_{t \geq 0}$ satisfies the (parabolic) Fokker-Planck equation (at least in the weak sense) $\partial_t \rho_t = \mathcal{L}^* \rho_t$, with the operator \mathcal{L}^* being the adjoint (with respect to the canonical dot product of $L^2(\mathbb{R}^d)$) of the infinitesimal generator of the dynamics that writes,

$$(\mathcal{L}f)(\theta) = -\langle \Sigma(\theta - \theta_*), \nabla f(\theta) \rangle + \frac{\gamma}{2} \sum_{i,j=1}^d [\sigma(\theta)\sigma(\theta)^\top]_{ij} \partial_{ij} f(\theta), \quad (26)$$

for all test functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ sufficiently smooth. Recall furthermore that the evolution of expectations of *observables* $(f(\theta_t))_{t \geq 0}$ is given by the action of the generator:

$$\frac{d}{dt} \mathbb{E}[f(\theta_t)] = \mathcal{L} \mathbb{E}[f(\theta_t)] . \quad (27)$$

5.2.1 INVARIANT MEASURE AND CONVERGENCE.

As the multiplicative noise does not cancel, \mathcal{L}^* is uniformly elliptic and there is existence and uniqueness of the invariant measure ρ_∞ of the SDE, which satisfies the PDE $\mathcal{L}^* \rho_\infty = 0$. Moreover, by ergodicity, the law of the iterates eventually converges towards this unique invariant measure. In order to show this quantitatively, we first prove a useful lemma on the Lipschitz behavior of the multiplicative noise:

Lemma 14 *There exists of constant $c > 0$ depending only on the distribution ρ , such that for all $\theta, \eta \in \mathbb{R}^d$, we have*

$$\|\sigma(\theta) - \sigma(\eta)\|_{\text{HS}}^2 \leq 2cK \langle \Sigma(\theta - \eta), \theta - \eta \rangle . \quad (28)$$

The proof is in the Appendix, Section B.2. We are now ready to state the main theorem of the section. Indeed, in the following result we also provide a quantitative statement, in Wasserstein distance, on the speed of convergence of the dynamics towards such a measure.

Theorem 15 *Let $(\theta_t)_{t \geq 0}$ follows the dynamics given by Eq.(13) initialized at $\theta_0 \in \mathbb{R}^d$, then for $\gamma < \frac{1}{Kc}$, there exists a unique stationary distribution $\rho^* \in \mathbb{P}_2(\mathbb{R}^d)$, and quantitatively,*

(i) **Parametric rate.** *For all $t \geq 0$, we have that*

$$W_2^2(\rho_t, \rho^*) \leq W_2^2(\rho_0, \rho^*) e^{-2\mu(1-\gamma Kc)t}, \quad (29)$$

where $\mu > 0$ is the smallest non zero eigenvalue of Σ .

(ii) **Non-parametric rate.** *Assume that we have the inequality for all $\alpha > 0$,*

$$\|\Sigma^{-\alpha/2}(\sigma(\theta) - \sigma(\eta))\|_{\text{HS}}^2 \leq 2c_\alpha K_\alpha \langle \Sigma(\theta - \eta), \theta - \eta \rangle , \quad (30)$$

then, for all $t \geq 0$, we have that $\forall \alpha > 0$,

$$W_2^2(\rho_t, \rho^*) \leq \left[\frac{1}{W_2^2(\rho_0, \rho^*)^{-2/\alpha} + C_\alpha t} \right]^\alpha , \quad (31)$$

with $C_\alpha = \frac{2(1-\gamma cK)}{\alpha} \left(\mathbb{E} [\langle \theta_0 - \Theta_*, \Sigma^{-\alpha}(\theta_0 - \Theta_*) \rangle] + \frac{\gamma c_\alpha K_\alpha}{1-\gamma cK} W_2^2(\rho_0, \rho^*) \right)^{-1/\alpha}$, and where the expectation in the first term is taken w.r.t. the optimal Wasserstein coupling between $\theta_0 \sim \rho_0$ and $\Theta_* \sim \rho_*$.

Similarly as before, the result is similar to the underparametrized regime for the training dynamics. This shows once again that what really matters is the ability of the model to be interpolated or not and not the finite sample size. The proof of the theorem can be found in Appendix, Section B.2.

5.2.2 FURTHER (EXPECTED) PROPERTIES.

In the following paragraph, we state what should be the properties of the dynamics by analogy with the SDE on the training loss case. Any proof could be adapted, but we take the option to avoid being redundant and state directly the expected properties without re-doing the work of the previous section. To be fair, we do not state the results as propositions and theorems. We will also abuse the notation $c > 0$ that plays the role of a universal constant in any of the expressions below.

It is possible to localize the invariant measure in a way similar to what has been done in Proposition 8. We give below estimates on its mean and standard deviation around it. In fact, if Θ_* is a random variable distributed according to ρ^* , then we expect that

$$\mathbb{E} [\Theta_*] = \theta_* , \text{ and } \mathbb{E} [\|\Theta_* - \theta_*\|^2] \leq c \frac{\gamma K \sigma^2}{\mu} . \quad (32)$$

We expect from ergodicity that time averages along a trajectory converge towards spatial mean taken according to the invariant measure, i.e. $\bar{\theta}_t := \frac{1}{t} \int_0^t \theta_s ds \rightarrow \mathbb{E}_{\rho^*} [\Theta_*] = \theta_*$, as t goes to infinity, as stated in the equation above. A quantification of this fact would give :

$$\mathbb{E} \|\Sigma(\bar{\theta}_t - \theta_*)\|^2 \leq c \frac{\gamma K \sigma^2 + \|\theta_0 - \theta_*\|^2}{t} .$$

Step-size decay would help canceling the noise in large times: indeed, choosing the step-size sequence as $\gamma_t = 1/(K + t^\alpha)$ for $\alpha > 1$, we get

$$\mathbb{E} [\|\theta_t - \theta_*\|^2] \leq \frac{c}{t^{\alpha-1}} ,$$

The heavy-tails phenomenon is also expected to hold similarly to the training case. More precisely, we show that if one fixes a step-size $\gamma > 0$, all moments up to a certain value $\alpha(\gamma)$ of the invariant distribution exist and all higher moments do not. In other words, the step size is a direct control of the tail of the asymptotic distribution.

6. Conclusion and Perspectives

In this article, we have shown how SGD could be efficiently modeled by a SDE to reflect its main qualitative and quantitative features: convergence speed, difference between noisy and noiseless settings and study of the asymptotic distribution. The specificity of the least-square set-up enabled us to show some localization of the invariant measure ρ_* in the noisy

context: however, it seems possible to improve this understanding to its precise *shape*, and some questions remain: is ρ_* log-concave? Is it possible to characterize its covariance in order to better apprehend its shape? Also, regarding its heavy-tail behavior, it would be great to have a precise estimate of the exponent, from which the moment of the stationary distribution explodes.

This work has been done to convey the idea that the SDE framework can improve the understanding of SGD dynamics. This is clear for the least-squares setup, yet the important question is to go beyond this and try to apply the same methodology for the non-convex dynamics arising from the training of non-linear neural networks. The study of *single or multi-index models* (Ben Arous et al., 2022; Bietti et al., 2023; Pillaud-Vivien et al., 2018) could be a first step toward broadening this systematic study.

Acknowledgments

AS acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – GZ 2047/1, Projekt-ID 390685813 and the support of the ERC MSCA grant SLOHD (101203974). AS and LP warmly thank the the Simons Foundation and especially the Flatiron Institute for its support as this research has been initiated while AS was invited to New York. Finally, the authors extend their gratitude to the *Incubateur de Fraîcheur* for hosting them and providing an ideal atmosphere that fostered exceptional discussions.

References

- Alnur Ali, Edgar Dobriban, and Ryan Tibshirani. The implicit regularization of stochastic gradient flow for least squares. In *International conference on machine learning*, pages 233–244. PMLR, 2020.
- Maksym Andriushchenko, Aditya Vardhan Varre, Loucas Pillaud-Vivien, and Nicolas Flammarion. Sgd with large step sizes learns sparse features. In *International Conference on Machine Learning*, pages 903–925. PMLR, 2023.
- Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Peter L Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. High-dimensional limit theorems for sgd: Effective dynamics and critical scaling. *Advances in Neural Information Processing Systems*, 35:25349–25362, 2022.
- Michel Benaïm. Dynamics of stochastic approximation algorithms. In *Seminaire de probabilités XXXIII*, pages 1–68. Springer, 2006.

- R. Berthier, F. Bach, and P. Gaillard. Tight nonparametric convergence rates for stochastic gradient descent under the noiseless linear model. In *Advances in Neural Information Processing Systems*, 2020.
- Alberto Bietti, Joan Bruna, and Loucas Pillaud-Vivien. On learning gaussian multi-index models with gradient flow. *arXiv preprint arXiv:2310.19793*, 2023.
- Blake Bordelon and Cengiz Pehlevan. Learning curves for sgd on structured features. *arXiv preprint arXiv:2106.02713*, 2021.
- L. Bottou and O. Bousquet. The tradeoffs of large scale learning. In *Advances in Neural Information Processing Systems*, 2008.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Patrick Cattiaux and Arnaud Guillin. A journey with the integrated γ 2 criterion and its weak forms. In *Geometric Aspects of Functional Analysis: Israel Seminar (GAFA) 2020-2022*, pages 167–208. Springer, 2023.
- Pratik Chaudhari and Stefano Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–10. IEEE, 2018.
- Hugo Cui, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. *Advances in Neural Information Processing Systems*, 34:10131–10143, 2021.
- A. Dieuleveut and F. Bach. Nonparametric stochastic approximation with large step-sizes. *Annals of Statistics*, 44(4):1363–1399, 2016.
- Xavier Fontaine, Valentin De Bortoli, and Alain Durmus. Convergence rates and approximation results for sgd and its continuous-time counterpart. In *Conference on Learning Theory*, pages 1965–2058. PMLR, 2021.
- Julie Lyng Forman and Michael Sørensen. The pearson diffusions: A class of statistically tractable diffusion processes. *Scandinavian Journal of Statistics*, 35(3):438–465, 2008.
- Avner Friedman. *Partial differential equations of parabolic type*. Courier Dover Publications, 2008.
- Mert Gurbuzbalaban, Umut Simsekli, and Lingjiong Zhu. The heavy-tail phenomenon in sgd. In *International Conference on Machine Learning*, pages 3964–3975. PMLR, 2021.
- Martin Hairer. Convergence of markov processes. *Lecture notes*, 18:26, 2010.

