



HAL
open science

Maximum entropy methods for texture synthesis: theory and practice

Valentin de Bortoli, Agnès Desolneux, Alain Durmus, Bruno Galerne, Arthur
Leclaire

► **To cite this version:**

Valentin de Bortoli, Agnès Desolneux, Alain Durmus, Bruno Galerne, Arthur Leclaire. Maximum entropy methods for texture synthesis: theory and practice. *SIAM Journal on Mathematics of Data Science*, 2020, 10.1137/19M1307731 . hal-03037099

HAL Id: hal-03037099

<https://hal.science/hal-03037099v1>

Submitted on 2 Dec 2020

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

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

Maximum entropy methods for texture synthesis: theory and practice.

Valentin De Bortoli¹, Agnès Desolneux¹, Alain Durmus¹, Bruno Galerne², and Arthur Leclaire³

¹Centre de mathématiques et de leurs applications, CNRS, ENS Paris-Saclay, Université Paris-Saclay, 94235, Cachan cedex, France.

²Institut Denis Poisson, Université d'Orléans, Université de Tours, CNRS

³Univ. Bordeaux, IMB, Bordeaux INP, CNRS, UMR 5251. F 33-400, TALENCE. FRANCE.

December 2, 2020

Abstract

Recent years have seen the rise of convolutional neural network techniques in exemplar-based image synthesis. These methods often rely on the minimization of some variational formulation on the image space for which the minimizers are assumed to be the solutions of the synthesis problem. In this paper we investigate, both theoretically and experimentally, another framework to deal with this problem using an alternate sampling/minimization scheme. First, we use results from information geometry to assess that our method yields a probability measure which has maximum entropy under some constraints in expectation. Then, we turn to the analysis of our method and we show, using recent results from the Markov chain literature, that its error can be explicitly bounded with constants which depend polynomially in the dimension even in the non-convex setting. This includes the case where the constraints are defined via a differentiable neural network. Finally, we present an extensive experimental study of the model, including a comparison with state-of-the-art methods and an extension to style transfer.

1 Introduction

Understanding texture formation is a crucial step towards a global theory of the human visual system as texture is an important perceptual cue. The more specific problem of exemplar-based texture synthesis arises in computer graphics where it is often desirable to be able to generate new large natural textures which look like an input image. This

application highlights the need of a mathematically sound model for texture generation as our only criterion for evaluating the performance of algorithms is via human inspection.

Two main approaches have been proposed in the literature: the patch-based methods [36, 35, 61, 62, 80, 60, 47, 58, 43] and the parametric ones [91, 40, 39, 15, 49, 96, 44, 55, 88, 89]. In the present paper, we are interested in the theoretical and visual properties of information-based parametric models. More precisely, we consider maximum entropy models. Indeed, the maximum entropy approach has the appealing property that the trade-off between innovation (maximizing the entropy) and the visual similarity with the input (geometrical or statistical feature constraints) is explicitly embedded in the model. There exist two main approaches for these maximum entropy formulation, the microcanonical model in which the constraints must be met almost surely and the macrocanonical model in which the constraints must be met in expectation [14]. Both share connections with statistical physics. In [14] the authors address the convergence of usual sampling scheme for the microcanonical model. One key contribution of the present paper is the derivation of a similar result for the macrocanonical model.

Contrary to the microcanonical model, the distribution of any macrocanonical model is a Gibbs measure, *i.e.* the exponential distribution of the features up to a scalar product with some parameters [53]. Our first contribution is to give explicit conditions on the features ensuring the existence of such a macrocanonical model, extending results from information geometry [20, 19, 18, 31].

Even if such a Gibbs measure exists we are facing two issues: 1) finding the optimal parameters, 2) sampling from the associated Gibbs measure. The first challenge can in fact be seen as the dual formulation of the maximum entropy problem under constraints and corresponds to the minimization of a convex functional over an open subset of \mathbb{R}^p with $p \in \mathbb{N}$. Therefore it is natural to consider gradient based method in order to find such parameters and this approach was considered in the seminal work of [96] which was the first to consider macrocanonical models in the context of image processing. However, the gradient of this functional is the expectation of the features with respect to the Gibbs measure. In this context, we turn ourselves to the Stochastic Approximation (SA) literature [81, 17]. More precisely, we use the Markov Chain Monte Carlo (MCMC) SA methodology proposed in [26] and referred as Stochastic Optimization with Unadjusted Langevin Algorithm (SOUL) which relies on the Langevin algorithm. This MCMC method has received a lot of attention in the recent years [23, 22, 32] since it exhibits desirable convergence properties and has been extensively used in machine learning applications [93, 86, 73, 66, 1]. Note that a similar methodology to SOUL was already used in a texture synthesis context in [65, 24].

Our second contribution is to establish the convergence of the methodology proposed in [26] in the context of macrocanonical texture synthesis and improve existing results on the dependency with respect to the hyperparameters in this specific case. In particular, the dependency in the dimension is polynomial even in the non-convex setting. This is in accordance with similar results for the convergence of diffusion processes with respect to

the Kantorovitch-Rubinstein distance which are known to be optimal [34].

The paper is organized as follows. Notations and image descriptors are introduced in Section 2. Related work on maximum entropy methods is discussed in Section 3.1. In Section 3.2, we give a mathematical presentation of the microcanonical model and the macrocanonical model. In Section 3.3 we extend results from information geometry to the context of exemplar-based texture synthesis. A Bayesian interpretation and some examples are given in Section 3.4. We then turn to the proposed algorithm for sampling macrocanonical models. The SOUL algorithm is exposed in Section 4.1 and the convergence results applied to our settings are presented in Section 4.2. In Section 4.3 we investigate the links between the microcanonical model and the macrocanonical one. Experiments are presented in Section 5. First we show that the SOUL algorithm numerically solves the problem at hand, considering a toy circular Gaussian case in Section 5.1. After discussing the parameters of the algorithm we turn to the challenging application of texture synthesis in Section 5.2. We study the advantages and the limitations of macrocanonical models and compare our visual results with existing algorithms. We conclude this section by presenting an extension of the studied framework to style transfer. Proofs and additional results are gathered in a supplementary document.

2 Notation and feature models

2.1 Notation

Let $d, p \in \mathbb{N}$. The complement of a set $A \subset \mathbb{R}^d$, is denoted by A^c . For any $A \in \mathbb{R}^d$, we denote $\text{int}(A)$ its interior, $\Delta_A = \{(x, x) \in \mathbb{R}^d \times \mathbb{R}^d : x \in A\}$ the diagonal of A , $\mathcal{B}(\mathbb{R}^d)$ the Borel σ -field of \mathbb{R}^d , $\mathbb{F}(\mathbb{R}^d, \mathbb{R}^p)$ the set of all \mathbb{R}^p -valued Borel measurable functions on \mathbb{R}^d . If $p = 1$, we write $\mathbb{F}(\mathbb{R}^d, \mathbb{R}^p) = \mathbb{F}(\mathbb{R}^d)$ and define for $f \in \mathbb{F}(\mathbb{R}^d)$,

$$\|f\|_\infty = \inf \left\{ t \geq 0 : \text{Leb}(\{x \in \mathbb{R}^d : |f(x)| > t\}) = 0 \right\} ,$$

where Leb is the Lebesgue measure over $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. An open ball of \mathbb{R}^d for the Euclidean distance with center $x_0 \in \mathbb{R}^d$ and radius $r > 0$ is denoted $B(x_0, \mathbb{R}^d)$. For μ a probability measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ and $f \in \mathbb{F}(\mathbb{R}^d, \mathbb{R}^p)$, a μ -integrable function, denote by $\mu(f)$ the integral of f with respect to (w.r.t.) μ , *i.e.*

$$\mu(f) = \int_{\mathbb{R}^d} f(x) d\mu(x) .$$

< If $f = \mathbb{1}_A$ for some measurable set A then we denote $\mu(\mathbb{1}_A) = \mu(A)$. Let $f \in \mathbb{F}(\mathbb{R}^d)$ then for any probability measure μ on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ we denote by $f_\# \mu$ the pushforward measure of μ by f .

Let U be an open set of \mathbb{R}^d . We denote by $C^k(U, \mathbb{R}^p)$ the set of \mathbb{R}^p -valued k -continuously differentiable functions. The differential of $f \in C^k(U, \mathbb{R}^p)$ is denoted df and its Jacobian

matrix Df . Let $C^k(\mathbf{U})$ stand for $C^k(\mathbf{U}, \mathbb{R})$. Let $f : \mathbf{U} \rightarrow \mathbb{R}$, we denote by ∇f , the gradient of f if it exists. f is said to be m -strongly convex with $m \geq 0$ if for all $x, y \in \mathbb{R}^d$ and $t \in [0, 1]$,

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y) - (m/2)t(1-t)\|x-y\|^2 .$$

We recall that if $f : \mathbf{U} \rightarrow \mathbb{R}$ is twice differentiable at point $a \in \mathbb{R}^d$, its Laplacian is given by $\Delta f(a) = \sum_{i=1}^d \partial^2 f(a) / \partial x_i^2$. For any $A \subset \mathbb{R}^d$, we denote by ∂A the boundary of A and $\text{Vol}(A) = \text{Leb}(A)$. For any $\alpha > 0$, let \mathcal{P}_α be the set of probability measures over $\mathcal{B}(\mathbb{R}^d)$ such that $\int_{\mathbb{R}^d} \|x\|^\alpha d\pi(x) < +\infty$. Let $(\Omega, \mathcal{G}, \mathbb{P})$ be a probability space, and $L^2(\Omega, \mathcal{G}) = \{X : X \text{ is a real-valued random variable on } \Omega \text{ such } \mathbb{E}[X^2] < +\infty\}$. Let μ, ν be two probability measures on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. We write $\mu \ll \nu$ if μ is absolutely continuous w.r.t. ν and $d\mu/d\nu$ an associated density. The Kullback-Leibler divergence, or relative entropy, of μ from ν is defined by

$$\text{KL}(\mu|\nu) = \begin{cases} \int_{\mathbb{R}^d} \frac{d\mu}{d\nu}(x) \log\left(\frac{d\mu}{d\nu}(x)\right) d\nu(x) & \text{if } \mu \ll \nu , \\ +\infty & \text{otherwise .} \end{cases}$$

If μ and ν are probability measures, the relative entropy takes values in $[0, +\infty]$. We take the convention that $\prod_{k=p}^n = 1$ and $\sum_{k=p}^n = 0$ for $n, p \in \mathbb{N}$, $n < p$. If $x, y \in \mathbb{C}^d$ with $d \in \mathbb{N}$ we define the periodic convolution between x and y and denote $z = x * y \in \mathbb{C}^d$, the element z such that for any $i \in \{0, \dots, d-1\}$, $z(i) = \sum_{k=0}^{d-1} x(k)y(i-k)$, where x and y are extended over \mathbb{Z} by periodicity. We also denote $\check{x} \in \mathbb{C}^d$ such that for any $i \in \{0, \dots, d-1\}$, $\check{x}(i) = x(-i)$ and x is extended over \mathbb{Z} by periodicity. For any $x \in \mathbb{C}^d$, $\mathcal{F}(x)$ (respectively $\mathcal{F}^{-1}(x) \in \mathbb{C}^d$) stands for the Fourier transform (respectively the inverse Fourier transform), defined for any $j \in \{0, \dots, d-1\}$ by

$$\mathcal{F}(x)(j) = \sum_{k=0}^{d-1} x(k)e^{-2i\pi jk/d} , \quad \mathcal{F}^{-1}(x)(j) = d^{-1} \sum_{k=0}^{d-1} x(k)e^{2i\pi jk/d} .$$

Note that we have $\mathcal{F}^{-1}(\mathcal{F}(x)) = x$. For any $z \in \mathbb{C}$, we denote by $\Re(z)$ the real part of z and by $\Im(z)$ its imaginary part. We denote by $\mathbb{A}_{n_2, n_1}(\mathbb{R})$ the vector space of affine operators from \mathbb{R}^{n_1} to \mathbb{R}^{n_2} and for any $A \in \mathbb{A}_{n_2, n_1}$, \tilde{A} is the linear part of A . Finally, $\mathbb{S}_d(\mathbb{R})$ is the space of $d \times d$ real symmetric matrices.