- Jeff Z HaoChen, Colin Wei, Jason Lee, and Tengyu Ma. Shape matters: Understanding the implicit bias of the noise covariance. In *Conference on Learning Theory*, pages 2315–2357. PMLR, 2021.
- J Harold, G Kushner, and George Yin. Stochastic approximation and recursive algorithm and applications. *Application of Mathematics*, 35, 1997.
- Liam Hodgkinson and Michael Mahoney. Multiplicative noise and heavy tails in stochastic optimization. In *International Conference on Machine Learning*, pages 4262–4274. PMLR, 2021.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Zhe Jiao and Martin Keller-Ressel. Emergence of heavy tails in homogenized stochastic gradient descent. *arXiv preprint arXiv:2402.01382*, 2024.
- Rafail Khasminskii. *Stochastic stability of differential equations*, volume 66. Springer Science & Business Media, 2011.
- Qianxiao Li, Cheng Tai, and E Weinan. Stochastic modified equations and dynamics of stochastic gradient algorithms i: Mathematical foundations. *The Journal of Machine Learning Research*, 20(1):1474–1520, 2019.
- Zhiyuan Li, Sadhika Malladi, and Sanjeev Arora. On the validity of modeling sgd with stochastic differential equations (sdes). *Advances in Neural Information Processing Systems*, 34:12712–12725, 2021.
- Junhong Lin and Lorenzo Rosasco. Optimal learning for multi-pass stochastic gradient methods. *Advances in Neural Information Processing Systems*, 29, 2016.
- Jonathan C Mattingly, Andrew M Stuart, and Michael V Tretyakov. Convergence of numerical time-averaging and stationary measures via poisson equations. *SIAM Journal on Numerical Analysis*, 48(2):552–577, 2010.
- Takashi Mori, Liu Ziyin, Kangqiao Liu, and Masahito Ueda. Power-law escape rate of sgd. In *International Conference on Machine Learning*, pages 15959–15975. PMLR, 2022.
- Gilles Pagès. Sur quelques algorithmes récurifs pour les probabilités numériques. *ESAIM: Probability and Statistics*, 5:141–170, 2001.
- Courtney Paquette, Elliot Paquette, Ben Adlam, and Jeffrey Pennington. Homogenization of sgd in high-dimensions: Exact dynamics and generalization properties. *arXiv preprint arXiv:2205.07069*, 2022.
- Scott Pesme, Loucas Pillaud-Vivien, and Nicolas Flammarion. Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity. *Advances in Neural Information Processing Systems*, 34:29218–29230, 2021.

- L. Pillaud-Vivien, A. Rudi, and F. Bach. Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes. *Advances in Neural Information Processing Systems*, 31:8114–8124, 2018.
- Loucas Pillaud-Vivien, Julien Reygner, and Nicolas Flammarion. Label noise (stochastic) gradient descent implicitly solves the lasso for quadratic parametrisation. In *Conference on Learning Theory*, pages 2127–2159. PMLR, 2022.
- H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Statistics*, 22:400–407, 1951.
- Aditya Vardhan Varre, Loucas Pillaud-Vivien, and Nicolas Flammarion. Last iterate convergence of sgd for least-squares in the interpolation regime. *Advances in Neural Information Processing Systems*, 34:21581–21591, 2021.
- S. Vaswani, F. Bach, and M. Schmidt. Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron. In *International Conference on Artificial Intelligence and Statistics*, pages 1195–1204. PMLR, 2019.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Max-K von Renesse and Karl-Theodor Sturm. Transport inequalities, gradient estimates, entropy and ricci curvature. *Communications on pure and applied mathematics*, 58(7):923–940, 2005.
- Stephan Wojtowytsch. Stochastic gradient descent with noise of machine learning type part i: Discrete time analysis. *Journal of Nonlinear Science*, 33(3):45, 2023.
- Stephan Wojtowytsch. Stochastic gradient descent with noise of machine learning type part ii: Continuous time analysis. *Journal of Nonlinear Science*, 34(1):16, 2024.
- Lei Wu, Chao Ma, et al. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. *Advances in Neural Information Processing Systems*, 31, 2018.
- Zeke Xie, Issei Sato, and Masashi Sugiyama. A diffusion theory for deep learning dynamics: Stochastic gradient descent exponentially favors flat minima. *arXiv preprint arXiv:2002.03495*, 2020.
- Chen Xing, Devansh Arpit, Christos Tsirigotis, and Yoshua Bengio. A walk with sgd. *arXiv preprint arXiv:1802.08770*, 2018.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- Liu Ziyin, Kangqiao Liu, Takashi Mori, and Masahito Ueda. Strength of minibatch noise in sgd. *arXiv preprint arXiv:2102.05375*, 2021.

ORGANIZATION OF THE APPENDIX

A. SGD on the training loss: Proof of Section 4.

A.1: The noiseless case. This includes the proof of the fact that θ_* is the orthogonal projection of θ_0 into \mathcal{I} (Lemma 3) and the proof of the convergence theorem of θ to θ_* (Theorem 4).

A.2: The noisy case. We prove here the existence of a stationary distribution (Lemma 6).

A.2.1: Invariant measure and convergence. In this subsection, we prove the quantitative convergence to the stationary distribution (Theorem 7) with the help of a technical Lemma (Lemma 16).

A.2.2: Localization of the invariant measure. We prove the insights given on the first and second moments of the stationary distribution (Proposition 8). Then, we show the moment explosion of the invariant distribution (Proposition 10).

A.2.3: Convergence of variance reduction techniques. In this subsection, we prove the convergence of the time-average of the iterates to θ_* , i.e. ergodicity (Proposition 11) as well as the convergence of θ_t to θ_* in the case of the step-size decay (Proposition 12).

B: Online SGD: Proofs of Section 5.

B.1: The noiseless case We give the proof of the convergence to θ_* with rates (Theorem 13).

B.2: The noisy case. We first prove that multiplicative noise carries some Lipschitz property (Lemma 14). This allow us to prove the quantitative convergence to the the stationary distribution (Theorem 15).

Appendix A. SGD on the training loss: Proofs of Section 4

A.1 The noiseless case

We begin by proving Lemma 3 on the fact that θ_* is the orthogonal projection of θ_0 into \mathcal{I} , that is $\theta_* = \operatorname{argmin}_{X\theta=y} \|\theta - \theta_0\|^2$.

Proof [Lemma 3] The proof of the lemma follows from Karush–Kuhn–Tucker conditions. Indeed, the argmin is unique, being the projection of θ_0 on a linear set, and it satisfies that there exists Lagrange multipliers $\lambda \in \mathbb{R}^n$ such that

$$\theta_* - \theta_0 = X^\top \lambda, \text{ and } X\theta_* = y.$$

This means that $\lambda = (\mathbf{X}\mathbf{X}^\top)^{-1}(y - \mathbf{X}\theta_0)$, and hence,

$$\begin{aligned}\theta_* &= \theta_0 + \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1}(y - \mathbf{X}\theta_0) \\ &= \mathbf{X}^\dagger y + (I - \mathbf{X}^\dagger \mathbf{X})\theta_0,\end{aligned}$$

which corresponds to the given definition of θ_* . ■

We now prove the main convergence theorem of this section. This is Theorem 4.

Proof [Theorem 4] (i) The key ingredient of the proof is the Gronwall Lemma. Combining the Itô formula with (13) gives us

$$\frac{d}{dt}\mathbb{E}\|\theta_t - \theta_*\|^2 = -2\mathbb{E}\langle \theta_t - \theta_*, \Sigma(\theta_t - \theta_*) \rangle + \frac{\gamma}{n}\mathbb{E}\text{Tr}\left(\mathbf{X}\mathbf{X}^\top R_x^2(\theta_t)\right).$$

By integrating the latter, we get

$$\begin{aligned}\mathbb{E}\|\theta_t - \theta_*\|^2 &= \|\theta_0 - \theta_*\|^2 - \frac{2}{n}\int_0^t \mathbb{E}\|\mathbf{X}(\theta_u - \theta_*)\|^2 du \\ &\quad + \frac{\gamma}{n}\int_0^t \mathbb{E}\text{Tr}\left(\mathbf{X}^\top \left(\text{diag}(r_x(\theta))^2 - \frac{1}{n}r_x(\theta)r_x(\theta)^\top\right)\mathbf{X}\right) du.\end{aligned}$$

Note that

$$\begin{aligned}\frac{\gamma}{n}\int_0^t \mathbb{E}\text{Tr}\left(\mathbf{X}^\top (\text{diag}(r_x(\theta))^2)\mathbf{X}\right) du &= \frac{\gamma}{n}\int_0^t \mathbb{E}\sum_{i=1}^n r_x^2(\theta)_i (\mathbf{X}\mathbf{X}^\top)_{ii} du \\ &= \frac{\gamma}{n}\int_0^t \mathbb{E}\sum_{i=1}^n r_x^2(\theta)_i \|x_i\|^2 du \\ &\leq \frac{\gamma}{n}\int_0^t \mathbb{E}\max_i \|x_i\|^2 \sum_{i=1}^n r_x^2(\theta)_i du \\ &\leq 2\gamma\mathbf{K}\int_0^t \mathbb{E}L(\theta_u) du,\end{aligned}$$

and

$$\begin{aligned}\frac{\gamma}{n^2}\int_0^t \mathbb{E}\text{Tr}\left(\mathbf{X}^\top r_x(\theta)r_x(\theta)^\top\mathbf{X}\right) du &= \frac{\gamma}{n^2}\int_0^t \mathbb{E}\text{Tr}\left(r_x(\theta_u)^\top\mathbf{X}\mathbf{X}^\top r_x(\theta_u)\right) du \\ &\geq \frac{\gamma}{n}\lambda_{\min}\left(\frac{1}{n}\mathbf{X}\mathbf{X}^\top\right)\int_0^t \mathbb{E}\left(r_x(\theta_u)^\top r_x(\theta_u)\right) du = 2\gamma\lambda_{\min}\left(\frac{1}{n}\mathbf{X}\mathbf{X}^\top\right)\int_0^t \mathbb{E}L(\theta_u) du.\end{aligned}$$

Collecting the estimates, we obtain

$$\mathbb{E}\|\theta_t - \theta_*\|^2 \leq \|\theta_0 - \theta_*\|^2 - \left(4 + 2\gamma\lambda_{\min}\left(\frac{1}{n}\mathbf{X}\mathbf{X}^\top\right) - 2\gamma\mathbf{K}\right)\int_0^t \mathbb{E}L(\theta_u) du.$$

We remark that $\theta_u - \theta_* \in \text{Ran}(\mathbf{X}^\top)$. Indeed, we have

$$\theta_u - \theta_* = (\theta_u - \theta_0) + (\theta_0 - \theta_*),$$

the first term on the r.h.s. is in $\text{Ran}(X^T)$ by (13) and the second term also by (14). We recall that $\mathbb{R}^d = \text{Ran}(X^T) \oplus \text{Ker}(X)$ and therefore note that $X|_{\text{Ran}(X^T)}$ is a bijection into its image. Combining the two last facts yield

$$\frac{\lambda_{\min}(\frac{1}{n}\mathbf{X}\mathbf{X}^\top)}{2}\mathbb{E}\|\theta_u - \theta_*\|^2 \leq \frac{1}{2n}\mathbb{E}\|\mathbf{X}(\theta_u - \theta_*)\|^2 = \mathbb{E}L(\theta_u)$$

because $\mathbf{X}\mathbf{X}^\top$ has the same spectrum as $\mathbf{X}^\top\mathbf{X}$. All in all, we have

$$\mathbb{E}\|\theta_t - \theta_*\|^2 \leq \|\theta_0 - \theta_*\|^2 - \lambda_{\min}(\frac{1}{n}\mathbf{X}\mathbf{X}^\top) \left(2 + \gamma\lambda_{\min}(\frac{1}{n}\mathbf{X}\mathbf{X}^\top) - \gamma\mathbf{K}\right) \int_0^t \mathbb{E}\|\theta_u - \theta_*\|^2 du,$$

and the first statement of the Theorem follows with Gronwall Lemma. We move to the proof of the second statement of the Theorem.