2.2 Image descriptors

In this work we are interested in sampling probability distributions derived from image models. Let $x_0 \in \mathbb{R}^d$ be an exemplar image and consider a set of constraints associated with some image descriptors $F : \mathbb{R}^d \rightarrow \mathbb{R}^p$. Assume that $F(x_0) = 0$, this can always be achieved upon subtracting $F(x_0)$ to the original features. The constraints on the target distribution are then given by $F = 0$ almost surely or in expectation.

In this section, we review some of the popular possibilities for the choice of the function F . In the literature, many different approaches such as Gaussian models, *i.e.* mean and correlation features [91, 40], wavelet-based descriptors [49, 77, 75, 30] or convolutional neural network features (CNN) [44, 90, 54] have been proposed to come up with visually satisfying image descriptors.

In our study we will focus on two sets of features: (i) Gaussian features ; (ii) CNN features. Gaussian features have the mathematical advantage of defining a strongly convex model, therefore allowing for strong convergence results to apply. However Gaussian textures do not exhibit sharp edges and lack long-range structures and as a consequence richer models should be investigated in order to obtain visually satisfying images. Similarly to [44, 90, 65] we consider features derived from a pretrained CNN. It has been observed that these features are efficient for describing a large variety of natural images. However, these improvements over the Gaussian model come at a high computational price. First, the features we end up with are no longer convex. Second, the dimension of the associated parameter space is usually high. An experimental investigation of the behavior of our proposed algorithm for these two sets of features is conducted in Section 5. We now describe precisely these two models.

Gaussian features Let $x_0 \in \mathbb{R}^d$ and consider $F(x) = x * \check{x} - x_0 * \check{x}_0$. In the Fourier domain, we have for any $i \in \{1, \dots, d\}$, $\mathcal{F}(F(x))(i) = |\mathcal{F}(x)(i)|^2 - |\mathcal{F}(x_0)(i)|^2$. Therefore, if $F(x) = 0$, x has same power spectrum, *i.e.* same autocorrelation, as x_0 , namely $\mathcal{F}(x)$ has the same modulus as $\mathcal{F}(x_0)$. However the equation $F(x) = 0$ gives no information on the phases of $\mathcal{F}(x)$.

Neural network features A Neural network is a series of affine operations (usually convolutions) followed at each step by a pointwise non-linearity. We define

$$(A_j^k)_{j \in \{1, \dots, M\}, k \in \{1, \dots, c_j\}} \in \prod_{j=1}^M \mathbb{A}_{n_j, c_{j-1} \times n_{j-1}}(\mathbb{R})^{c_j}, \quad (n_j, c_j)_{j \in \{0, \dots, M\}} \in \mathbb{N}^{M+1} \times \mathbb{N}^{M+1}, \quad (1)$$

with $M \in \mathbb{N}$, $n_0 = d$ and $c_0 = 1$. For each $j \in \{0, \dots, M\}$, n_j is the *dimension* of the j -th layer and c_j is the *number of channels* of the j -th layer, and for any $k \in \{1, \dots, c_j\}$, $A_j^k \in \mathbb{A}_{n_j, c_{j-1} \times n_{j-1}}$. Namely for any layer $j \in \{1, \dots, M\}$ and channel $k \in \{1, \dots, c_j\}$, $A_j^k : \mathbb{R}^{c_{j-1} \times n_{j-1}} \rightarrow \mathbb{R}^{n_j}$ is the affine operator which maps the $(j-1)$ -th layer to the j -th layer and channel k before the non-linear operation. With our notations, the 0-th layer corresponds to the original image. We recall that for any $j \in \{1, \dots, M\}$ and $k \in \{1, \dots, c_j\}$, \tilde{A}_j^k denotes the linear part of A_j^k . We also define for any $j \in \{1, \dots, M\}$ and $x \in \mathbb{R}^{c_{j-1} \times n_{j-1}}$, $A_j(x) = (A_j^k(x))_{k \in \{1, \dots, c_j\}}$, *i.e.* $A_j \in \mathbb{A}_{c_j \times n_j, c_{j-1} \times n_{j-1}}$ is the affine operator which maps the $j-1$ -th layer to the j -th layer before the non-linear operation.

Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function. By a slight abuse of notation we denote for

any $d \in \mathbb{N}$ and $x \in \mathbb{R}^d$, $\varphi(x) = (\varphi(x(0)), \dots, \varphi(x(d-1)))$. We assume that φ satisfies the following conditions:

- (a) for any $d \in \mathbb{N}$, there exists $C_{d,\varphi} \geq 0$ such that for any $x \in \mathbb{R}^d$, $\|\varphi(x)\| \leq C_{d,\varphi}(1+\|x\|)$,
- (b) φ is non-decreasing,
- (c) $\lim_{t \rightarrow +\infty} \varphi(t) = +\infty$.

We define for any $j \in \{1, \dots, M\}$ and $k \in \{1, \dots, c_j\}$, the (j, k) -th *layer-channel feature* $\mathcal{G}_j^k : \mathbb{R}^d \rightarrow \mathbb{R}^{n_j}$, for any $x \in \mathbb{R}^d$, by

$$\mathcal{G}_j^k(x) = \left(\varphi \circ A_j^k \circ \varphi \circ A_{j-1} \circ \dots \circ \varphi \circ A_1 \right) (x), \quad \mathcal{G}_0(x) = x.$$

For any layer $j \in \{1, \dots, M\}$ and channel $k \in \{1, \dots, c_j\}$, $\mathcal{G}_j^k(x)$ is the neural network response of x at layer j and channel k . We also define for any $j \in \{1, \dots, M\}$, the j -th *layer feature* $\mathcal{G}_j : \mathbb{R}^d \rightarrow \mathbb{R}^{c_j \times n_j}$, for any $x \in \mathbb{R}^d$, by $\mathcal{G}_j(x) = (\mathcal{G}_j^k(x))_{k \in \{1, \dots, c_j\}}$. Let $\mathbf{j} \subset \{1, \dots, M\}$ then we can define $F(x) \in \mathbb{R}^p$ for any $x \in \mathbb{R}^d$ by

$$F(x) = \left(\overline{\mathcal{G}}_j^k(x) - \overline{\mathcal{G}}_j^k(x_0) \right)_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}}, \quad \overline{\mathcal{G}}_j^k(x) = n_j^{-1} \sum_{\ell=1}^{n_j} \mathcal{G}_j^k(x)(\ell), \quad p = \sum_{j \in \mathbf{j}} c_j. \quad (2)$$

A few remarks are in order regarding the dimension of the associated parameter space. In our applications, we will use the **VGG19** convolutional neural network [85], see Appendix D for details on the structure of **VGG19** network. Note that since **VGG19** is a convolutional neural network, *i.e.* the linear part of the affine operators is given by a convolutional operator, and since we average the neural network response, the output dimension p is independent of the input dimension d . Selecting the layers $\mathbf{j} = \{1, 3, 6, 8, 11, 13, 15, 24, 26, 31\}$ we have $p = 2688 \approx 10^3$. Usually we will consider images of size at least 512×512 for which $d = 262144 \approx 10^5$. Therefore the features described by (2) performs a dimension reduction. In [44] similar image descriptors are considered but Gram matrices are used instead of $(\overline{\mathcal{G}}_j)_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}}$. This leads to a parameter space with dimension $352256 \approx 10^5$, see [79].

3 Statistical texture models

3.1 Previous work

As emphasized in Section 1, there have been two main approaches to address the exemplar-based texture synthesis problem. First, non-parametric methods produce an output image following a statistical process, e.g. a Markov random field [36, 16, 78]. These methods do not require an explicit texture model. Most of these algorithms are based on patch

information, see the review paper [79]. Indeed, in order to update the current image, the patches of the input texture are rearranged in order to generate a new element (a pixel or a block of pixels) which is locally coherent with the pre-existing structure. The seminal work of [36] paved the way for the use of such methods and was later extended in [35, 61] to handle blockwise updates instead of pixelwise ones. The statistical model of [36] was analyzed in [62] in which the authors reformulate the original algorithm as a bootstrap scheme. Since these methods duplicate some part of the input image in order to sample the new image, their innovation capability is limited as they might suffer from a copy-paste problem. Some recent patch-based methods deal with these issues by introducing randomness either in the update [80]. In [43], starting from a random microtexture initial image, the patches are rearranged using optimal transport. [60] reformulates a patch-based synthesis algorithm as an optimization procedure, therefore yielding a global texture model defined by the patch information. This model was later extended in [47, 58].

For the seconde type of approaches, *i.e.* parametrics ones, in the early work of [38], textures were described as fractional Brownian motions. It was later noticed in [91] that a large class of textures could be generated using spot noise models whose normalized limit for a large number of spots is a Gaussian random field with a circulant covariance matrix [40]. In this works the underlying image model is Gaussian and the pixel distribution of the output image has, in expectation, the same moment of order 1 and 2 as the input texture. It was shown that these algorithms are able to sample a large class of textures. Different features are designed in [74]. All the images in this class share the property that they do not exhibit salient spatial structures implying that the knowledge of the second-order moments was not enough to reproduce natural images. In [39, 15, 49] the authors remark that structured textures could be obtained using hand-selected features. These ideas are extended in [77] by designing a bank of wavelet features. In addition the authors rewrite parametric exemplar-based texture synthesis algorithms as maximum entropy problems under constraints. Similarly, in the seminal paper [96], the authors derive the first texture synthesis methodology based on a maximum entropy approach with constraints in expectation. Later, replacing wavelet features in [77] by convolutional neural network features, [44] was able to obtain striking visual results. Combining these neural network features with pixel-based constraints yields improvements of the original work of [44]. For instance, in [64] spectral constraints are added to the convolutional neural network ones. [55, 88, 89] design parametric methods relying on convolutional neural network features. The sampling procedure does not depend on any gradient-based method but instead is performed in a feed-forward manner. More recently, many papers investigate the use of Generative Adversarial Networks (GAN) [46] in texture synthesis giving promising results [54, 9, 63, 95]. In GAN [54, 46] the structure constraint is encoded in the loss on the generator, *i.e.* the samples must look like the input image. The innovation constraint is encoded in the loss on the discriminator, *i.e.* the samples must be diverse enough for the discrimination task to be hard. This formulation can be rewritten as an optimization problem on probability measures [3, 10]. In this case, the target distribution

on natural textures is the solution of the training problem. Note that in this case, though the synthesis is performed in a feed-forward manner, the neural network must be retrained when presented a new class of textures. In the following approaches, the generation is not feed-forward but the natural texture distribution is designed so that the same features are generically used for all textures.