(ii) We define $\eta_t := \theta_t - \theta_*$ and recall that

$$\mathbb{E}\|\eta_t\|^2 \leq \|\theta_0 - \theta_*\|^2 - 2(2 - \gamma\mathbf{K}) \int_0^t \mathbb{E}L(\theta_u) du.$$

The first consequence of this inequality is that

$$\int_0^t \mathbb{E}L(\theta_u) du \leq \frac{1}{2(2 - \gamma\mathbf{K})} \|\theta_0 - \theta_*\|^2.$$

We want to lower bound the term $\mathbb{E}L(\theta_u)$ without using the smallest non negative eigenvalue of Σ , that is supposedly infinitely small. First note $\mathbf{X}^\top\mathbf{X}$ is symmetric and thus diagonalizable in an orthonormal basis (v_i) . We steadily check that

$$\text{Ker}(\mathbf{X}^\top\mathbf{X}) = \text{Ker}(\mathbf{X})$$

by invertibility of $\mathbf{X}\mathbf{X}^\top$. We thus have $\mathbb{R}^d = \text{Ran}(X^T) \oplus \text{Ker}(\mathbf{X}^\top\mathbf{X})$. We denote by $\lambda_1 \leq \dots \leq \lambda_{d^*}$ the non zero eigenvalue of $\mathbf{X}^\top\mathbf{X}$ where $d^* \leq d$ is the number of non-zero eigenvalues and (v_1, \dots, v_{d^*}) the corresponding eigenvectors. For all $t > 0$, we thus have the following decomposition $\eta_t = \sum_{k=1}^{d^*} \eta_k^t v_k$. Thanks to the Hölder inequality, we claim that the following holds true:

Let $p, q \in (0, 1)$, $p + q = 1$, we have, for all $t \geq 0$,

$$\left(\mathbb{E}[\|\eta_t\|^2]\right)^{\frac{1}{p}} \leq 2\mathbb{E}[L(\theta_t)] \left(\mathbb{E}[\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle]\right)^{q/p}. \quad (33)$$

We move to the proof of the above claim. To lighten the notation, we write $\eta_k^t = \eta_k$ in the following to wit

$$\mathbb{E}[\|\eta_t\|^2] = \sum_{k=1}^{d^*} \mathbb{E}[\eta_k^2] = \sum_{k=1}^{d^*} \mathbb{E}[\eta_k^{2p} \lambda_k^p \eta_k^{2q} \lambda_k^{-p}] \leq \sum_{k=1}^{d^*} \left(\mathbb{E}[\eta_k^2 \lambda_k]\right)^p \left(\mathbb{E}[\eta_k^2 \lambda_k^{-p/q}]\right)^q,$$

thanks to Hölder inequality w.r.t. the expectation. Then, applying once again the Hölder inequality for the sum, we have

$$\begin{aligned} \mathbb{E}[\|\eta_t\|^2] &\leq \left(\sum_{k=1}^{d^*} \mathbb{E}[\eta_k^2 \lambda_k]\right)^p \left(\sum_{k=1}^{d^*} \mathbb{E}[\eta_k^2 \lambda_k^{-p/q}]\right)^q \\ &= (\mathbb{E}2L(\theta))^p \left(\mathbb{E}[\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle]\right)^q, \end{aligned}$$

and the claim (33) follows. We now prove that $t \mapsto \mathbb{E}(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle)$ is bounded. Indeed, we proceed as before with Itô formula to get

$$\begin{aligned}
 \frac{d}{dt} \mathbb{E}(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle) &= 2\mathbb{E}\langle \frac{d\eta_t}{dt}, \Sigma^{-p/q} \eta_t \rangle + \mathbb{E}\langle \frac{d\eta_t}{dt}, \Sigma^{-p/q} \frac{d\eta_t}{dt} \rangle \\
 &= -2\mathbb{E}\langle -\Sigma \eta_t, \Sigma^{-p/q} \eta_t \rangle + \mathbb{E}\langle \sqrt{\frac{\gamma}{n}} \mathbf{X}^\top R_x(\theta_t) dB_t, \Sigma^{-p/q} \sqrt{\frac{\gamma}{n}} \mathbf{X}^\top R_x(\theta_t) dB_t \rangle \\
 &= -2\mathbb{E}\langle \eta_t, \Sigma^{1-p/q} \eta_t \rangle + \frac{\gamma}{n} \mathbb{E} \text{Tr} \left((\mathbf{X}^\top R)^T \Sigma^{-p/q} \mathbf{X}^\top R \right) \\
 &\leq \frac{\gamma}{n} \mathbb{E} \text{Tr} \left(\Sigma^{-p/q} \mathbf{X}^\top \left(\text{diag}(r_x(\theta))^2 - \frac{1}{n} r_x(\theta) r_x(\theta)^\top \right) \mathbf{X} \right) \\
 &\leq \frac{\gamma}{n} \mathbb{E} \sum_{i=1}^n \langle x_i, \Sigma^{-p/q} x_i \rangle (\langle \eta_t, x_i \rangle - y_i)^2 \\
 &\leq 2\gamma \mathbf{K}_{p/q} L(\theta_t) ,
 \end{aligned}$$

where we have used that for all $i \in \llbracket 1, n \rrbracket$, we have $\langle x_i, \Sigma^{-p/q} x_i \rangle \leq \mathbf{K}_{p/q}$. Then, by integrating with respect to t , it yields

$$\begin{aligned}
 \mathbb{E}(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle) &\leq \langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + 2\gamma \mathbf{K}_{p/q} \int_0^t \mathbb{E} L(\theta_u) du \\
 &\leq \langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + \frac{\gamma \mathbf{K}_{p/q}}{2 - \gamma \mathbf{K}} \|\theta_0 - \theta_*\|^2 .
 \end{aligned}$$

Hence, calling $\mathbf{C} = \frac{1}{2}(\langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + \frac{\gamma \mathbf{K}_{p/q}}{2 - \gamma \mathbf{K}} \|\theta_0 - \theta_*\|^2)^{-p/q}$, we have the inequality, for all $t \geq 0$,

$$\mathbb{E} L(\theta_t) \geq \mathbf{C} (\mathbb{E} \|\eta_t\|^2)^{1/p} ,$$

and this yields the inequality

$$\mathbb{E} \|\eta_t\|^2 \leq \|\eta_0\|^2 - \mathbf{C} \int_0^t (\mathbb{E} \|\eta_u\|^2)^{1/p} du , \quad (34)$$

that implies, from a slight modification of Gronwall Lemma that for all $t \geq 0$, we have

$$\mathbb{E} \|\eta_t\|^2 \leq \left[\frac{1}{\frac{1}{\|\eta_0\|^{2(1/p-1)}} + (1/p-1) \mathbf{C} t} \right]^{\frac{1}{1/p-1}} ,$$

this gives the result claim in the theorem. To see how the last inequality goes, we define $g(t) = \|\eta_0\|^2 - \mathbf{C} \int_0^t (\mathbb{E} \|\eta_u\|^2)^{1/p} du$ (which is positive) and we rewrite (34) as

$$\left(\frac{g'(t)}{-\mathbf{C}} \right)^p \leq g(t) \iff \frac{g'(t)}{g(t)^{1/p}} \geq -\mathbf{C} \implies \frac{1}{-1/p+1} (g(t)^{-1/p+1} - g(0)^{-1/p+1}) \geq -\mathbf{C} t ,$$

and thus

$$g(t) \leq \left[-\mathbf{C} t \left(-\frac{1}{p} + 1 \right) + \|\eta_0\|^{2(-1/p+1)} \right]^{\frac{1}{-1/p+1}} ,$$

and we conclude with (34).

To prove the convergence almost surely, we use the Itô Formula to obtain

$$\begin{aligned} \mathbb{E} (\|\theta_t - \theta_*\|^2 | \mathcal{F}_s) &= \|\theta_0 - \theta_*\|^2 + 2 \int_0^s \sqrt{\frac{\gamma}{n}} \langle R_x(\theta_u)^\top \mathbf{X}(\theta_u - \theta_*), dB_u \rangle \\ &\quad + \mathbb{E} \left(\int_0^t \frac{\gamma}{n} \text{Tr} \left(\mathbf{X}^\top \left(\text{diag}(r_x(\theta))^2 du - \frac{1}{n} r_x(\theta_u) r_x(\theta_u)^\top \right) \mathbf{X} \right) \right. \\ &\quad \left. - \frac{2}{n} \langle (\theta_u - \theta_*), \mathbf{X}^\top \mathbf{X}(\theta_t - \theta_*) \rangle du \middle| \mathcal{F}_s \right), \end{aligned}$$

For $\gamma < \frac{1}{2\mathbf{K}}$, we have proven in (i) that the term inside the conditional expectation value/integral is negative. We can thus overestimate the latter by integrating from 0 to s to obtain

$$\mathbb{E} (\|\theta_t - \theta_*\|^2 | \mathcal{F}_s) \leq \|\theta_s - \theta_*\|^2,$$

and we deduce that $\|\theta_t - \theta_*\|^2$ is a positive supermartingale, L^1 bounded and therefore converges almost surely to 0. \blacksquare

A.2 The noisy case

We begin by showing why $V(\theta) := \frac{1}{2}\|\theta - \theta_*\|^2$ is a Lyapunov function for the dynamics. This is Lemma 6.

Proof [Lemma 6] First, as L is a quadratic function, we have that for all $\theta \in \mathbb{R}^d$, we have

$$\langle \nabla L(\theta), \theta - \theta_* \rangle = 2(L(\theta) - L(\theta_*)).$$

We recall that $\mathbf{X}^\top \mathbf{X} \theta_* = \mathbf{X}^\top y$ and it follows that $\nabla L(\theta) = \Sigma(\theta - \theta_*)$. We thus deduce that for all $\theta \in \mathbb{R}^d$,

$$\begin{aligned} \mathcal{L}V(\theta) &= -\langle \Sigma(\theta - \theta_*), (\theta - \theta_*) \rangle + \frac{\gamma}{2n} \text{Tr} \left[\mathbf{X}^\top R_x^2(\theta) \mathbf{X} \right] \\ &= 2(L(\theta_*) - L(\theta)) + \frac{\gamma}{2n} \text{Tr} \left[\mathbf{X}^\top R_x^2(\theta) \mathbf{X} \right] \\ &\leq 2(L(\theta_*) - L(\theta)) + \frac{\gamma}{2n} \sum_{i=1}^n (\langle x_i, \theta \rangle - y_i)^2 \|x_i\|^2 \\ &\leq 2 \left(L(\theta_*) - L(\theta) \left(1 - \frac{\gamma \mathbf{K}}{2} \right) \right). \end{aligned}$$

This Lyapunov identity implies the existence of a stationary distribution as explained in (Khasminskii, 2011, Theorem 3.7). This implies that for all $t \geq 0$, the dynamics does not explode. \blacksquare

A.2.1 INVARIANT MEASURE AND CONVERGENCE

We now turn into proving the main theorem of this part on the quantitative convergence to the stationary distribution.