3.2 Maximum entropy probability measures

We now define microcanonical and macrocanonical models as introduced by [14]. Let μ be a probability measure over $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ with $d \in \mathbb{N}$. The measure μ will be referred to as the *reference probability measure*. Let $F : \mathbb{R}^d \rightarrow \mathbb{R}^p$ with $p \in \mathbb{N}$ be a measurable mapping. This mapping will be referred to as the *statistical constraints* of the model. A discussion on the choice of F was conducted in Section 2.2. From now on we assume that we observe an input texture x_0 such that $F(x_0) = 0$. Once again, note that this can always be achieved upon subtracting $F(x_0)$ to the original features. Given a set of features and a target image, we are now interested in the probability distributions which have maximum entropy (innovation constraint) and such that the features are equal to the ones of the target image (structure constraint) almost surely. In order for the model to be well-defined we replace the maximization of the entropy by a minimization of the Kullback-Leibler divergence $\text{KL}(\cdot|\mu)$ where μ is the reference probability measure. This methodology is called the *microcanonical model*, see [14].

Definition 1. *A probability measure π is a microcanonical model with respect to the constraints F and the reference measure μ , if $\pi(\{F = 0\}) = 1$ and if for any probability distribution ν which satisfies the previous assumption we have $\text{KL}(\pi|\mu) \leq \text{KL}(\nu|\mu)$.*

This model was already considered in [39, 15]. Steerable pyramids constraints are used in [49]. In [77] more than 700 cross-correlations, autocorrelations and statistical moments of wavelet coefficients are selected in order to design the statistical constraint features. With the rise of convolutional neural networks, striking visual results were obtained by [44]. This line of work has then been extended by several authors, adding other constraints such as spectrum projection [64]. However the microcanonical model distribution is untractable for most statistical constraints. Therefore in order to sample from this distribution the following heuristics is used: 1) sample a white noise image; 2) perform a gradient descent on the square norm of the features. We shall see in Section 4.3 that this algorithm, while providing satisfying visual results for neural network constraints, does not sample from the microcanonical distribution. We now consider a relaxation of the *microcanonical model*: the *macrocanonical model*, see [14]. Given a set of features and a target image, we are now interested in the probability distributions which have maximum entropy (innovation constraint) and such that expectation of the features is equal to the features of the target image (structure constraint).

Definition 2. A probability measure π is a macrocanonical model with respect to the constraints F and the reference measure μ , if F is π -integrable, $\pi(F) := \int_{\mathbb{R}^d} F(x) d\pi(x) = 0$ and if for any probability distribution ν which satisfies the previous assumption we have $\text{KL}(\pi|\mu) \leq \text{KL}(\nu|\mu)$.

This maximum entropy model was introduced by [53] and used in texture synthesis by [96]. Under some assumptions the macrocanonical model can be written in a closed-form. In [96], the authors used a dictionary of quantiles for various filters, linear and non-linear, as statistical constraint features. This work has been extended, using convolutional neural network features in [65]. In [14] conditions under which the macrocanonical and the microcanonical models coincide when the dimension of the image space grows towards infinity are identified.

In the next section we study the existence and uniqueness of the macrocanonical model introducing a dual problem and using the point of view of information geometry.

3.3 Existence, uniqueness and dual formulation

Considering the model defined by Definition 2 it is natural to ask the following questions:

- (a) When does a macrocanonical model exist? Can we identify explicit conditions for its existence?
- (b) If such a model exists, is it unique?
- (c) Can we find closed forms for the probability distribution functions of macrocanonical models?

We will answer positively to (b). In the case where the problem is non degenerate, *i.e.* there exists a probability measure ν such that $\text{KL}(\nu|\mu) < +\infty$ and $\nu(F) = 0$, then the macrocanonical model exists and is given by a parametric measure, answering both (a) and (c). However, checking that the problem is indeed non degenerate can be as hard as finding the macrocanonical model.

To show the existence of a macrocanonical model we give a dual, convex and finite dimensional formulation. This problem is then solved under the following conditions on F and μ . Let $\alpha > 0$ and $\beta > 0$:

A1 (α). F is continuous and there exists $C_\alpha \geq 0$ with $\sup_{x \in \mathbb{R}^d} \{ \|F(x)\| (1 + \|x\|^\alpha)^{-1} \} \leq C_\alpha < +\infty$.

A2 (β). There exists $\eta > 0$, such that $\int_{\mathbb{R}^d} \exp[\eta \|x\|^\beta] d\mu(x) < +\infty$.

Let \mathcal{P}_α be the set of probability measures over $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ such that $\int_{\mathbb{R}^d} \|x\|^\alpha d\pi(x) < +\infty$. We define the set of *admissible probability measures* by $\mathcal{P}_\alpha^F = \{ \pi \in \mathcal{P}_\alpha : \pi(F) = 0 \}$ and consider the following problem:

$$\text{minimize } \text{KL}(\pi|\mu) \text{ subject to } \pi \in \mathcal{P}_\alpha^F. \tag{P}$$

We denote $v(\mathbf{P}) = \inf_{\mathcal{P}_\alpha^F} \text{KL}(\pi|\mu)$. It is clear that any solution of (P) is a macrocanonical model with respect to the constraints F and the reference measure μ . First, we assert that if the solution of (P) exists and is non-degenerate, *i.e.* $v(\mathbf{P}) < +\infty$, then it is unique. Let π_1^* and π_2^* be two solutions of (P) with $v(\mathbf{P}) < +\infty$ and $\phi(t) = t \log(t)$, defined on $[0, +\infty)$ with the convention that $\phi(0) = 0$. Since $v(\mathbf{P}) < +\infty$ we have that $\pi_1^* \ll \mu$ and $\pi_2^* \ll \mu$. Since \mathcal{P}_α^F is convex, π^* defined for any $x \in \mathbb{R}^d$ by $\frac{d\pi^*}{d\mu}(x) = \left(\frac{d\pi_1^*}{d\mu}(x) + \frac{d\pi_2^*}{d\mu}(x) \right) / 2$ belongs to \mathcal{P}_α^F . Using that ϕ is strictly convex we have

$$2v(\mathbf{P}) = \text{KL}(\pi_1^*|\mu) + \text{KL}(\pi_2^*|\mu) = \int_{\mathbb{R}^d} \left\{ \phi\left(\frac{d\pi_1^*}{d\mu}(x)\right) + \phi\left(\frac{d\pi_2^*}{d\mu}(x)\right) \right\} d\mu(x) \geq 2\text{KL}(\pi^*|\mu) ,$$

with equality if and only if for μ almost every $x \in \mathbb{R}^d$ we have $\frac{d\pi_1^*}{d\mu}(x) = \frac{d\pi_2^*}{d\mu}(x)$. As a consequence, since $\text{KL}(\pi^*|\mu) = v(\mathbf{P})$, $\pi_1^* = \pi_2^*$.

We now introduce another optimization problem.

$$\text{maximize } \inf_{\pi \in \mathcal{P}_\alpha} \{ \text{KL}(\pi|\mu) + \langle \theta, \pi(F) \rangle \} \text{ subject to } \theta \in \Theta_F , \quad (\text{Q})$$

with

$$\Theta_F = \left\{ \theta \in \mathbb{R}^p : \int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) < +\infty \right\} . \quad (3)$$

Using Hölder's inequality, one can show that Θ_F is convex. In addition, if **A1**(α) and **A2**(α) hold with $\alpha > 0$ then $\bar{\mathbf{B}}(0, \eta/C_\alpha) \subset \text{int}(\Theta_F)$. Similarly to (P), we denote $v(\mathbf{Q}) = \sup_{\Theta_F} \inf_{\pi \in \mathcal{P}_\alpha} \{ \text{KL}(\pi|\mu) + \langle \theta, \pi(F) \rangle \}$. Introducing the Lagrangian $\mathcal{L}(\pi, \theta) = \text{KL}(\pi|\mu) + \langle \theta, \pi(F) \rangle$ defined over $\mathcal{P}_\alpha \times \Theta_F$, we have

$$v(\mathbf{Q}) = \sup_{\Theta_F} \inf_{\mathcal{P}_\alpha} \mathcal{L} \leq \inf_{\mathcal{P}_\alpha} \sup_{\Theta_F} \mathcal{L} \leq v(\mathbf{P}) . \quad (4)$$

We denote $d(\mathbf{P}, \mathbf{Q})$ the *duality gap* $d(\mathbf{P}, \mathbf{Q}) = v(\mathbf{P}) - v(\mathbf{Q})$ with the convention that $\infty - \infty = 0$. Let the log-partition function $L : \Theta_F \rightarrow \mathbb{R}$ be given for any $\theta \in \Theta_F$ by

$$L(\theta) = \log \left[\int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) \right] . \quad (5)$$

We also define for any $\theta \in \Theta_F$, the probability measure π_θ whose density with respect to μ is given for any $x \in \mathbb{R}^d$ by

$$\frac{d\pi_\theta}{d\mu}(x) = \exp[-\langle \theta, F(x) \rangle - L(\theta)] . \quad (6)$$

Using Proposition 12, (Q) is equivalent to

$$\text{minimize } L(\theta) \text{ subject to } \theta \in \Theta_F . \quad (\text{Q}')$$

More precisely θ^* is a solution of (Q) if and only if it is a solution of (Q'). Assuming that \mathbb{R}^d is replaced by \mathbf{X} , a finite set, in (P), that μ is the counting measure on \mathbf{X} and identifying the probability distributions over \mathbf{X} and the $|\mathbf{X}|$ -dimensional simplex, we obtain that the primal problem (P) can be rewritten as a finite dimensional convex optimization problem under linear constraints. In this case $\Theta_F = \mathbb{R}^p$ and we have that (Q) is the Lagrangian dual problem of (P), where θ are the Lagrange multipliers. In this setting and under some identifiability conditions, the solution of (P) is given by some Gibbs measure, *i.e.* for any $x \in \mathbf{X}$, $\pi_{\theta^*}(x) = \exp[-\langle \theta^*, F(x) \rangle - L(\theta^*)]$ with $\theta^* \in \mathbb{R}^p$, see [69]. In what follows we investigate how the results obtained in the discrete setting can be extended to the general case. The next proposition is an extension of [18, Theorem 3.1] in the case where F is not bounded, which characterizes the solutions of (P). Under **A2**(α') with $\alpha' > \alpha$ the existence of a solution of (P) is ensured as soon as the set of admissible probability measures is not empty.

Proposition 3. *Assume **A1**(α) with $\alpha \geq 0$. The following holds.*

- (a) *If there exists a solution of (P) such that $v(\text{P}) < +\infty$ then there exists $\theta^* \in \mathbb{R}^p$ such that the solution of (P) is given by π_{θ^*} defined for μ almost any $x \in \mathbb{R}^d$ by*

$$\begin{cases} \frac{d\pi_{\theta^*}}{d\mu}(x) = \exp[-\langle \theta^*, F(x) \rangle] / \int_{\mathbb{R}^d \setminus \mathbf{N}} \exp[-\langle \theta^*, F(y) \rangle] d\mu(y) & \text{if } x \notin \mathbf{N}, \\ \frac{d\pi_{\theta^*}}{d\mu}(x) = 0 & \text{otherwise,} \end{cases} \quad (7)$$

with $\mathbf{N} \in \mathcal{B}(\mathbb{R}^d)$ such that for all $\pi \in \mathcal{P}_\alpha^F$ with $\text{KL}(\pi|\mu) < +\infty$, $\pi(\mathbf{N}) = 0$. If there exists $\pi \in \mathcal{P}_\alpha^F$ with $\text{KL}(\pi|\mu) < +\infty$ such that $\mu \ll \pi$ then θ^ is a solution of (Q), π_{θ^*} given by (6) is a solution of (P) and $v(\text{Q}) = v(\text{P})$, where $v(\text{Q})$ and $v(\text{P})$ are given in (4).*

- (b) *Assume **A2**(α') with $\alpha' > \alpha$. There exists π_{θ^*} solution of (P) with $v(\text{P}) < +\infty$ if and only if there exists $\pi \in \mathcal{P}_\alpha^F$ such that $\text{KL}(\pi|\mu) < +\infty$.*

Proof. The proof is postponed to Appendix A.1. □

Note that this result was extended to a large class of divergences, see [2, Theorem 3.8]. In [18, 19, 20] different sufficient conditions for solving (P) are derived. In particular, in [18, Theorem 3.3], the existence of a solution to (P) is shown if F is measurable, **A2**(α) holds and there exists $\mathbf{U} \subset \mathbb{R}^p$ open with $0 \in \mathbf{U}$ such that for any $a \in \mathbf{U}$ there exists ν_a with $\text{KL}(\nu_a|\mu) < +\infty$ and $\nu_a(F) = a$. In [87, Theorem 3], the authors show similar results in the case where the Kullback-Leibler divergence is given by the true entropy. In practice it is difficult to check the conditions of Proposition 3. Indeed, even if **A2**(α') holds with $\alpha' > \alpha$ it is not trivial to find an element $\pi \in \mathcal{P}_\alpha^F$ such that $\text{KL}(\pi|\mu) < +\infty$. We now turn to the dual formulation (Q) which is easier to deal with since it is a finite dimensional and convex optimization problem. Under **A2**(α), any stationary point of the log-partition function yields a solution of the primal formulation.

Proposition 4. Assume **A1**(α) and **A2**(α) with $\alpha \geq 0$. Then, $L \in C^\infty(\text{int}(\Theta_F))$, where L is given in (5), and for any $\theta \in \text{int}(\Theta_F)$, $\nabla L(\theta) = \pi_\theta(F)$ with π_θ given by (6). In addition, if there exists $\theta^* \in \text{int}(\Theta_F)$ such that $\nabla L(\theta^*) = 0$, then π_{θ^*} is the solution of (P).

Proof. The proof is postponed to Appendix A.2. □

A similar result was derived in [56], in the case where μ is no longer a probability measure but only sigma-finite. The study of the log-partition function L is fundamental in information theory. Its main properties are summarized in [7, Chapter 9] and in [21] the authors show that the log-partition is a special case of a more general information theoretical model, where the Kullback-Leibler divergence is replaced by another notion of entropy. In fact, the log-partition function is a convex function, hence any stationary point is a global minimizer. We exploit this fact in the following proposition.

Proposition 5. Assume **A1**(α), **A2**(α') with $\alpha \geq 0$, $\alpha' > \alpha$.

(a) If for any $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$, we have

$$\mu\left(\left\{x \in \mathbb{R}^d : \langle F(x), \theta \rangle < 0\right\}\right) > 0, \quad (8)$$

then θ^* , solution of (Q), exists and π_{θ^*} given by (6) is the solution of (P).

(b) In particular, (8) is satisfied if $\mu(A) > 0$ for any non-empty open set $A \subset \mathbb{R}^d$, F is continuous and there exists $x \in F^{-1}(\{0\})$ such that F is differentiable at x and we have $\det(DF(x)DF(x)^T) > 0$.

Proof. The proof is postponed to Appendix A.3. □

In Section 3.4, we apply Proposition 5 for some example of functions F of the form (2). However, note that it does not apply to the case where φ is the rectified linear unit (RELU), *i.e.* for any $t \in \mathbb{R}$, $\varphi(t) = \max(0, t)$. Still we give a similar result, valid for RELU, in Proposition 7. In [51, Theorem 3.5] the authors derive analogous results in the case where μ is the Lebesgue measure, *i.e.* in the case where the Kullback-Leibler is replaced by the true entropy.

We are also able to show that the condition (8) is almost necessary in the next result.

Proposition 6. Assume **A1**(α), **A2**(α') with $\alpha \geq 0$, $\alpha' > \alpha$ and that there exists $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$ such that

$$\mu\left(\left\{x \in \mathbb{R}^d : \langle F(x), \theta \rangle \leq 0\right\}\right) = 0.$$

Then δ_{x_0} solves (P) and $v(P) = +\infty$.

Proof. The proof is postponed to Appendix A.4. □

We have seen that, under assumptions on the reference distribution μ and the statistical constraints F , the macrocanonical model is the distribution π_{θ^*} with the following parametric form: for any $x \in \mathbb{R}^d$, $(d\pi_{\theta^*}/d\mu)(x) = \exp[-\langle \theta^*, F(x) \rangle - L(\theta^*)]$, with θ^* which satisfies $\theta^* \in \operatorname{argmin}_{\theta \in \Theta_F} L(\theta)$. These exponential families can also be retrieved in the following Bayesian framework. Assume that the likelihood of texture images associated with parameters $\theta \in \mathbb{R}^p$ is given for any $x \in \mathbb{R}^d$ by

$$p(x|\theta) = \exp[-\langle \theta, F(x) \rangle] \Big/ \int_{\mathbb{R}^d} \exp[-\langle \theta, F(y) \rangle] d\mu(y) .$$

Assume that x_0 is a sample from this distribution and that $F(x_0) = 0$. Using Bayes' formula we obtain that for any $\theta \in \mathbb{R}^p$

$$p(\theta|x_0) = \frac{p(x_0|\theta)p(\theta)}{p(x_0)} \propto p(\theta) \Big/ \int_{\mathbb{R}^d} \exp[-\langle \theta, F(y) \rangle] d\mu(y) .$$

In order to compute the maximum a posteriori estimator θ_{MAP} we need to set a prior distribution $p(\theta)$. Choosing the non-informative improper prior $p(\theta) = \mathbb{1}_{\Theta_F}(\theta)$ we get that

$$\theta_{\text{MAP}} \in \operatorname{argmin}_{\theta \in \Theta_F} \log \left[\int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) \right] ,$$

which corresponds to the dual formulation (Q). However, other prior distributions could be considered, yielding hierarchical Bayesian models [92].

3.4 Illustrative examples

Gaussian features We consider $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$, *i.e.* $d = p$, given for any $x \in \mathbb{R}^d$ by $F(x) = x * \check{x} - x_0 * \check{x}_0$ and μ a Gaussian probability measure with zero mean and diagonal covariance matrix with diagonal coefficient σ^2 with $\sigma > 0$. Assume in addition that for any $\ell \in \{0, \dots, d-1\}$, $\mathcal{F}(x_0)(\ell) \neq 0$. We have that, **A1**(α) and **A2**(α) hold with $\alpha = 2$ and $\eta \in (0, (2\sigma^2)^{-1})$. Using the Fourier-Plancherel formula we get that for any $x \in \mathbb{R}^d$ and $\theta = (\theta(0), \dots, \theta(d-1)) \in \mathbb{R}^d$,

$$\begin{aligned} \langle \theta, F(x) \rangle + \|x\|^2 / (2\sigma^2) &= d^{-1} \left[\langle \mathcal{F}(\theta), |\mathcal{F}(x)|^2 - |\mathcal{F}(x_0)|^2 \rangle + \|\mathcal{F}(x)\|^2 / (2\sigma^2) \right] \\ &= \sum_{i=0}^{d-1} \left\{ d^{-1} (\mathcal{F}(\theta)(i) + (2\sigma^2)^{-1}) |\mathcal{F}(x)(i)|^2 \right\} - d^{-1} \langle \mathcal{F}(\theta), |\mathcal{F}(x_0)|^2 \rangle . \end{aligned}$$

This implies that, $\Theta_F = \mathcal{F}^{-1}[\mathfrak{R}^{-1}\{(-2\sigma^2)^{-1}, +\infty\}^d] \cap \mathbb{R}^d$. In addition, for any $\theta \in \Theta_F$ with $\mathcal{F}(\theta) \in \mathbb{R}^d$ and X distributed according to π_θ , we obtain that $\mathcal{F}(X)$ is a d -dimensional complex Gaussian random variable on $\mathcal{F}(\mathbb{R}^d)$ with zero mean and diagonal covariance matrix with diagonal coefficients given by $d(\mathcal{F}(\theta) + (2\sigma^2)^{-1})^{-1}/2$. Similarly, we obtain that X distributed according to π_θ is a d -dimensional Gaussian random variable with zero

mean and invertible circulant covariance matrix $\mathbf{C}_\theta \in \mathbb{S}_d(\mathbb{R})$ whose inverse is given for any $i, j \in \{0, \dots, d-1\}$ by

$$\mathbf{C}_\theta^{-1}(i, j) = 2\theta(i - j) + \sigma^{-2} . \quad (9)$$

Note that in this case, since $\mathcal{F}(\theta) \in \mathbb{R}^d$, $\theta = \check{\theta}$. Let

$$\mathcal{F}(\theta^*) = (d|\mathcal{F}(x_0)|^{-2} - \sigma^{-2})/2 \in \mathcal{F}(\Theta_F) . \quad (10)$$

In this case, for any X distributed according to π_{θ^*} we obtain that $\mathcal{F}(X)$ is a d -dimensional Gaussian random variable with zero mean and diagonal covariance matrix with diagonal coefficients given by $|\mathcal{F}(x_0)|^2$. Therefore we have

$$\mathbb{E} [X * \check{X}] = \mathcal{F}^{-1} \left(\mathbb{E} [|\mathcal{F}(X)|^2] \right) = \mathcal{F}^{-1} (|\mathcal{F}(x_0)|^2) = x_0 * \check{x}_0 ,$$

which implies that $\pi_{\theta^*}(F) = 0$ and that $\theta^* = \mathcal{F}^{-1}(d|\mathcal{F}(x_0)|^{-2} - \sigma^{-2})/2$ is a solution of (Q), since $\theta^* \in \Theta_F$. Using Proposition 4 we get that π_{θ^*} is a solution of (P). Therefore the solution of (P) is the Gaussian probability measure with zero mean which satisfies for any $m, n \in \{0, \dots, d-1\}$, $\mathbb{E}[X(m)X(n)] = d^{-1}(x_0 * \check{x}_0)(m - n)$. This distribution is invariant by spatial translation, *i.e.* denoting $\tau : \mathbb{R}^d \rightarrow \mathbb{R}^d$, defined for any $x \in \mathbb{R}^d$ and $i \in \{1, \dots, d-1\}$ by $\tau(x)(i) = x(i+1)$ and $\tau(x)(d-1) = x(0)$, we have for any $\mathbf{A} \in \mathcal{B}(\mathbb{R}^d)$

$$\pi^*(\mathbf{A}) = \pi^*(\tau(\mathbf{A})) . \quad (11)$$

Note that the distribution of the random variable X is the same as the one of $d^{-1/2}(x_0 * Z)$ with Z a d -dimensional Gaussian random variable with zero mean and covariance matrix identity, see [40].