Proof [Theorem 7] (i) Wassertein contraction comes essentially from *coupling arguments*. Let $\gamma \leq \mathsf{K}^{-1}$ and $\rho_0^1, \rho_0^2 \in \mathbb{P}_2(\mathbb{R}^d)$ two possible initial distributions. Then by (Villani et al., 2009, Theorem 4.1), there exists a couple of random variables θ_0^1, θ_0^2 such that $W_2^2(\rho_0^1, \rho_0^2) = \mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2]$. Let $(\theta_t^1)_{t \geq 0}$ (resp. $(\theta_t^2)_{t \geq 0}$) be the solution of the SDE (13), sharing the same Brownian motion $(B_t)_{t \geq 0}$. Then, for all $t \geq 0$, the random variable (θ_t^1, θ_t^2) is a coupling between ρ_t^1 and ρ_t^2 , and hence

$$W_2^2(\rho_t^2, \rho_t^1) \leq \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2].$$

Moreover, we denote by $\|\cdot\|_{\text{HS}}$ the Frobenius norm and by Itô formula, we have

$$\begin{aligned} \frac{d}{dt} \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] &= -\frac{2}{n} \mathbb{E} \langle \theta_t^1 - \theta_t^2, \mathbf{X}^\top (\mathbf{X} \theta_t^1 - y) - \mathbf{X}^\top (\mathbf{X} \theta_t^2 - y) \rangle + \frac{\gamma}{n} \mathbb{E} \|\mathbf{X}^\top R_x(\theta_t^1) - \mathbf{X}^\top R_x(\theta_t^2)\|_{\text{HS}}^2 \\ &= -\frac{2}{n} \mathbb{E} \|\mathbf{X}(\theta_t^1 - \theta_t^2)\|^2 + \frac{\gamma}{n} \mathbb{E} \|\mathbf{X}^\top (R_x(\theta_t^1) - R_x(\theta_t^2))\|_{\text{HS}}^2 \\ &\leq -\frac{2}{n} \mathbb{E} \|\mathbf{X}(\theta_t^1 - \theta_t^2)\|^2 \\ &\quad + \frac{2\gamma}{n} \mathbb{E} \left(\|\mathbf{X}^\top \text{diag}(\langle \theta_t^1 - \theta_t^2, x_i \rangle_i)\|_{\text{HS}}^2 + \frac{1}{n^2} \|\mathbf{X}^\top (\langle \theta_t^1 - \theta_t^2, x_i \rangle_i) \mathbf{1}^\top\|_{\text{HS}}^2 \right). \end{aligned}$$

Furthermore, we have for all $\theta^1, \theta^2 \in \mathbb{R}^d$ that

$$\begin{aligned} \|\mathbf{X}^\top \text{diag}(\langle \theta^1 - \theta^2, x_i \rangle_i)\|_{\text{HS}}^2 &= \text{Tr} \left(\mathbf{X} \mathbf{X}^\top \text{diag}(\langle \theta^1 - \theta^2, x_i \rangle_i^2) \right) \\ &= \sum_{i=1}^n \|x_i\|^2 \langle \theta^1 - \theta^2, x_i \rangle^2 \\ &\leq \mathsf{K} \|\mathbf{X}(\theta^1 - \theta^2)\|^2, \end{aligned}$$

and

$$\begin{aligned} \frac{1}{n^2} \|\mathbf{X}^\top (\langle \theta^1 - \theta^2, x_i \rangle_i) \mathbf{1}^\top\|_{\text{HS}}^2 &= \frac{1}{n} \text{Tr} \left(\mathbf{X} \mathbf{X}^\top (\langle \theta^1 - \theta^2, x_i \rangle_i) (\langle \theta^1 - \theta^2, x_i \rangle_i)^\top \right) \\ &\leq \frac{1}{n} \text{Tr} \left(\mathbf{X} \mathbf{X}^\top \right) \text{Tr} \left((\langle \theta^1 - \theta^2, x_i \rangle_i) (\langle \theta^1 - \theta^2, x_i \rangle_i)^\top \right) \\ &\leq \mathsf{K} \|\mathbf{X}(\theta^1 - \theta^2)\|^2, \end{aligned}$$

where we use from the first to the second line the inequality $\text{Tr}(AB) \leq \text{Tr}(A)\text{Tr}(B)$ for any A and B positive semi-definite. Altogether, this gives the inequality:

$$\frac{d}{dt} \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq -\frac{2-4\gamma\mathsf{K}}{n} \mathbb{E} \|\mathbf{X}(\theta_t^1 - \theta_t^2)\|^2 \quad (35)$$

$$\leq -2\mu(1-2\gamma\mathsf{K}) \mathbb{E} \|\theta_t^1 - \theta_t^2\|^2, \quad (36)$$

and by Gronwall Lemma, denoting $c_\gamma = 2\mu(1 - 2\gamma\mathbf{K})$ this gives that

$$W_2^2(\rho_t^1, \rho_t^2) \leq \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq e^{-c_\gamma t} \mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2] = e^{-c_\gamma t} W_2^2(\rho_0^2, \rho_0^1) .$$

Now, for all $s \geq 0$ setting $\rho_0^1 = \rho_0 \in \mathcal{P}_2(\mathbb{R}^d)$ and $\rho_0^2 = \rho_s \in \mathbb{P}_2(\mathbb{R}^d)$, we have for all $t \geq 0$,

$$W_2^2(\rho_t, \rho_{t+s}) \leq e^{-c_\gamma t} W_2^2(\rho_0, \rho_s) ,$$

which shows that the process $(\rho_t)_{t \geq 0}$ is of Cauchy type, and since $(\mathbb{P}_2(\mathbb{R}^d), W_2)$ is a Polish space, $\rho_t \rightarrow \rho_* \in \mathbb{P}_2(\mathbb{R}^d)$ as t grows to infinity. Now, since there exists a stationary solution to the process, let us fix $\rho_0^1 = \rho^* \in \mathbb{P}_2(\mathbb{R}^d)$ and $\rho_0^2 = \rho_0 \in \mathbb{P}_2(\mathbb{R}^d)$. We have then,

$$W_2^2(\rho_t, \rho^*) \leq e^{-c_\gamma t} W_2^2(\rho_0, \rho^*) ,$$

which concludes the first part of the Theorem.

(ii) We will use the same steps as for the proof of Theorem 4 (ii). Again, one steadily checks that for $p, q \in (0, 1)$, $p + q = 1$, we have, for all $t \geq 0$,

$$\frac{1}{n} \mathbb{E} [\|\mathbf{X}(\theta_t^1 - \theta_t^2)\|^2] \geq \frac{\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2]^{1/p}}{\mathbb{E} [\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q}(\theta_t^1 - \theta_t^2) \rangle]^{q/p}} . \quad (37)$$

By Ito formula, skipping the details, we get

$$\begin{aligned} & \frac{d}{dt} \mathbb{E} \left(\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q}(\theta_t^1 - \theta_t^2) \rangle \right) \\ &= -2\mathbb{E} \langle \theta_t^1 - \theta_t^2, \Sigma^{1-p/q}(\theta_t^1 - \theta_t^2) \rangle + \frac{\gamma}{n} \mathbb{E} \text{Tr} \left(\Sigma^{-p/q} \mathbf{X}^T (R_x(\theta_1) - R_x(\theta_2)) (R_x(\theta_1) - R_x(\theta_2))^\top \mathbf{X} \right) \\ &\leq \frac{\gamma}{n} \mathbb{E} \text{Tr} \left(\Sigma^{-p/q} \mathbf{X}^T (R_x(\theta_1) - R_x(\theta_2)) (R_x(\theta_1) - R_x(\theta_2))^\top \mathbf{X} \right) \\ &= \frac{\gamma}{n} \|(R_x(\theta_1) - R_x(\theta_2))^\top \mathbf{X} \Sigma^{-p/2q}\|_{\text{HS}}^2 \\ &\leq \frac{2\gamma}{n} \mathbb{E} (\|\Sigma^{-p/2q} \mathbf{X}^\top \text{diag}(\langle \theta_t^1 - \theta_t^2, x_i \rangle_i)\|_{\text{HS}}^2 + \frac{1}{n^2} \|\Sigma^{-p/2q} \mathbf{X}^\top (\langle \theta_t^1 - \theta_t^2, x_i \rangle_i) \mathbb{1}^\top\|_{\text{HS}}^2) \\ &= \frac{2\gamma}{n} \mathbb{E} [\text{Tr} \left(\mathbf{X} \Sigma^{-p/q} \mathbf{X}^\top \text{diag}(\langle \theta^1 - \theta^2, x_i \rangle_i^2) \right) + \frac{1}{n} \text{Tr} \left(\mathbf{X} \Sigma^{-p/q} \mathbf{X}^\top (\langle \theta_t^1 - \theta_t^2, x_i \rangle_i) (\langle \theta_t^1 - \theta_t^2, x_i \rangle_i)^\top \right)] \\ &\leq \frac{2\gamma}{n} \mathbb{E} \left[\sum_{i=1}^n \langle x_i, \Sigma^{-p/q} x_i \rangle \langle \theta^1 - \theta^2, x_i \rangle^2 + \frac{1}{n} \text{Tr} \left(\mathbf{X} \Sigma^{-p/q} \mathbf{X}^\top \right) \text{Tr} \left((\langle \theta_t^1 - \theta_t^2, x_i \rangle_i) (\langle \theta_t^1 - \theta_t^2, x_i \rangle_i)^\top \right) \right] , \end{aligned}$$

the last line by the fact that $\text{Tr}(AB) \leq \text{Tr}(A)\text{Tr}(B)$ for A and B positive semi-definite. We thus conclude that

$$\frac{d}{dt} \mathbb{E} \left(\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q}(\theta_t^1 - \theta_t^2) \rangle \right) \leq \frac{4\gamma}{n} \mathbb{E} \|\mathbf{X}(\theta^1 - \theta^2)\|^2 \mathbf{K}_{p/q} ,$$

where we have used that for all $i \in \llbracket 1, n \rrbracket$, we have $\langle x_i, \Sigma^{-p/q} x_i \rangle \leq \mathbf{K}_{p/q}$. In addition, by (35), we get

$$\int_0^t \mathbb{E} [\|\mathbf{X}(\theta_u^1 - \theta_u^2)\|^2] du \leq \frac{n}{2 - 4\gamma\mathbf{K}} \mathbb{E} \|\theta_0^1 - \theta_0^2\|^2 .$$

Collecting the estimates, we thus get

$$\frac{1}{n} \mathbb{E} [\|\mathbf{X}(\theta_t^1 - \theta_t^2)\|^2] \geq C \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2]^{1/p},$$

with $C = \left(\mathbb{E} [\langle \theta_0^1 - \theta_0^2, \Sigma^{-p/q}(\theta_0^1 - \theta_0^2) \rangle] + 2\gamma \frac{\mathbf{K}_{p/q}}{1-2\gamma\mathbf{K}} \mathbb{E}(\|\theta_0^1 - \theta_0^2\|^2) \right)^{-q/p}$, which combined with (35), gives

$$\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq \mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2] - C(2 - 4\gamma\mathbf{K}) \int_0^t \mathbb{E} \|\theta_u^1 - \theta_u^2\|^2 du,$$

that implies, from a slight modification of Gronwall Lemma that for all $t \geq 0$ (for $C = 2C(1 - 2\gamma\mathbf{K})$), we have

$$\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq \left[\mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2]^{1-1/p} + (1/p - 1)Ct \right]^{\frac{1}{1-1/p}},$$

and we conclude as for (i). ■

A.2.2 LOCALIZATION OF THE INVARIANT MEASURE.

We now turn in the localization of the invariant measure presented in Proposition 8.

Proof [Proposition 8] We proved that θ_t tends weakly to Θ_* for $\gamma < 1/\mathbf{K}$. To insure that the first and second moment of θ_t tends to the first and two moments of Θ_* , we first prove that there exists a M such that for all t , $\mathbb{E} \|\theta_t - \theta_*\|^4 \leq M$. By the Itô formula,

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[V(\theta)] &= \frac{d}{dt} \frac{1}{2} \mathbb{E}[\|\theta_t - \theta_*\|^2] \leq \mathbb{E} \left(-\frac{1}{n} \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{\gamma\mathbf{K}}{2n} \|\mathbf{X}\theta_t - y\|^2 \right) \\ &\leq \mathbb{E} \left(\left(-\frac{1}{n} + \frac{\gamma\mathbf{K}}{n}\right) \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{\gamma\mathbf{K}}{n} \|\mathbf{X}\theta_t^* - y\|^2 \right) \\ &\leq (-2 + 2\gamma\mathbf{K})\mu V(\theta) + 2\gamma\mathbf{K}\sigma^2, \end{aligned}$$

the second line by triangle inequality and the third line because $L(\theta_*) = \sigma^2$. This implies by Gronwall Lemma that

$$\mathbb{E}[V(\theta)] \leq e^{-2(1-\gamma\mathbf{K})t} \left(\frac{1}{2} \|\theta_0 - \theta_*\|^2 - \frac{\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu} \right) + \frac{\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu}, \quad (38)$$

thus $t \mapsto \mathbb{E}[\|\theta_t\|^2]$ is uniformly upper bounded for $\gamma < 1/\mathbf{K}$. Again, by Ito formula, we get

$$\begin{aligned}
 \frac{d}{dt} \mathbb{E} \|\theta_t - \theta_*\|^4 &= -4 \mathbb{E} \|\theta_t - \theta_*\|^2 \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta_*) \rangle \\
 &\quad + \frac{\gamma}{2n} \mathbb{E} \text{Tr} \left(\mathbf{X}^\top R_{\mathbf{x}}(\theta_t) R_{\mathbf{x}}(\theta_t)^\top \mathbf{X} \left(8(\theta_t - \theta_*)(\theta_t - \theta_*)^\top + 4 \|\theta_t - \theta_*\|^2 I_d \right) \right) \\
 &\leq -\frac{4}{n} \mathbb{E} \|\theta_t - \theta_*\|^2 \|\mathbf{X}(\theta_t - \theta_*)\|^2 \\
 &\quad + \frac{\gamma}{2n} \mathbb{E} \text{Tr} \left(\mathbf{X}^\top \text{diag}(r_{\mathbf{x}}(\theta))^2 \mathbf{X} \left(8(\theta_t - \theta_*)(\theta_t - \theta_*)^\top + 4 \|\theta_t - \theta_*\|^2 I_d \right) \right) \\
 &\leq -\frac{4}{n} \mathbb{E} \|\theta_t - \theta_*\|^2 \|\mathbf{X}(\theta_t - \theta_*)\|^2 \\
 &\quad + \frac{8\gamma}{2n} \mathbb{E} \text{Tr} \left(\mathbf{X}^\top \text{diag}(r_{\mathbf{x}}(\theta))^2 \mathbf{X} \left((\theta_t - \theta_*)(\theta_t - \theta_*)^\top \right) \right) + \frac{4\gamma\mathbf{K}}{2n} \mathbb{E} \|\theta_t - \theta_*\|^2 \|\mathbf{X}\theta_t - y\|^2 \\
 &\leq -\frac{4}{n} \mathbb{E} \|\theta_t - \theta_*\|^2 \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{6\gamma\mathbf{K}}{n} \mathbb{E} \|\mathbf{X}\theta_t - y\|^2 \|\theta_t - \theta_*\|^2,
 \end{aligned}$$

the last line by the fact that $\text{Tr}(AB) \leq \text{Tr}(A)\text{Tr}(B)$ for A and B positive semi-definite and we deduce with the triangle inequality that the latter is less than

$$\begin{aligned}
 \frac{d}{dt} \mathbb{E} \|\theta_t - \theta_*\|^4 &\leq \left(-\frac{4}{n} + \frac{12\gamma\mathbf{K}}{n} \right) \mathbb{E} \|\theta_t - \theta_*\|^2 \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{12\gamma\mathbf{K}}{n} \mathbb{E} \|\mathbf{X}\theta_t - y\|^2 \|\theta_t - \theta_*\|^2 \\
 &\leq (-4 + 12\gamma\mathbf{K})\mu \mathbb{E} \|\theta_t - \theta_*\|^4 + 24\gamma\mathbf{K}\sigma^2 \mathbb{E} \|\theta_t - \theta_*\|^2,
 \end{aligned}$$

the last term is bounded by (38) and we conclude with Gronwall Lemma for $\gamma < \frac{1}{3\mathbf{K}}$ the claim on the fourth moment.

For the mean, we consider the equation satisfies by $\mathbb{E}(\theta_t)$: $\frac{d}{dt} \mathbb{E}(\theta_t) = -\nabla L(\mathbb{E}(\theta_t))$, by linearity of $\theta \mapsto \nabla L(\theta)$. Hence, we have

$$\frac{d}{dt} \frac{1}{2} \|\mathbb{E}(\theta_t) - \theta_*\|^2 = -\langle \Sigma(\mathbb{E}(\theta_t) - \theta_*), \mathbb{E}(\theta_t) - \theta_* \rangle \leq -\mu \|\mathbb{E}(\theta_t) - \theta_*\|^2,$$

which gives that $\mathbb{E}(\theta_t) \rightarrow \theta_*$ as t goes to infinity. Combining the latter with the boundedness of the fourth moment (One actually needs only a $1 + \epsilon$ moment) and the weak convergence of θ_t to Θ_* imply the first claim of the proposition. By (38), we recall that

$$\mathbb{E}[V(\theta)] \leq e^{-2(1-\gamma\mathbf{K})t} \left(\frac{1}{2} \mathbb{E} \|\theta_0 - \theta_*\|^2 - \frac{\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu} \right) + \frac{\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu},$$

which combined with the boundedness of the fourth moment implies by taking the limit in the inequality, and weak convergence of ρ_t towards ρ^* that $\theta_* \sim \rho^*$ satisfies the claimed inequality in the proposition. \blacksquare

We turn into proving the moment explosion of the invariant distribution. This corresponds to proving Proposition 10. First, we need the following lemma that aims at lower bounding some quartic form of θ :

Lemma 16 *For $n \geq 2d$, there exists a constant $c > 0$ that depends only on \mathbf{X}, y such that $\forall \eta \in \mathbb{R}^d$,*

$$\langle \mathbf{X}\eta, R_{\mathbf{x}}^2 \mathbf{X}\eta \rangle \geq c \|\eta\|^2 \|\mathbf{X}\eta\|^2.$$

We assume the validity of Lemma 16 for the time being and turn into the proof of Proposition 10.

Proof of Proposition 10 Using Itô's Lemma, we obtain that

$$\begin{aligned} \frac{d}{dt} \mathbb{E} \left[\frac{1}{2} \|\eta_t\|^{2\alpha} \right] &= -\alpha \mathbb{E} \|\eta_t\|^{2(\alpha-1)} \langle \Sigma \eta_t, \eta_t \rangle + \mathbb{E} \frac{\gamma \alpha (\alpha - 1) \|\eta_t\|^{2(\alpha-2)}}{n} \text{Tr} \left(\mathbf{X}^\top R_x^2 \mathbf{X} \eta_t \eta_t^\top \right) \\ &\quad + \mathbb{E} \frac{\gamma \alpha \|\eta_t\|^{2(\alpha-1)}}{2n} \text{Tr} \left(\mathbf{X}^\top R_x^2 \mathbf{X} \right). \end{aligned}$$

We suppose that $\|\Theta_*\|$ has moments of order 2α . We take $\eta_0 = \Theta_*$ and thus obtain

$$\mathbb{E} \|\Theta_*\|^{2(\alpha-1)} \langle \Sigma \Theta_*, \Theta_* \rangle \geq \frac{\gamma(\alpha-1)}{n} \mathbb{E} \|\Theta_*\|^{2(\alpha-2)} \text{Tr} \left(\mathbf{X}^\top R_x^2 \mathbf{X} \Theta_* \Theta_*^\top \right).$$

Using the Lemma 16, we deduce that

$$\mathbb{E} \|\Theta_*\|^{2\alpha} \geq c \frac{\mu}{\lambda} \gamma (\alpha - 1) \mathbb{E} \|\Theta_*\|^{2\alpha},$$

which leads to a contradiction for α large enough and we thus deduce that $\|\Theta_*\|$ has no moments of order 2α for $\alpha > \frac{2}{\gamma} \left(1 + \frac{\lambda}{\mu c} \right)$. \blacksquare

It remains to prove Lemma 16. For convenience, we now denote the residual $r = X\theta - y \in \mathbb{R}^n$, and the constant noise vector $\sigma = X\theta^* - y \in \mathbb{R}^n$, and note in this proof $R_x^2 \equiv R_r^2$ and $X\eta = r - \sigma$ to emphasize that everything here can be expressed as a function of the residuals r .

Proof of Lemma 16. We first give the kernel of R^2 as a function of r . Let $I = \{i \in \llbracket 1, n \rrbracket, r_i = 0\}$ and $\alpha_r = (1/r_1, 1/r_2, \dots, 1/r_n)^\top$, then

- (i) if $I \neq \emptyset$, $\ker(R_r^2) = \text{span}(e_i)_{i \in I}$
- (ii) if $I = \emptyset$, $\ker(R_r^2) = \text{span}(\alpha_r)$.

We fix $r \in \mathbb{R}^n$ and let $z \in \ker(R_r^2)$, that is

$$\left(\text{diag}(r)^2 - \frac{1}{n} r r^\top \right) z = 0,$$

i.e. for all $i \in \llbracket 1, n \rrbracket$,

$$r_i^2 z_i = \frac{\langle r, z \rangle}{n} r_i.$$

If $I = \emptyset$, we have the relationship for all $i, j \in \llbracket 1, n \rrbracket$, $r_i z_i = \frac{\langle r, z \rangle}{n} = r_j z_j$, which instantly gives the result. Otherwise if $I \neq \emptyset$, for all $i \in I^c$, we have $r_i z_i = \frac{\langle r, z \rangle}{n}$ and summing all these equality of $i \in I^c$, we have

$$\langle r, z \rangle = \sum_{i \in I^c} r_i z_i = (n - |I|) \frac{\langle r, z \rangle}{n},$$

which implies, as $|I| \neq n$, that $z \perp r$. Recalling that for $i \in \mathcal{I}^c$, we have $z_i = \frac{\langle r, z \rangle}{r_i n} = 0$ and (i),(ii) follow.

Let us remark that $r \mapsto R_r^2$ carries the homogeneity property that for all $r \in \mathbb{R}^n$,

$$R_{r/\|r\|}^2 = R_r^2 / \|r\|^2 ,$$

so that denoting the sphere $\mathcal{S}_{n-1} := \{v \in \mathbb{R}^n, \text{ s.t. } \|v\| = 1\}$, we have for $u = r/\|r\| \in \mathcal{S}_{n-1}$ and $v = (r - \sigma)/\|r - \sigma\| \in \mathcal{S}_{n-1} \cap \text{Ran}(\mathbf{X})$ (note that v is defined i.f.f. η is not equal to 0), we have

$$\frac{\langle \mathbf{X}\eta, R_x^2 \mathbf{X}\eta \rangle}{\|\mathbf{X}\eta\|^2 \|\mathbf{X}\eta + \sigma\|^2} = \langle v, R_u^2 v \rangle . \quad (39)$$

Now we also need the following intermediate result: Let Π_u be the orthogonal projector into $\ker(R_u^2)^\perp$, we have that there exists a constant $c > 0$ that depends only on the data such that for all $\eta \in \mathbb{R}^d$,

$$\|\Pi_u v\| \geq c .$$

There are two cases given by (i) and (ii), that allow to compute explicitly Π_u for all $u \in \mathbb{R}^n$:

(i) If $I \neq \emptyset$, we have

$$\|\Pi_u v\|^2 = \frac{\langle \sum_{i \in I^c} x_i x_i^\top \eta, \eta \rangle}{\langle \sum_{i=1}^n x_i x_i^\top \eta, \eta \rangle} \geq c ,$$

because any covariance matrix of the form $\sum_{i \in I^c} x_i x_i^\top$ is invertible if $|I^c| \geq d$, which is verified for $n \geq 2d$.