CNN Features We now turn to the case where F is given by (2) for $\mathbf{j} \subset \{1, \dots, M\}$ and the reference measure μ is a Gaussian probability measure with zero mean and symmetric positive covariance matrix. **A1**(1) and **A2**(α') hold for any $\alpha' \in [0, 2)$. Note that in the case, where φ is differentiable, the results of Proposition 4 hold assuming that there exists a point $x \in \mathbb{R}^d$ such that $F(x) = F(x_0)$ and $dF(x)$ is surjective. In the case where for all $t \in \mathbb{R}$, $\varphi(t) = \max(0, t)$ we can define a certificate ensuring the existence of a macrocanonical model. We now introduce some preliminary notations. Let $(A_j)_{j \in \{1, \dots, M\}}$ be given by (1) and $(\mathcal{G}_j)_{j \in \{1, \dots, M\}}$ be given by (2). For any $j \in \{1, \dots, M\}$, let $A_{j,+} \in \mathbb{A}_{c_j \times n_j, c_{j-1} \times n_{j-1}}(\mathbb{R})$ defined by

$$A_{j,+} = D_j(x_0)A_j , \quad (12)$$

with $D_j(x_0) \in \mathbb{R}^{c_j \times n_j} \times \mathbb{R}^{c_j \times n_j}$ a diagonal matrix such that for any $i \in \{1, \dots, c_j \times n_j\}$, $D_j(x_0)(i, i) = \mathbb{1}_{(0, +\infty)}(\mathcal{G}_j(x_0)(i))$. Note that we have for any $j \in \{1, \dots, M\}$, $\tilde{A}_{j,+} = D_j(x_0)\tilde{A}_j$. Let $(\tilde{v}_{j,k})_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}}$ such that for any $j, j' \in \mathbf{j}$ and $k, k' \in \{1, \dots, c_j\}$, $\tilde{v}_{j,k} \in \mathbb{R}^{c_j \times n_j}$ and

$$\begin{aligned} \tilde{v}_{j,c}(j', c') &= 0 , \text{ if } c' \neq c , \\ \tilde{v}_{j,c}(j', c') &= n_j^{-1} \text{ otherwise .} \end{aligned}$$

Note that for any $x \in \mathbb{R}^d$, $F(x) = (\tilde{v}_{j,k}^T \mathcal{G}_j(x))_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}}$. In addition, define for any $j \in \mathbf{j}$ and $k \in \{1, \dots, c_j\}$

$$v_{j,k} = \tilde{A}_{1,+}^T \dots \tilde{A}_{j,+}^T \tilde{v}_{j,k} .$$

Note that for any $j \in \mathbf{j}$ and $k \in \{1, \dots, c_j\}$, $v_{j,k} \in \mathbb{R}^d$. The next proposition highlights the role of $(v_{j,k})_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}}$ as a certificate for the existence of the macrocanonical model.

Proposition 7. *Assume that **A2**(α) holds with $\alpha > 1$, that $\mu(\mathbf{A}) > 0$ for every open set $\mathbf{A} \subset \mathbb{R}^d$ with $\mathbf{A} \neq \emptyset$ and that F is given by (2) with $\varphi(t) = \max(0, t)$ for any $t \in \mathbb{R}$, $x_0 \in \mathbb{R}^d$ and $\mathbf{j} \subset \{1, \dots, M\}$. Moreover assume that for any $j \in \{1, \dots, M\}$ and $k \in \{1, \dots, c_j \times n_j\}$*

$$e_{k,j}^T A_j \mathcal{G}_{j-1}(x_0) \neq 0 , \quad (13)$$

where for any $j \in \{1, \dots, M\}$, $(e_{k,j})_{k \in \{1, \dots, c_j \times n_j\}}$ is the canonical basis of $\mathbb{R}^{c_j \times n_j}$. In addition, assume that the family $(v_{j,c})_{j \in \mathbf{j}, c \in \{1, \dots, c_j\}}$ is linearly independent. Then there exists a solution θ^* to (Q) and π_{θ^*} defined given by (6) is the solution of (P).

Proof. The proof is postponed to Appendix A.5. □

Note that since $(v_{j,c})_{j \in \mathbf{j}, c \in \{1, \dots, c_j\}}$ has a closed form, the independence condition in Proposition 7 can be explicitly checked given a trained neural network. The proof of Proposition 7 exploits the fact that neural networks are locally linear under mild assumptions. Finally, since the family $(v_{j,c})_{j \in \mathbf{j}, c \in \{1, \dots, c_j\}}$ has cardinality $p = \sum_{j \in \mathbf{j}} c_j$, Proposition 7 never applies when $p > d$.

Finally, note that here we propose to study Gaussian features and CNN features but other choices are possible such as wavelet features or scattering transform features [67, 13].

4 Sampling from macrocanonical models

In this section, our objective is twofold. First, we want to find a sequence $(\theta_n)_{n \in \mathbb{N}}$ which converges almost surely to θ^* , the solution of (Q). Second, we aim at sampling from the macrocanonical model π_{θ^*} defined by (6). We present a Stochastic Approximation (SA) algorithm addressing simultaneously these two problems in Section 4.1. Our main result are summarized in Section 4.2. In Section 4.3 we draw a qualitative link between macrocanonical and microcanonical models.

4.1 The Stochastic Optimization via Unadjusted Langevin method

In this section, we introduce a SA algorithm to minimize L . In Section 4.1.1, we recall some basic facts on convex stochastic optimization and in Section 4.1.2 we present the Stochastic Optimization with Unadjusted Langevin (SOUL) Algorithm applied to the maximum entropy problem.

4.1.1 Stochastic approximation

Let $K \subset \text{int}(\Theta_F)$ be a non-empty compact convex set such that $K \cap \text{argmin}_{\Theta_F} L \neq \emptyset$ with L the log-partition given in (5). Since $\theta \mapsto L(\theta)$ is a convex mapping we obtain that the sequence $(\tilde{\theta}_n)_{n \in \mathbb{N}}$ defined by $\tilde{\theta}_0 \in K$ and for any $n \in \mathbb{N}$, $\tilde{\theta}_{n+1} = \Pi_K[\tilde{\theta}_n - \delta \nabla L(\tilde{\theta}_n)]$ where $\delta > 0$ is a stepsize and Π_K is the projection onto K , converges under mild assumptions to $\theta^* \in \text{argmin}_{\Theta_F} L$, since L is convex, see [70]. However, for any $\theta \in \Theta_F$, $\nabla L(\theta) = \pi_\theta(F)$ and evaluating this quantity is generally unfeasible. In order to approximate this high-dimensional integral, nested Laplace approximations were recently introduced [84, 72]. However, it is difficult to give upper-bounds on the bias of these approximations. In what follows, we rely on Monte-Carlo approximations for which we derive explicit upper-bounds on the bias. More precisely, assuming that it is possible to sample from π_θ then $\nabla L(\theta)$ can be approximated by $N^{-1} \sum_{k=1}^N F(X_k)$, where $(X_k)_{k \in \{1, \dots, N\}}$ are independently sampled from π_θ . In most of our applications it is not feasible to sample directly from π_θ , but we can construct Markov chains for which π_θ is an invariant probability measure. Then, under assumptions and using classical Markov chain theory, it can be shown that $N^{-1} \sum_{k=1}^N F(X_k)$ converges almost surely to $\pi_\theta(F)$, [29, Theorem 11.3.1]. Such examples of Markov chains include the Metropolis-Hastings algorithm [48], which uses a rejection step. In a high-dimensional setting, the acceptance ratio can be extremely low and the proposed new iterate is then always discarded. Hence, we focus on Markov chains without rejection step. In this scenario, π_θ is not an invariant measure of the Markov chain in general. However, for an appropriate choice of Markov chain, the bias between its actual invariant probability measure and the target probability measure π_θ can be explicitly controlled.

4.1.2 SOUL algorithm

First, we consider some regularity assumption on the measure μ with respect to the Lebesgue measure.

B1. $\mu \ll \text{Leb}$ and its Radon-Nikodym density w.r.t. to the Lebesgue measure is given for almost every $x \in \mathbb{R}^d$ by $\exp[-r(x)] / \int_{\mathbb{R}^d} \exp[-r(y)] dy$, where $r : \mathbb{R}^d \rightarrow \mathbb{R}$ is measurable.

Let $\theta = (\theta(0), \dots, \theta(p-1)) \in \mathbb{R}^p$ and consider the overdamped unadjusted Langevin algorithm, called ULA in [82], defined by $(\tilde{X}_n)_{n \in \mathbb{N}}$ with $\tilde{X}_0 = x \in \mathbb{R}^d$ and for any $n \in \mathbb{N}$

$$\tilde{X}_{n+1} = \tilde{X}_n - \gamma \left(\sum_{i=1}^p \theta(i) \nabla F_i(\tilde{X}_n) + \nabla r(\tilde{X}_n) \right) + \sqrt{2\gamma} Z_{n+1}, \quad (14)$$

where $\gamma > 0$ is a stepsize and $(Z_n)_{n \in \mathbb{N}}$ is a sequence of independent d -dimensional Gaussian random variables with zero mean and identity covariance matrix. This algorithm is the Euler-Maruyama discretization of the overdamped Langevin stochastic differential equation [33] for which π_θ is the invariant probability measure. The study of the geometric convergence of this Markov chain under various metrics was conducted in [32, 33, 23]. In

[26] a SA scheme, the Stochastic Optimization with Unadjusted Langevin (SOUL) Algorithm, is proposed in order to construct a sequence $(\theta_n)_{n \in \mathbb{N}}$ such that $(\theta_n)_{n \in \mathbb{N}}$ converges almost surely and in L^1 to some $\theta^* \in \operatorname{argmin}_{\Theta_F \cap \mathbb{K}} L$. Let $\theta_0 \in \mathbb{K}$ and $X_0^0 \in \mathbb{R}^d$. For any $n \in \mathbb{N}$ and $k \in \{0, \dots, m_n - 1\}$ we define

$$\begin{aligned} X_{k+1}^n &= X_k^n - \gamma_n \left(\sum_{i=1}^p \theta_n(i) \nabla F_i(X_k^n) + \nabla r(X_k^n) \right) + \sqrt{2\gamma_n} Z_{k+1}^n \quad \text{and } X_0^n = X_{n-1}^{m_n-1} ; \\ \theta_{n+1} &= \Pi_{\mathbb{K}} \left[\theta_n + \delta_{n+1} m_n^{-1} \sum_{k=1}^{m_n} F(X_k^n) \right] , \end{aligned} \tag{15}$$

where $(\delta_n)_{n \in \mathbb{N}^*}$ and $(\gamma_n)_{n \in \mathbb{N}}$ are sequence of positive stepsizes and the sequence $(Z_k^n)_{n \in \mathbb{N}, k \in \{1, \dots, m_n\}}$ is a sequence of independent d -dimensional Gaussian random variables with zero mean and covariance identity. By convention, $X_{-1}^{m_n-1} = X_0^0$. The condition $X_0^n = X_{n-1}^{m_n-1}$ for all $n \in \mathbb{N}$ is referred to as a warm-start condition.

To illustrate the expected behavior of the proposed SOUL algorithm (15), we consider the toy example where $x_0 \in \mathbb{R}$, $F(x) = x^2 - 4$ and $F(x_0) = 0$. In this case the maximum entropy distribution is given by the Gaussian distribution with zero mean and variance 4, see [69, Section 4.6.2]. One shows that, the optimal weight θ^* is given by $\theta^* = 1/8$. In Figure 1, we experimentally check the convergence of $(\theta_n)_{n \in \mathbb{N}}$. We set $r(x) = 0$ and observe that the sequence $(\theta_n)_{n \in \mathbb{N}}$ as well as the sequence $(\bar{\theta}_n)_{n \in \mathbb{N}}$ define for any $n \in \mathbb{N}$ by

$$\bar{\theta}_n = 0 \text{ if } n < N , \quad \bar{\theta}_n = \sum_{k=N}^n \delta_k \theta_k \Big/ \sum_{k=N}^n \delta_k \text{ otherwise } , \tag{16}$$

where $N \in \mathbb{N}$ is a fixed parameter, converge to θ^* . We now state our main results on the dependency on the dimension in the explicit error in SOUL.

4.2 Main results

In [26], the convergence of the sequence $(\theta_n)_{n \in \mathbb{N}}$ is studied under general assumptions. In this section, we complement these results in our setting. In particular, we show that the error bound in L^1 norm between $L(\theta_n)$ and $L(\theta^*)$ is upper bounded by a constant which depends polynomially in the dimension d . Let $\alpha \geq 1$, we consider the following assumption:

B2 (α). *There exists $\mathbb{K} \subset \mathbb{R}^p$ such that:*

- (a) *\mathbb{K} is a non-empty convex compact set with $\mathbb{K} \subset \operatorname{int}(\Theta_F)$ and we denote $R_{\mathbb{K}} > 0$ such that $\mathbb{K} \subset \bar{B}(0, R_{\mathbb{K}})$;*
- (b) *F is differentiable and there exists $M \geq 0$ such that for any $x, y \in \mathbb{R}^d$*

$$\|F(x) - F(y)\| \leq M(1 + \|x\|^{\alpha-1} + \|y\|^{\alpha-1}) \|x - y\| ;$$

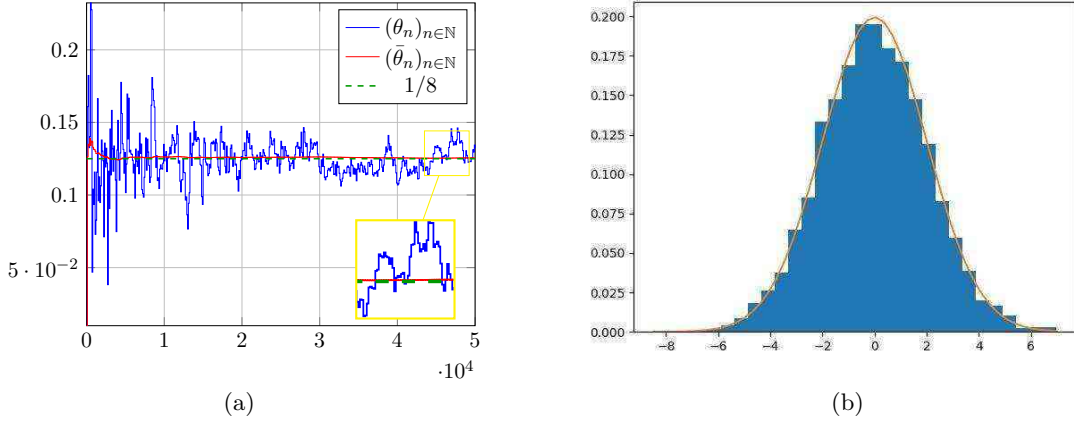


Figure 1: **Variance estimation.** In (a), the sequence of parameters $(\theta_n)_{n \in \mathbb{N}}$ (blue curve) and the sequence of average parameters $(\bar{\theta}_n)_{n \in \mathbb{N}}$ (red curve) converge to the optimal value $\theta^* = 1/8$. In (b) we illustrate empirically the convergence of the sequence $(X_0^n)_{n \in \mathbb{N}}$ to the Gaussian distribution with zero mean and variance 4 (orange curve) by plotting its histogram. In this example $N = 0$, $\delta_n = 0.1 \times n^{-0.7}$, $\gamma_n = 0.1 \times n^{-0.3}$ and $m_n = 10 \times \lceil n^{0.6} \rceil$.

(c) there exists $\theta^* \in \mathbf{K}$ solution of (Q).

Under **B2**(α) and **A2**(α) with $\alpha \geq 1$, a solution of (P) exists and is given by (6), see Proposition 4. In addition, L is differentiable on \mathbf{K} and we show in Proposition 4 that $\nabla L \in C^1(\mathbb{R}^d, \mathbb{R}^d)$, hence Lipschitz continuous over \mathbf{K} with constant $\sup_{\theta \in \mathbf{K}} \|\nabla^2 L(\theta)\|$. Conditions under which **B2**(α)-(c) is satisfied are given in Proposition 5.

Condition **B1** implies that the density of π_θ , see (6), with respect to the Lebesgue measure, is given for any $x \in \mathbb{R}^d$ by $(d\pi_\theta/d\text{Leb})(x) = \exp[-U(\theta, x)] / \int_{\mathbb{R}^d} \exp[-U(\theta, y)] dy$ with U defined for any $\theta \in \mathbf{K}$ and $x \in \mathbb{R}^d$ by

$$U(\theta, x) = \langle \theta, F(x) \rangle + r(x). \quad (17)$$

The mapping $U : \mathbf{K} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is referred to as the potential function. Consider the following assumption on U .

B3. There exist $U_i : \mathbf{K} \times \mathbb{R}^d \rightarrow \mathbb{R}$ with $i \in \{1, 2\}$ such that for any $\theta \in \mathbf{K}$ and $x \in \mathbb{R}^d$ $U(\theta, x) = U_1(\theta, x) + U_2(\theta, x)$. In addition,

- (a) there exists $N \geq 0$ such that for any $i \in \{1, 2\}$, $x \mapsto U_i(\theta, x)$ is continuously differentiable and for any $x, y \in \mathbb{R}^d$, $\|\nabla_x U_i(\theta, x) - \nabla_x U_i(\theta, y)\| \leq N \|x - y\|$;
- (b) there exists $m_1 > 0$ and $x^* \in \mathbb{R}^d$ such that for any $\theta \in \mathbf{K}$, $U_1(\theta, \cdot)$ is m_1 -strongly convex and $x^* \in \text{argmin}_{x \in \mathbb{R}^d} U_1(\theta, x)$;

(c) there exists $\mathcal{M} \geq 0$ such that for any $\theta \in \mathbf{K}$ and $x \in \mathbb{R}^d$, $\|\nabla_x U_2(\theta, x)\| \leq \mathcal{M}$;

We can relax the assumption that for any $\theta \in \mathbf{K}$, $x^* \in \operatorname{argmin}_{x \in \mathbb{R}^d} U_1(\theta, x)$ by the following: there exists $R \geq 0$ such that for any $\theta \in \mathbf{K}$, there exists $x_\theta^* \in \operatorname{argmin}_{x \in \mathbb{R}^d} U_1(\theta, x)$ and $x_\theta^* \in \bar{\mathbf{B}}(0, R)$. But for the sake of simplicity we do not consider this assumption. The general assumption **B3** is satisfied for both the Gaussian features and the CNN features introduced in Section 2.2. Indeed, if the features are Gaussian and the reference measure is Gaussian we recall that $\Theta_F = \mathcal{F}^{-1}[\mathfrak{A}^{-1}\{(-2\sigma^2)^{-1}, +\infty\}^d] \cap \mathbb{R}^d$ containing θ^* with θ^* given in (10), see Section 3.4. Then, $x \mapsto U(\theta, x)$ is a definite positive quadratic form associated with the symmetric matrix \mathbf{C}_θ , see (9). Setting \mathbf{N} and \mathbf{m} respectively the largest and lowest eigenvalues of \mathbf{C}_θ over \mathbf{K} we obtain that **B3** is satisfied with $U_1 = U$ and $U_2 = 0$.

In the case of CNN features, if the reference measure is a Gaussian distribution with zero mean and invertible covariance matrix \mathbf{C} , we obtain that for any $\theta \in \mathbb{R}^p$ and $x \in \mathbb{R}^d$

$$U(\theta, x) = \langle \theta, F(x) \rangle + x^T \mathbf{C}^{-1} x / 2 .$$

If in addition, φ is differentiable with Lipschitz derivative and for any $t \in \mathbb{R}$, $\sup_{t \in \mathbb{R}} |\varphi'(t)| < +\infty$, we have that **B3** is satisfied with for any $\theta \in \mathbf{K}$ and $x \in \mathbb{R}^d$

$$U_1(\theta, x) = x^T \mathbf{C}^{-1} x / 2 , \quad U_2(\theta, x) = \langle \theta, F(x) \rangle .$$

In particular the fact that U_2 is gradient-Lipschitz and Lipschitz is ensured by a straightforward recursion since for any $f \in C^1(\mathbb{R}^{d_3}, \mathbb{R}^{d_2})$ and $g \in C^1(\mathbb{R}^{d_2}, \mathbb{R}^{d_1})$, $x \mapsto d(g \circ f)(x)$ and $g \circ f$ Lipschitz if f, g, df and dg are Lipschitz. Note that the differentiability assumption is not met in classical convolutional neural networks such as VGG19 . Therefore, in all of our experiments we replace the max-pooling operator by a mean-pooling operator and the RELU function by a Continuously Differentiable Exponential Linear Unit (CELU), see [8].

We now state our main results in the case where U is a strongly convex potential, *i.e.* $U_2 = 0$.