(ii) If $I = \emptyset$,

$$\|\Pi_u v\|^2 = 1 - \frac{\langle v, \alpha_r \rangle^2}{\|\alpha_r\|^2} = 1 - \varphi(\eta) ,$$

where

$$\varphi(\eta) = \frac{\left(\sum_{i=1}^n \frac{\langle x_i, \eta \rangle}{\langle x_i, \eta \rangle - \sigma_i} \right)^2}{\left(\sum_{i=1}^n \langle x_i, \eta \rangle^2 \right) \left(\sum_{i=1}^n (\langle x_i, \eta \rangle - \sigma_i)^{-2} \right)} .$$

From the Cauchy-Schwarz inequality, $\varphi \in [0, 1]$ and let $m = \sup_{\eta \in \mathbb{R}^d} \varphi(\eta)$. We show below that $m < 1$.

First, let's study the situation *at infinity*. For this, we write $\eta = \bar{\eta} \|\eta\|$, where $\bar{\eta} = \eta / \|\eta\| \in \mathcal{S}_{d-1}$ and call $J = \{i \in \llbracket 1, n \rrbracket, \text{ such that } \langle x_i, \bar{\eta} \rangle = 0\}$. We know that $|J| \leq d$. We fix $\bar{\eta}$, and study the limit of φ as $\|\eta\|$ grows.

If $J \neq \emptyset$, we have

$$\varphi(\eta) \sim_{\|\eta\| \rightarrow \infty} \frac{(n - |J|)^2}{\left(\sum_{i \in J^c} \langle x_i, \eta \rangle^2 \right) \left(\sum_{i \in J} \sigma_i^{-2} \right)} .$$

Hence, if $J \neq \emptyset$, then $\varphi(\eta) \rightarrow 0$ when $\|\eta\| \rightarrow \infty$. Now if $J = \emptyset$,

$$\varphi(\eta) \sim_{\|\eta\| \rightarrow \infty} \frac{n^2}{\left(\sum_{i=1}^n \langle x_i, \bar{\eta} \rangle^2 \right) \left(\sum_{i=1}^n \langle x_i, \bar{\eta} \rangle^{-2} \right)} := \phi(\bar{\eta}) .$$

Moreover, $M := \sup_{\bar{\eta} \in \mathcal{S}_{d-1}} \phi(\bar{\eta}) < 1$, because if $M = 1$, it would be attained for a certain $\bar{\eta}_*$ by compactness of \mathcal{S}_{d-1} and continuity of ϕ , and if $M = 1$, then $\phi(\bar{\eta}_*) = 1$ and by the

equality case of Cauchy-Schwarz, this corresponds to the fact that there exists $\lambda \in \mathbb{R}$ such that $|\langle \mathbf{X}\bar{\eta}_* | \rangle| = \lambda$, which has no solution for $n > d$ (recall that because $\bar{\eta}_* \in \mathcal{S}_{d-1}$, $\lambda = 0$ offer no solution either). We conclude from this that $M < 1$.

Now either the supremum of φ is M , and in this case we have $m = M < 1$ and we are done with the proof. Either, $m > M$, and in this case, the supremum is attained in a compact set of \mathbb{R}^d , hence attained for a certain $\eta_* \in \mathbb{R}^d$ by continuity of φ . As previously, if $m = 1$, this corresponds to the equality case of Cauchy-Schwarz and hence the existence of $\lambda \in \mathbb{R}$ such that

$$(\langle x_i, \eta \rangle - \sigma_i)^2 \langle x_i, \eta \rangle^2 = \lambda ,$$

which has no solution generically for $n > d$. Finally $m = 1$ is impossible and we have proven that $m < 1$, which corresponds to $c = 1 - m > 0$.

Now if we take back Eq.(39), and remark that

$$\langle v, R_u^2 v \rangle = \langle \Pi_u v, R_u^2 \Pi_u v \rangle = \|\Pi_u\|^2 \langle \bar{\Pi}_u v, R_u^2 \bar{\Pi}_u v \rangle \geq c \langle \bar{\Pi}_u v, R_u^2 \bar{\Pi}_u v \rangle \geq c \lambda_{\min}(R_{u|_{\ker(R_u^2)^\perp}}^2).$$

Now define $\Xi : \mathcal{S}_{n-1} \rightarrow \mathbb{R}_{\geq 0}$ as $\Xi(u) = \lambda_{\min}(R_{u|_{\ker(R_u^2)^\perp}}^2)$. We look for $\ell := \min_{u \in \mathcal{S}_{n-1}} \Xi(u)$. As Ξ is a continuous function defined on a compact, it attains its minimum within this compact. If $\ell = 0$, then this means that there exists $u^* \in \mathcal{S}_{n-1}$ such that $\Xi(u^*) = \lambda_{\min}(R_{u^*|_{\ker(R_{u^*}^2)^\perp}}^2) = 0$, which is impossible by the fact that we have restricted R_u^2 orthogonally to its kernel. Hence $\ell > 0$ and finally we have

$$\frac{\langle \mathbf{X}\eta, R_x^2 \mathbf{X}\eta \rangle}{\|\mathbf{X}\eta\|^2 \|\mathbf{X}\eta + \sigma\|^2} = \langle v, R_u^2 v \rangle \geq c\ell > 0 ,$$

which leads to the proof of the lemma considering that there exists $K > 0$ such that $\|\mathbf{X}\eta + \sigma\|^2 \geq K\|\eta\|^2$ and noting $c = c\ell K$. Indeed, combining the fact that $\langle \Sigma\eta_t, \eta_t \rangle \leq 4(L(\theta_t) + L(\theta_*))$ by the triangle inequality and the fact that $L(\theta^*) \leq L(\theta_t)$, we obtain that $\langle \Sigma\eta_t, \eta_t \rangle \leq 8L(\theta_t)$, which implies that

$$\frac{\mu n}{4} \langle \eta_t, \eta_t \rangle \leq \|\mathbf{X}\eta + \sigma\|^2.$$

■

A.2.3 CONVERGENCE OF VARIANCE REDUCTION TECHNIQUES

Now, we turn into the convergence of the time-average of the iterates. This corresponds to the proof of Proposition 11.

Proof [Proposition 11] The idea we use to show the quantitative convergence of the average iterates comes from Mattingly et al. (2010) and the use of Poisson equation. Indeed, let us define the map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\varphi(\theta) = \Sigma^{-1}\theta$, and note $(\varphi_i(\theta))_{i \in \llbracket 1, d \rrbracket}$ its coordinates. Note that, for all $i \in \llbracket 1, d \rrbracket$, we have $\nabla \varphi_i(\theta) = \Sigma^{-1}e_i$, where $(e_i)_{i \in \llbracket 1, d \rrbracket}$ is the canonical basis of \mathbb{R}^d . Hence, we have the Poisson equation, for all $i \in \llbracket 1, d \rrbracket$ and $\theta \in \mathbb{R}^d$,

$$\mathcal{L}\varphi_i(\theta) = -\langle \Sigma(\theta - \theta_*), \Sigma^{-1}e_i \rangle = \theta_i - \theta_{*,i} .$$

Hence considering the action of \mathcal{L} on fields of \mathbb{R}^d applied coordinate-wise, we have more generally $\mathcal{L}\varphi(\theta) = \theta - \theta_*$. Thus, we have by Itô calculus that for all $s \geq 0$

$$d\varphi(\theta_s) = (\theta_s - \theta_*)ds + \sqrt{\frac{\gamma}{n}}\Sigma^{-1}\mathbf{X}^\top R_x(\theta_s)dB_s$$

and thus integrating from 0 to t and dividing by t , and defining the martingale $M_t = -\frac{1}{\sqrt{n}}\mathbf{X}^\top \int_0^t R_x(\theta_s)dB_s$ we have,

$$\bar{\theta}_t - \theta_* = \frac{1}{t}(\varphi(\theta_t) - \varphi(\theta_0)) + \frac{\sqrt{\gamma}}{t}\Sigma^{-1}M_t = \frac{\Sigma^{-1}(\theta_t - \theta_0)}{t} + \frac{\sqrt{\gamma}}{t}\Sigma^{-1}M_t .$$

Hence, multiplying by Σ , taking norms and then expectations, we have

$$\mathbb{E}\|\Sigma(\bar{\theta}_t - \theta_*)\|^2 \leq \frac{2\mathbb{E}\|\theta_t - \theta_0\|^2}{t^2} + \frac{2\gamma}{t^2}\mathbb{E}\|M_t\|^2 .$$

Then, by Itô isometry, we have

$$\begin{aligned} \mathbb{E}\|M_t\|^2 &= \frac{1}{n}\mathbb{E}\left[\int_0^t \|\mathbf{X}^\top R_x(\theta_s)\|_{\text{HS}}^2 ds\right] \\ &= \frac{1}{n}\mathbb{E}\left[\int_0^t \text{Tr}\left(\mathbf{X}\mathbf{X}^\top R_x^2(\theta_s)\right) ds\right] \\ &\leq 2\mathbf{K}\int_0^t \mathbb{E}[L(\theta_s)] ds . \end{aligned}$$

Moreover, we have seen that in Lemma 6 that for $\gamma \leq \mathbf{K}^{-1}$, we have

$$\frac{d}{dt}\frac{1}{2}\mathbb{E}\|\theta_t - \theta_*\|^2 \leq -\mathbb{E}L(\theta_t) + 2\sigma^2 ,$$

and hence we have the upper bound on the loss $\int_0^t \mathbb{E}[L(\theta_s)] ds \leq \frac{1}{2}\|\theta_0 - \theta_*\|^2 + 2\sigma^2 t$. Finally, this gives that

$$\mathbb{E}\|M_t\|^2 \leq \mathbf{K}\|\theta_0 - \theta_*\|^2 + 4\mathbf{K}\sigma^2 t .$$

On the other hand, we have for all $t \geq 0$

$$\begin{aligned} \mathbb{E}\|\theta_t - \theta_0\|^2 &\leq 2\mathbb{E}\|\theta_t - \theta_*\|^2 + 2\mathbb{E}\|\theta_* - \theta_0\|^2 \\ &\leq 4e^{-2(1-\gamma\mathbf{K})t}\left(\frac{1}{2}\|\theta_0 - \theta_*\|^2 - \frac{\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu}\right) + \frac{4\gamma\mathbf{K}\sigma^2}{(1-\gamma\mathbf{K})\mu} + 2\|\theta_* - \theta_0\|^2 \\ &\leq 4\|\theta_0 - \theta_*\|^2, \end{aligned}$$

the second line by (38) for $\gamma < 1/\mathbf{K}$. Overall, we have the bound,

$$\mathbb{E}\|\Sigma(\bar{\theta}_t - \theta_*)\|^2 \leq \frac{8\gamma\mathbf{K}\sigma^2}{t} + \frac{10\|\theta_0 - \theta_*\|^2}{t^2} .$$

■

Now, we turn into the convergence of θ_t in the case of the step-size decay. This corresponds to the proof of Proposition 12.