Theorem 8. *Let $\alpha \geq 1$. Assume **A2**(α), **B1**, **B2**(α), **B3** with $U_2 = 0$. Let $(\gamma_n)_{n \in \mathbb{N}}$, $(\delta_n)_{n \in \mathbb{N}}$ be sequences of non-increasing positive real numbers and $(m_n)_{n \in \mathbb{N}}$ a sequence of positive integers satisfying $\delta_n < 1/(\sup_{\theta \in \mathbf{K}} \|\nabla^2 L(\theta)\|)$ and $\gamma_n < \min(\mathbf{m}_1/(2\mathbf{N}^2), 1/2)$ for any $n \in \mathbb{N}$. Then, there exists $(E_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}^*$*

$$\mathbb{E} \left[\left\{ \frac{\sum_{k=1}^n \delta_k L(\theta_k)}{\sum_{k=1}^n \delta_k} \right\} - \min_{\mathbf{K}} L \right] \leq E_n / \left(\sum_{k=1}^n \delta_k \right) ,$$

with for any $n \in \mathbb{N}^*$,

(a) if $m_n = m_0$ for all $n \in \mathbb{N}$ and $\sup_{n \in \mathbb{N}} |\delta_{n+1} - \delta_n| \delta_n^{-2} < +\infty$

$$E_n = C_1(1 + d^\varpi) \left(1 + \sum_{k=0}^{n-1} \delta_{k+1} \gamma_k^{1/2} + \sum_{k=0}^{n-1} \delta_{k+1} \gamma_{k+1}^{-5/2} (\gamma_k - \gamma_{k+1})^{1/2} + \sum_{k=0}^{n-1} \delta_{k+1}^2 / \gamma_k^{3/2} + \delta_{n+1} / \gamma_n \right) ;$$

(b) otherwise

$$E_n = C_2(1 + d^\varpi) \left(1 + \sum_{k=0}^{n-1} \delta_{k+1} \gamma_k^{1/2} + \sum_{k=0}^{n-1} \delta_{k+1} / (m_k \gamma_k) \right. \\ \left. + \sum_{k=0}^{n-1} \delta_{k+1}^2 \gamma_k + \sum_{k=0}^{n-1} \delta_{k+1}^2 / (m_k \gamma_k)^2 \right),$$

with $C_1, C_2, \varpi \geq 0$ which do not depend on the dimension d .

Proof. The proof is postponed to Appendix B.1. \square

It should be noted that Theorem 8-(b) applies even if for any $n \in \mathbb{N}$, X_0^n is distributed according to ν with $\nu \in \mathcal{P}_\alpha$, i.e. the warm-start procedure can be avoided.

In the case where for any $n \in \mathbb{N}$, $m_n = m_0$, $\gamma_n = \gamma_0$ and $\lim_{n \rightarrow +\infty} \delta_n = 0$ with $\sum_{k=0}^{+\infty} \delta_k = +\infty$, we obtain using [76, Problem 80, Part I] that, $\lim_{n \rightarrow +\infty} \sum_{k=0}^n \delta_k^2 / \sum_{k=0}^n \delta_k = 0$. Therefore, using Theorem 8-(a) we get that

$$\limsup_{n \rightarrow +\infty} \mathbb{E} \left[L(\bar{\theta}_n) \right] - \min_{\mathbb{K}} L \leq C_1(1 + d^\varpi) \gamma_0^{1/2},$$

with $\bar{\theta}_n = \sum_{k=1}^n \delta_k \theta_k / \sum_{k=1}^n \delta_k$ and using Theorem 8-(b) we get

$$\limsup_{n \rightarrow +\infty} \mathbb{E} \left[L(\bar{\theta}_n) \right] - \min_{\mathbb{K}} L \leq C_2(1 + d^\varpi) \left[\gamma_0^{1/2} + (m_0 \gamma_0)^{-1} \right].$$

Therefore, the minimum of L can be reached with arbitrary precision. Note that the constants C_1, C_2 do not have the same dependency with respect to the problem parameters and that C_2 is usually better than C_1 .

We now state our main results in the case where the potential is not convex anymore. We consider the following additional regularity assumption on F .

B4. $F \in C^1(\mathbb{R}^d, \mathbb{R}^p)$ and there exists $B \geq 0$ such that for any $x, y \in \mathbb{R}^d$

$$\|dF(x) - dF(y)\| \leq B \|x - y\|.$$

Theorem 9. Assume **A2**(1), **B1**, **B2**(1), **B3** and **B4**. Let $(\gamma_n)_{n \in \mathbb{N}}$, $(\delta_n)_{n \in \mathbb{N}}$ be sequences of non-increasing positive real numbers and $(m_n)_{n \in \mathbb{N}}$ a sequence of positive integers satisfying $\delta_n < 1/(\sup_{\theta \in \mathbb{K}} \|\nabla^2 L(\theta)\|)$ and $\gamma_n < \min(m_1/(8M^2), 1/2)$ for any $n \in \mathbb{N}$. Then, there exists $(E_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}^*$

$$\mathbb{E} \left[\left\{ \sum_{k=1}^n \delta_k L(\theta_k) / \sum_{k=1}^n \delta_k \right\} - \min_{\mathbb{K}} L \right] \leq E_n / \left(\sum_{k=1}^n \delta_k \right),$$

with for any $n \in \mathbb{N}^*$,

(a) if $m_n = m_0$, $\gamma_n = \gamma_0$ for all $n \in \mathbb{N}$ and $\sup_{n \in \mathbb{N}} |\delta_{n+1} - \delta_n| \delta_n^{-2} < +\infty$

$$E_n = C_1(1 + d^\varpi) \left(1 + \sum_{k=0}^{n-1} \delta_{k+1} \gamma_0^{1/2} + \sum_{k=0}^{n-1} \delta_{k+1}^2 / \gamma_0 + \delta_n / \gamma_0 \right) ;$$

(b) else

$$E_n = C_2(1 + d^\varpi) \left(1 + \sum_{k=0}^{n-1} \delta_{k+1} \gamma_k^{1/2} + \sum_{k=0}^{n-1} \delta_{k+1} / (m_k \gamma_k) + \sum_{k=0}^{n-1} \delta_{k+1}^2 \right) ,$$

with $C_1, C_2, \varpi \geq 0$ which do not depend on the dimension d .

Proof. The proof is postponed to Appendix B.2 □

The discussion conducted after Theorem 8 is still valid in this case. Also, note that the results of Theorem 8-(b) are covered by Theorem 9-(b) under B4. However, in Theorem 9-(a), $\gamma_n = \gamma_0$ for all $n \in \mathbb{N}$ whereas in Theorem 8-(a), $(\gamma_n)_{n \in \mathbb{N}}$ is any arbitrary non-increasing sequence of positive numbers. Theorem 9 follows from more general results derived in Theorem 20 and Theorem 24.

4.3 Links with macrocanonical models

In this section, we present qualitative results on the microcanonical model and the asymptotic behavior of the macrocanonical model for specific geometrical constraints. We start by recalling a result concerning the convergence of the sampler of the microcanonical model from [14, Theorem 3.7-(i)].

Let $\nu_0 \ll \text{Leb}$ be an initial probability measure. We consider the sequence of probability measures $(\nu_n)_{n \in \mathbb{N}}$ defined by the following recursion: for any $n \in \mathbb{N}$,

$$\nu_{n+1} = \Phi_{\sharp}(\nu_n) , \tag{18}$$

where $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined for any $x \in \mathbb{R}^d$ by $\Phi(x) = x - \gamma \text{d}F(x)^T F(x)$, with $\gamma > 0$ a stepsize. Namely, for all $n \in \mathbb{N}$, ν_n is the pushforward measure of ν_0 by n steps of the gradient descent for the the loss function $x \mapsto \|F(x)\|^2$.

Theorem 10 ([14]). *Let $F \in C^2(\mathbb{R}^d, \mathbb{R}^p)$ such that for any compact set $K \subset \mathbb{R}^p$, $F^{-1}(K)$ is compact. Assume B4 and that for any $x \in F^{-1}(\{0\})$, $\det(\text{D}F(x)\text{D}F(x)^T) > 0$. In addition, assume that F satisfies the strict saddle property, i.e. defining $M_v \in \text{S}_d(\mathbb{R})$ for any $v = (v_1, \dots, v_p) \in \text{Ker}(\text{d}F(x)^T) \setminus \{0\}$, by*

$$M_v = \sum_{k=1}^p v_k \nabla^2 F_k(x) + \text{D}F(x)^T \text{D}F(x) ,$$

M_ν admits at least one negative eigenvalue. Then, if $\gamma \in (0, \mathbf{B}^{-1})$ and X_0 is distributed according to ν_0 , $(X_n)_{n \in \mathbb{N}}$ defined for any $n \in \mathbb{N}$ by $X_{n+1} = \Phi(X_n)$ converges almost surely to a random variable X_∞ . In addition, $\nu_\infty(F^{-1}(\{0\})^c) = 0$.

Let $\mathbf{A} = F^{-1}(\{0\})$. If \mathbf{A} is compact, the microcanonical model, see Definition 1, associated with the reference measure Leb and the constraints F , is given by the uniform distribution over \mathbf{A} , denoted $\nu_{\mathbf{A}}$. If ν_∞ were the microcanonical model associated with F then we should have $\nu_\infty = \nu_{\mathbf{A}}$. However, as illustrated in Figure 2, ν_∞ strongly depends on the initial probability measure ν_0 . Let $(\Phi_n)_{n \in \mathbb{N}}$ be defined by $\Phi_0 = \text{Id}$ and the following recursion: $\Phi_{n+1} = \Phi \circ \Phi_n$. Note that X_n in Theorem 10 is distributed according to $(\Phi_n)_\#(\nu_0)$. It is shown in the proof of [14, Theorem 3.7] that $\Phi_\infty = \lim_n \Phi_n$ is well-defined. Since for any $n \in \mathbb{N}$, Φ_n is Lebesgue measurable so is Φ_∞ . Therefore we have that $\nu_\infty = (\Phi_\infty)_\#(\nu_0)$.

In what follows we prove that considering specific constraint functions f_ε , one can construct an explicit probability measure π_∞ such that π_∞ is supported on $F^{-1}(\{0\})$ and π_∞ is the limit of macrocanonical models associated with f_ε and some reference probability measure μ . Let $\varepsilon > 0$. We define $f_\varepsilon : \mathbb{R}^d \rightarrow \mathbb{R}$ such that for any $x \in \mathbb{R}^d$, $f_\varepsilon(x) = \|F(x)\|^2 - \varepsilon$. We denote π_ε the macrocanonical model, see Definition 2, associated with f_ε when it exists.