Proof [Proposition 12] By Ito formula, we obtain that

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[\|\theta_t - \theta_*\|^2] &\leq \mathbb{E} \left(-\frac{2}{n} \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{\gamma_t \mathbf{K}}{n} \|\mathbf{X}\theta_t - y\|^2 \right) \\ &\leq \mathbb{E} \left(\left(-\frac{2}{n} + \frac{2\gamma_t \mathbf{K}}{n} \right) \|\mathbf{X}(\theta_t - \theta_*)\|^2 + \frac{2\gamma_t \mathbf{K}}{n} \|\mathbf{X}\theta_t^* - y\|^2 \right) \\ &\leq (-2 + 2\gamma_t \mathbf{K}) \mu \mathbb{E}[\|\theta_t - \theta_*\|^2] + 4\gamma_t \mathbf{K} \sigma^2 \\ &\leq -\mu \mathbb{E}[\|\theta_t - \theta_*\|^2] + 4\gamma_t \mathbf{K} \sigma^2 . \end{aligned}$$

Applying the Gronwall lemma, we get

$$\begin{aligned} \mathbb{E}[\|\theta_t - \theta_*\|^2] &\leq \mathbb{E}[\|\theta_0 - \theta_*\|^2] e^{-\mu t} + \int_0^t e^{-\mu(t-s)} 4\mathbf{K} \sigma^2 / (2\mathbf{K} + s^\alpha) ds \\ &\leq \mathbb{E}[\|\theta_0 - \theta_*\|^2] e^{-\mu t} + e^{-\mu t/2} \int_0^{t/2} 4\mathbf{K} \sigma^2 / (2\mathbf{K}) ds + \int_{t/2}^t 4\mathbf{K} \sigma^2 / s^\alpha ds \\ &\leq (\mathbb{E}[\|\theta_0 - \theta_*\|^2] + t\sigma^2) e^{-\mu t/2} + 4\mathbf{K} \sigma^2 \frac{(t/2)^{1-\alpha}}{\alpha - 1} \\ &\leq \frac{1}{t^{\alpha-1}} \left(e^{-\alpha} (\mathbb{E}[\|\theta_0 - \theta_*\|^2]) e^{2(\alpha-1)/\mu} + (2\alpha/\mu)^\alpha \sigma^2 \right) + \frac{2^{1+\alpha} \mathbf{K} \sigma^2}{(\alpha-1)} . \end{aligned}$$

■

Appendix B. Proofs of Section 5 : online SGD

B.1 The noiseless case

We begin this proof's section with the proof of convergence in the noiseless case. This corresponds to Theorem 13.

Proof [Theorem 13] (i) The key ingredient of the proof is the Gronwall Lemma. Combining the Itô formula with (13) gives us

$$\begin{aligned} \frac{d}{dt} \mathbb{E} \|\theta_t - \theta_*\|^2 &= -2\mathbb{E} \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta^*) \rangle + \gamma \mathbb{E} \text{Tr} \left(\sigma(\theta_t) \sigma(\theta_t)^\top \right) \\ &\leq -2\mathbb{E} \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta^*) \rangle + \gamma \mathbb{E} \left[\text{Tr} \left((\langle \theta_t, X \rangle - Y)^2 X X^\top \right) \right] \\ &\leq -2\mathbb{E} \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta^*) \rangle + \gamma \mathbf{K} \mathbb{E} \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta^*) \rangle \\ &\leq -(2 - \gamma \mathbf{K}) \mathbb{E} \langle \theta_t - \theta_*, \Sigma(\theta_t - \theta^*) \rangle \\ &\leq -\mu(2 - \gamma \mathbf{K}) \mathbb{E} \|\theta_t - \theta_*\|^2 . \end{aligned}$$

By integrating the latter thanks to Gronwall Lemma, we get the result claim for (i).

(ii) We define $\eta_t := \theta_t - \theta_*$ and recall that, thanks to the penultimate inequality of (i) we have by integration that:

$$\mathbb{E}\|\eta_t\|^2 \leq \|\theta_0 - \theta_*\|^2 - 2(2 - \gamma\mathbf{K}) \int_0^t \mathbb{E}L(\theta_u) du.$$

The first consequence of this inequality is that

$$\int_0^t \mathbb{E}L(\theta_u) du \leq \frac{1}{2(2 - \gamma\mathbf{K})} \|\theta_0 - \theta_*\|^2.$$

We want to lower bound the term $\mathbb{E}L(\theta_u)$ without using the smallest eigenvalue of Σ , that is supposedly infinitely small. Similarly to what was done before, thanks to the Hölder inequality, if $p, q \in (0, 1)$, $p + q = 1$, we have, for all $t \geq 0$,

$$\left(\mathbb{E}[\|\eta_t\|^2]\right)^{\frac{1}{p}} \leq 2\mathbb{E}[L(\theta_t)] \left(\mathbb{E}\left[\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle\right]\right)^{q/p}. \quad (40)$$

We now prove that $t \mapsto \mathbb{E}\left(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle\right)$ is bounded. Indeed, we proceed as before with Itô formula to get

$$\begin{aligned} \frac{d}{dt} \mathbb{E}\left(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle\right) &= -2\mathbb{E}\langle \eta_t, \Sigma^{1-p/q} \eta_t \rangle + \gamma \mathbb{E} \text{Tr}\left(\sigma \sigma^\top \Sigma^{-p/q}\right) \\ &\leq \gamma \mathbb{E} \text{Tr}\left(\Sigma^{-p/q} (\langle \theta_t, X \rangle - Y)^2 X X^\top\right) \\ &\leq \gamma \mathbb{E}\left[\langle X, \Sigma^{-p/q} X \rangle (\langle \theta_t, X \rangle - Y)^2\right] \\ &\leq 2\gamma \mathbf{K}_{p/q} L(\theta_t), \end{aligned}$$

where we have used that almost surely, we have $\langle X, \Sigma^{-p/q} X \rangle \leq \mathbf{K}_{p/q}$. Then, by integrating with respect to t , it yields

$$\begin{aligned} \mathbb{E}\left(\langle \eta_t, \Sigma^{-p/q} \eta_t \rangle\right) &\leq \langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + 2\gamma \mathbf{K}_{p/q} \int_0^t \mathbb{E}L(\theta_u) du \\ &\leq \langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + \frac{\gamma \mathbf{K}_{p/q}}{2 - \gamma\mathbf{K}} \|\theta_0 - \theta_*\|^2. \end{aligned}$$

Hence, calling $\mathbf{C} = \frac{1}{2}(\langle \eta_0, \Sigma^{-p/q} \eta_0 \rangle + \frac{\gamma \mathbf{K}_{p/q}}{2 - \gamma\mathbf{K}} \|\theta_0 - \theta_*\|^2)^{-p/q}$, we have the inequality, for all $t \geq 0$,

$$\mathbb{E}L(\theta_t) \geq \mathbf{C} (\mathbb{E}\|\eta_t\|^2)^{1/p},$$

and this yields the inequality

$$\mathbb{E}\|\eta_t\|^2 \leq \|\eta_0\|^2 - \mathbf{C} \int_0^t (\mathbb{E}\|\eta_u\|^2)^{1/p} du, \quad (41)$$

that implies, from a slight modification of Gronwall Lemma that for all $t \geq 0$, we have

$$\mathbb{E}\|\eta_t\|^2 \leq \left[\frac{1}{\frac{1}{\|\eta_0\|^{2(1/p-1)}} + (1/p-1)Ct} \right]^{\frac{1}{1/p-1}},$$

this gives the result claim in the theorem. To see how the last inequality goes, we define $g(t) = \|\eta_0\|^2 - C \int_0^t (\mathbb{E}\|\eta_u\|^2)^{1/p} du$ (which is positive) and we rewrite (41) as

$$\left(\frac{g'(t)}{-C} \right)^p \leq g(t) \iff \frac{g'(t)}{g(t)^{1/p}} \geq -C \implies \frac{1}{-1/p+1} (g(t)^{-1/p+1} - g(0)^{-1/p+1}) \geq -Ct,$$

and thus

$$g(t) \leq \left[-Ct \left(-\frac{1}{p} + 1\right) + \|\eta_0\|^{2(-1/p+1)} \right]^{\frac{1}{-1/p+1}},$$

and we conclude with (41).

To prove the convergence almost surely, we use the Itô Formula to obtain

$$\begin{aligned} \mathbb{E}(\|\theta_t - \theta_*\|^2 | \mathcal{F}_s) &= \|\theta_0 - \theta_*\|^2 + 2\sqrt{\gamma} \int_0^s \langle (\theta_u - \theta_*), \sigma(\theta_u) dB_u \rangle \\ &\quad + \gamma \mathbb{E} \left[\int_0^t \left[\text{Tr}(\sigma \sigma^\top(\theta_u)) - 4L(\theta_u) \right] du | \mathcal{F}_s \right], \end{aligned}$$

For $\gamma < \frac{1}{2K}$, we have proven in (i) that the term inside the conditional expectation value/integral is negative. We can thus overestimate the latter by integrating from 0 to s to obtain

$$\mathbb{E}(\|\theta_t - \theta_*\|^2 | \mathcal{F}_s) \leq \|\theta_s - \theta_*\|^2,$$

and we deduce that $\|\theta_t - \theta_*\|^2$ is a positive supermartingale, L^1 bounded and therefore converges almost surely to 0. \blacksquare

B.2 The noisy case

We begin this section by proving that the multiplicative noise carries some Lipschitz property. This corresponds to Lemma 14.

Proof [Lemma 14] Recall that the diffusion matrix writes

$$\begin{aligned} \sigma(\theta) &= (f(\theta) - g(\theta))^{1/2}, \quad \text{with} \\ f(\theta) &= \mathbb{E}_\rho \left[(\langle \theta, X \rangle - Y)^2 X X^\top \right], \quad \text{and} \quad g(\theta) = \mathbb{E}_\rho [(\langle \theta, X \rangle - Y) X] \mathbb{E}_\rho [(\langle \theta, X \rangle - Y) X]^\top. \end{aligned}$$

Let us introduce for all $\theta \in \mathbb{R}^d$, the differential operator of σ , that is $d\sigma_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ such that for all $h \in \mathbb{R}^d$,

$$\sigma(\theta + h) = \sigma(\theta) + d\sigma_\theta(h) + o(\|h\|).$$

We calculate, for all $\theta, h \in \mathbb{R}^d$,

$$\begin{aligned} \sigma(\theta + h) &= (f(\theta + h) - g(\theta + h))^{1/2} \\ &= (f(\theta) - g(\theta) + l_\theta(h) + o(\|h\|))^{1/2} \end{aligned}$$

where $l_\theta(h) = df_\theta(h) - dg_\theta(h)$ and

$$\begin{aligned} df_\theta(h) &= 2\mathbb{E}_\rho \left[(\langle \theta, X \rangle - Y) \langle h, X \rangle X X^\top \right] \\ dg_\theta(h) &= \mathbb{E}_\rho [\langle h, X \rangle X] \mathbb{E}_\rho [(\langle \theta, X \rangle - Y) X]^\top + \mathbb{E}_\rho [(\langle \theta, X \rangle - Y) X] \mathbb{E} [\langle h, X \rangle X]^\top . \end{aligned}$$

Hence, introducing the square root operator $\psi : \mathbb{S}_+^d \rightarrow \mathbb{S}_+^d$, such that for all $M \in \mathbb{S}_+^d$, $\psi(M) = M^{1/2}$. We have, for all $\theta, h \in \mathbb{R}^d$,

$$d\sigma_\theta(h) = d\psi_{\sigma^2(\theta)}(l_\theta(h)),$$

that is that $d\sigma_\theta(h)$ is the unique solution to the matrix Lyapunov equation

$$\sigma(\theta) d\sigma_\theta(h) + d\sigma_\theta(h) \sigma(\theta) = l_\theta(h) ,$$

or equivalently, expressed in close form as the expression

$$d\sigma_\theta(h) = \int_0^{+\infty} e^{-s\sigma(\theta)} l_\theta(h) e^{-s\sigma(\theta)} ds .$$

That being written. Let us pose for $t \in [0, 1]$, the function $\Psi(t) = \sigma(t\theta + (1-t)\eta)$, we have

$$\sigma(\theta) - \sigma(\eta) = \Psi(1) - \Psi(0) = \int_0^1 \Psi'(u) du ,$$

and knowing that $\Psi'(u) = d\sigma_{m_u}(\theta - \eta)$, where $m_u = u\theta + (1-u)\eta$, and hence using the triangular inequality for the Hilbert-Schmidt norm, we have

$$\|\sigma(\theta) - \sigma(\eta)\|_{\text{HS}} \leq \int_0^1 \|d\sigma_{m_u}(\theta - \eta)\|_{\text{HS}} du$$

where $(m_u)_{0 \leq u \leq 1}$ is a parametrization of the line joining θ and η : $m_u = u\theta + (1-u)\eta$.

Hence, it remains to upper bound the Hilbert-Schmidt norm of the differential of σ thanks to its integral representation presented above. In fact for all $\theta, h \in \mathbb{R}^d$, we have

$$\|d\sigma_\theta(h)\|_{\text{HS}} \leq \|l_\theta(h)\| \int_0^{+\infty} \|e^{-s\sigma(\theta)}\|_{\text{HS}}^2 ds .$$

Moreover, we have $\|l_\theta(h)\| \leq \|df_\theta(h)\| + \|dg_\theta(h)\|$ and using that $\|\cdot\|_{\text{HS}}$ is sub-multiplicative, we have

$$\begin{aligned} \|df_\theta(h)\| &\leq 2\mathbb{E} \left[|\langle \theta, X \rangle - Y| |\langle h, X \rangle| \|X X^\top\| \right] \\ &\leq 2\mathbf{K} \sqrt{\mathbb{E} [|\langle \theta, X \rangle - Y|^2]} \sqrt{\mathbb{E} [|\langle h, X \rangle|^2]} \\ &= 2\sqrt{2}\mathbf{K} \sqrt{L(\theta)} \sqrt{\langle \Sigma h, h \rangle} , \end{aligned}$$

and similarly,

$$\|dg_\theta(h)\| \leq 2\sqrt{2}\mathbf{K} \sqrt{L(\theta)} \sqrt{\langle \Sigma h, h \rangle} .$$

Finally, we have the bound

$$\int_0^{+\infty} \|e^{-s\sigma(\theta)}\|_{\text{HS}}^2 ds = \int_0^{+\infty} \text{Tr}(e^{-2s\sigma(\theta)}) ds = \frac{1}{2} \text{Tr}(\sigma^{-1}(\theta)) .$$

As $\sigma^2(\theta) \succeq a^2 L(\theta) I_d$, we deduce that

$$L(\theta) \text{Tr}(\sigma^{-1}(\theta))^2 \leq \text{Tr}(L(\theta)\sigma^{-2}(\theta)) d \leq \frac{d^2}{a^2} ,$$

the first inequality by Cauchy-Schwarz and the Lemma 14 follows. ■

We can now turn into the proof of the main theorem of this section: the proof of convergence in the noisy case. This corresponds to proving Theorem 15.

Proof [Theorem 15] (i) Wassertein contraction comes essentially from *coupling arguments*. Let $\gamma \leq K^{-1}$ and $\rho_0^1, \rho_0^2 \in \mathbb{P}_2(\mathbb{R}^d)$ two possible initial distributions. Then by (Villani et al., 2009, Theorem 4.1), there exists a couple of random variables θ_0^1, θ_0^2 such that $W_2^2(\rho_0^1, \rho_0^2) = \mathbb{E}[\|\theta_0^1 - \theta_0^2\|^2]$. Let $(\theta_t^1)_{t \geq 0}$ (resp. $(\theta_t^2)_{t \geq 0}$) be the solution of the SDE (13), sharing the same Brownian motion $(B_t)_{t \geq 0}$. Then, for all $t \geq 0$, the random variable (θ_t^1, θ_t^2) is a coupling between ρ_t^1 and ρ_t^2 , and hence

$$W_2^2(\rho_t^2, \rho_t^1) \leq \mathbb{E}[\|\theta_t^1 - \theta_t^2\|^2] .$$

Moreover, we denote by $\|\cdot\|_{\text{HS}}$ the Frobenius norm and by Itô formula, we have

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[\|\theta_t^1 - \theta_t^2\|^2] &= -2\mathbb{E}[\langle \theta_t^1 - \theta_t^2, \Sigma(\theta_t^1 - \theta_*) - \Sigma(\theta_t^2 - \theta_*) \rangle] + \gamma \mathbb{E}\|\sigma(\theta_t^1) - \sigma(\theta_t^2)\|_{\text{HS}}^2 \\ &\leq -2\mathbb{E}[\langle \Sigma(\theta_t^1 - \theta_t^2), \theta_t^1 - \theta_t^2 \rangle] + 2\gamma \text{Kc} \mathbb{E}[\langle \Sigma(\theta_t^1 - \theta_t^2), \theta_t^1 - \theta_t^2 \rangle] , \end{aligned}$$

thanks to Lemma 14. Hence, this gives the inequality:

$$\frac{d}{dt} \mathbb{E}[\|\theta_t^1 - \theta_t^2\|^2] \leq -2(1 - \gamma \text{Kc}) \mathbb{E}[\langle \Sigma(\theta_t^1 - \theta_t^2), \theta_t^1 - \theta_t^2 \rangle] \quad (42)$$

$$\leq -2\mu(1 - \gamma \text{Kc}) \mathbb{E}\|\theta_t^1 - \theta_t^2\|^2 , \quad (43)$$

and by Gronwall Lemma, denoting $c_\gamma = 2\mu(1 - \gamma \text{Kc})$ this gives that

$$W_2^2(\rho_t^1, \rho_t^2) \leq \mathbb{E}[\|\theta_t^1 - \theta_t^2\|^2] \leq e^{-c_\gamma t} \mathbb{E}[\|\theta_0^1 - \theta_0^2\|^2] = e^{-c_\gamma t} W_2^2(\rho_0^2, \rho_0^1) .$$

Now, for all $s \geq 0$ setting $\rho_0^1 = \rho_0 \in \mathcal{P}_2(\mathbb{R}^d)$ and $\rho_0^2 = \rho_s \in \mathbb{P}_2(\mathbb{R}^d)$, we have for all $t \geq 0$,

$$W_2^2(\rho_t, \rho_{t+s}) \leq e^{-c_\gamma t} W_2^2(\rho_0, \rho_s) ,$$

which shows that the process $(\rho_t)_{t \geq 0}$ is of Cauchy type, and since $(\mathbb{P}_2(\mathbb{R}^d), W_2)$ is a Polish space, $\rho_t \rightarrow \rho_* \in \mathbb{P}_2(\mathbb{R}^d)$ as t grows to infinity. Now, since there exists a stationary solution to the process, let us fix $\rho_0^1 = \rho_* \in \mathbb{P}_2(\mathbb{R}^d)$ and $\rho_0^2 = \rho_0 \in \mathbb{P}_2(\mathbb{R}^d)$. We have then,

$$W_2^2(\rho_t, \rho_*) \leq e^{-c_\gamma t} W_2^2(\rho_0, \rho_*) ,$$

which concludes the first part of the Theorem.

(ii) We will use the same steps as for the proof of Theorem 4 (ii). Again, one steadily checks that for $p, q \in (0, 1)$, $p + q = 1$, we have, for all $t \geq 0$,

$$\mathbb{E} \left[\|\Sigma^{1/2} (\theta_t^1 - \theta_t^2)\|^2 \right] \geq \frac{\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2]^{1/p}}{\mathbb{E} [\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q} (\theta_t^1 - \theta_t^2) \rangle]^{q/p}} . \quad (44)$$

By Ito formula, skipping the details, we get

$$\begin{aligned} \frac{d}{dt} \mathbb{E} \left(\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q} (\theta_t^1 - \theta_t^2) \rangle \right) &\leq \gamma \mathbb{E} \text{Tr} \left(\Sigma^{-p/q} (\sigma(\theta_1) - \sigma(\theta_2))^2 \right) \\ &\leq 2\gamma c_{p/q} \mathbf{K}_{p/q} \langle \Sigma(\theta_1 - \theta_2), \theta_1 - \theta_2 \rangle , \end{aligned}$$

thanks to assumption (30). In addition, by (42), we get

$$\int_0^t \mathbb{E} \left[\|\Sigma^{1/2} (\theta_u^1 - \theta_u^2)\|^2 \right] du \leq \frac{1}{2(1 - \gamma \mathbf{K}c)} \mathbb{E} \|\theta_0^1 - \theta_0^2\|^2 .$$

Collecting the estimates, we thus get

$$\mathbb{E} \left(\langle \theta_t^1 - \theta_t^2, \Sigma^{-p/q} (\theta_t^1 - \theta_t^2) \rangle \right) \leq \frac{\gamma c_{p/q} \mathbf{K}_{p/q}}{1 - \gamma \mathbf{K}c} \mathbb{E} \|\theta_0^1 - \theta_0^2\|^2 + \mathbb{E} \left(\langle \theta_0^1 - \theta_0^2, \Sigma^{-p/q} (\theta_0^1 - \theta_0^2) \rangle \right) .$$

That is to say that

$$\mathbb{E} \left[\|\Sigma^{1/2} (\theta_t^1 - \theta_t^2)\|^2 \right] \geq C \mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2]^{1/p} , \quad (45)$$

with $C = \left(\mathbb{E} [\langle \theta_0^1 - \theta_0^2, \Sigma^{-p/q} (\theta_0^1 - \theta_0^2) \rangle] + \frac{\gamma c_{p/q} \mathbf{K}_{p/q}}{1 - \gamma c \mathbf{K}} \mathbb{E} (\|\theta_0^1 - \theta_0^2\|^2) \right)^{-q/p}$, which combined with equation (42), gives

$$\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq \mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2] - 2C(1 - \gamma c \mathbf{K}) \int_0^t \mathbb{E} [\|\theta_u^1 - \theta_u^2\|^2]^{1/p} du ,$$

that implies, from a slight modification of Gronwall Lemma that for all $t \geq 0$ (for $C = 2C(1 - \gamma c \mathbf{K})$), we have

$$\mathbb{E} [\|\theta_t^1 - \theta_t^2\|^2] \leq \left[\mathbb{E} [\|\theta_0^1 - \theta_0^2\|^2]^{1-1/p} + (1/p - 1)Ct \right]^{\frac{1}{1-1/p}} ,$$

and we conclude as for (i). ■