Proposition 11. *Assume A2(2) and that for any non-empty open set $\mathbf{A} \subset \mathbb{R}^d$, $\mu(\mathbf{A}) > 0$. Let F be given by (2), assume that $1 \in \mathbf{j}$ and that there exists $k \in \{1, \dots, c_1\}$ such that for any $x \in \mathbb{R}^d$ with $x \neq 0$ there exists $\ell \in \{1, \dots, n_1\}$ with $e_\ell^T \tilde{A}_1^k x > 0$. Then there exists $\varepsilon_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0)$, π_ε exists. In addition, the following propositions hold:*

(a) *Assume that $\mu(F^{-1}(\{0\})) > 0$ then $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon = \pi_\infty$, with for any $x \in \mathbb{R}^d$*

$$\frac{d\pi_\infty}{d\mu}(x) = \frac{\mathbb{1}_{F^{-1}(\{0\})}(x)}{\mu(F^{-1}(\{0\}))}.$$

(b) *Assume that $F^{-1}(\{0\}) = \{x_1, \dots, x_K\}$ with $(x_i)_{i \in \{1, \dots, K\}} \in (\mathbb{R}^d)^K$, $K \in \mathbb{N}^*$, $\varphi \in C^3(\mathbb{R})$, $x \mapsto (d\mu/d\text{Leb})(x)$ is continuous and $(d\mu/d\text{Leb})^{-1}(F^{-1}(\{0\})) \neq \{0\}$. Let $H(x) = \nabla^2(\|F(\cdot)\|^2)(x)$ and assume that for any $x \in F^{-1}(\{0\})$, $\det H(x) \neq 0$. Then $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon = \pi_\infty$ with*

$$\pi_\infty = \sum_{i=1}^K \frac{\frac{d\mu}{d\text{Leb}}(x_i) \det(H(x_i))}{\sum_{j=1}^K \frac{d\mu}{d\text{Leb}}(x_j) \det(H(x_j))} \delta_{x_i}.$$

(c) *Assume that $F^{-1}(\{0\})$ is a smooth compact manifold, $\varphi \in C^3(\mathbb{R})$, $x \mapsto (d\mu/d\text{Leb})(x)$ is continuous and $(d\mu/d\text{Leb})^{-1}(F^{-1}(\{0\})) \neq \{0\}$. Let $H(x) = \nabla^2(\|F(\cdot)\|^2)(x)$ and assume that for any $x \in F^{-1}(\{0\})$, $\det H(x) \neq 0$. Then $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon = \pi_\infty$ with for any $x \in \mathbb{R}^d$*

$$\frac{d\pi_\infty}{d\mathcal{H}}(x) = \frac{\mathbb{1}_{F^{-1}(\{0\})}(x) \frac{d\mu}{d\text{Leb}}(x) \det H(x)}{\int_{F^{-1}(\{0\})} \frac{d\mu}{d\text{Leb}}(y) \det H(y) d\mathcal{H}(y)},$$

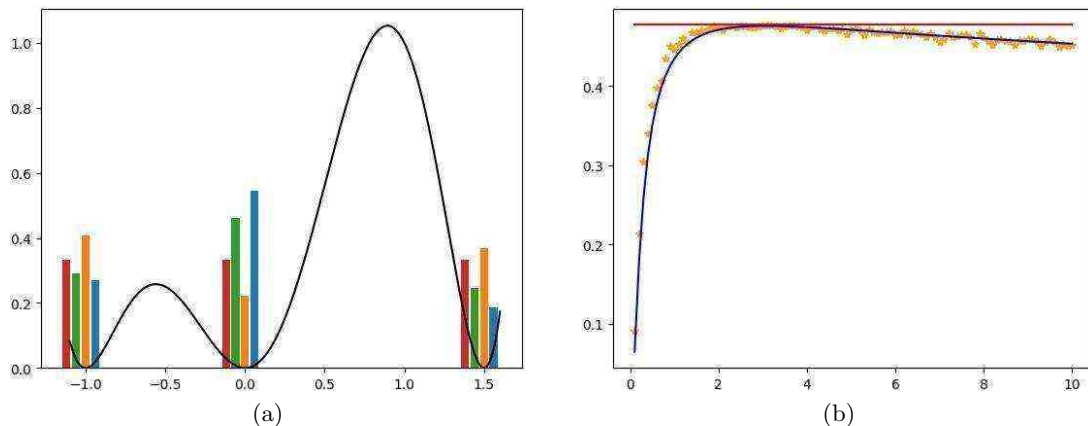


Figure 2: **Microcanonical sampling scheme.** In this one-dimensional toy example the features are given by $F(x) = (x + 1)^2 x^2 (x - 1.5)^2$ (black curve). The microcanonical model associated with these features is the uniform distribution over $\{-1, 0, 1.5\}$ (red bars). In (a), we plot, the distribution $(\Phi_\infty)_\#(\nu_0)$ for different initial distributions ν_0 , standard Gaussian (blue bars), uniform over $[-3, 3]$ (orange bars) and standard Cauchy (green bars). The distribution $(\Phi_\infty)_\#(\nu_0)$ is approximated by sampling 10^3 points according to ν_0 and performing the recursion associated with (18) for these points for 10^4 iterations. None of the initial distribution yields a distribution $(\Phi_\infty)_\#(\nu_0)$ which is the uniform distribution. Let ν_0 be a Gaussian distribution with zero mean and variance σ^2 with $\sigma > 0$. In (b), we show the dependency of the entropy of $(\Phi_\infty)_\#(\nu_0)$ with respect to σ^2 (orange points). The distribution $(\Phi_\infty)_\#(\nu_0)$ is approximated by sampling 10^3 points according to ν_0 and performing the recursion associated with (18) for these points for 10^3 iterations. Then, we compute its entropy and show that it is close to the one given by numerical integration (blue curve). We also plot the entropy upper-bound $\log_{10}(3)$ (red curve) given by the uniform distribution on $\{-1, 0, 1.5\}$.

where \mathcal{H} is the intrinsic measure on $F^{-1}(\{0\})$, see [11, Chapter 6].

Proof. The proof is postponed to Appendix B.3 □

In Proposition 11, if $F^{-1}(\{0\}) \subset \bar{B}(0, R)$, with $R > 0$ then letting μ such that for any $x \in \mathbb{R}^d$, $\frac{d\mu}{d\text{Leb}}(x) = \mathbb{1}_{\bar{B}(0, R)}(x)/\text{Leb}(\bar{B}(0, R))$, we have $\pi_\infty \ll \text{Leb}$ and for any $x \in \mathbb{R}^d$

$$\frac{d\pi_\infty}{d\text{Leb}}(x) = \left(\frac{d\pi_\infty}{d\mu}(x) \right) \left(\frac{d\mu}{d\text{Leb}}(x) \right) = \frac{\mathbb{1}_{F^{-1}(\{0\})}(x)}{\mu(F^{-1}(\{0\}))\text{Leb}(\bar{B}(0, R))} = \frac{\mathbb{1}_{F^{-1}(\{0\})}(x)}{\text{Leb}(F^{-1}(\{0\}))},$$

hence π_∞ is a microcanonical model. Finally, note that the conclusions of Proposition 11-(a)-(b) do not hold if $d > p$, since in this case $\det(H(x)) = 0$ for any $x \in \mathbb{R}^d$ such that $F(x) = 0$.

5 Experiments

In this section, we assess the computational efficiency of SOUL algorithm (15) for texture synthesis. Variants of the original methodology are presented in Appendix C.1.

5.1 Periodic Gaussian model

First, we consider the toy problem of periodic Gaussian texture synthesis, see Section 2.2 for details. Note that the presentation of the model was conducted for one dimensional signals. The extension of our findings two dimensional signals is straightforward upon replacing the one dimensional Fourier transform by its two dimensional counterpart. We recall that in this case the macrocanonical model is explicit and given by a measure π_{θ^*} which is the probability distribution of $X = d^{-1/2}(x_0 * Z)$ where Z is a standard d -dimensional Gaussian random variable, see Section 3.4. This model was introduced in the context of computer graphics in [91] and its mathematical study was conducted in [41, 40, 42].

5.1.1 Empirical convergence

We consider a 8×8 image, denoted x_0 , corrupted by some noise, so that $\mathcal{F}(x_0)$ is non-zero everywhere on the 8×8 grid. The reference measure μ is a Gaussian distribution with zero mean and diagonal covariance matrix with diagonal coefficients given by σ^2 . In this setting, $d = p = 64$ and using (10) we have

$$\theta^* = \mathcal{F}^{-1}(d|\mathcal{F}(x_0)|^{-2} - \sigma^{-2})/2.$$

Using the spatial translation invariance property of π_{θ^*} , see (11), we identify four configurations which are equally likely to be sampled by π_{θ^*} , see Figure 3.

The images $(X_0^n)_{n \in \mathbb{N}}$ generated by the SOUL algorithm (15) are approximate samples of π_{θ^*} for n large enough. The configurations identified in Figure 3 are recovered during

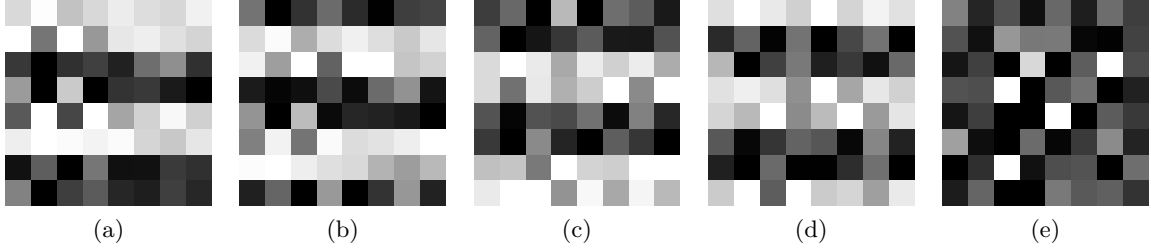


Figure 3: **Exemplar images and parameters.** The exemplar image x_0 is shown in (a). Translated versions of this image, which are equally likely to be sampled by π_{θ^*} are presented in (b), (c) and (d). The target parameter θ^* is shown in (d).

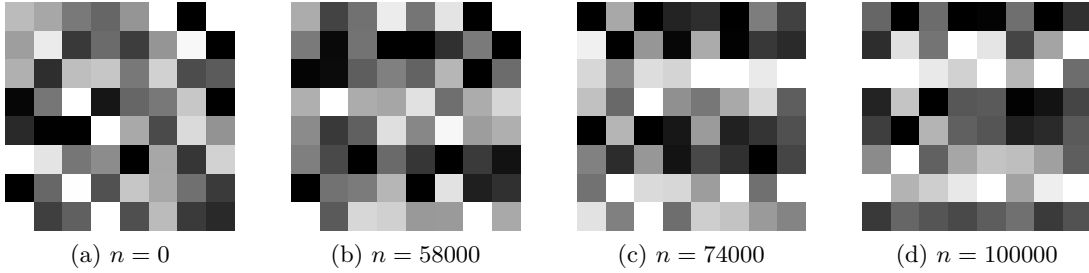


Figure 4: **Sequence of images.** The initialization (a) of the algorithm is some white noise, *i.e.* the realization of a standard Gaussian random variable on the 8×8 grid. We then show some selected samples (b)-(d) of the sequence generated with fixed parameters $\delta_n = 10^{-1}$, $\gamma_n = 10^{-4}$ and $m_n = 1$. Note that these samples are visually close to the ones presented in Figure 3.

one run of the algorithm, see Figure 4. A video of the evolution of the sequence $(X_0^n)_{n \in \mathbb{N}}$ is available at https://vdeborto.github.io/publication/texture_soul/.

The main theoretical results in Theorem 8 deal with the error between $L(\bar{\theta}_n)_{n \in \mathbb{N}}$ and $\operatorname{argmin}_{\theta \in \Theta} L(\theta)$, where $(\bar{\theta}_n)_{n \in \mathbb{N}}$ is given by (16). Selecting fixed parameters, $\gamma_n = 10^{-4}$, $\delta_n = 10^{-1}$ and $m_n = 1$ we observe the convergence of the sequence $(\theta_n)_{n \in \mathbb{N}}$ towards a biased estimate of θ^* . The Normalized Root Mean Square Error (NRMSE) defined for any $n \in \mathbb{N}$ by

$$N(\theta_n) = \|\theta_n - \theta^*\|_2 / \|\theta^*\|_2 ,$$

is upper bounded by 0.2 for $n \geq 4 \times 10^4$, see Figure 5. In Figure 6, we show that this error level yields satisfactory parameters from a visual point of view. We highlight that $\mathcal{F}(\theta^*)$ corresponds to the precision matrix (up to a constant factor) of the Gaussian model under study.

The previous experiment suggests to set $\gamma_n = \gamma > 0$, $\delta_n = \delta > 0$ and $m_n = m \in \mathbb{N}^*$, at

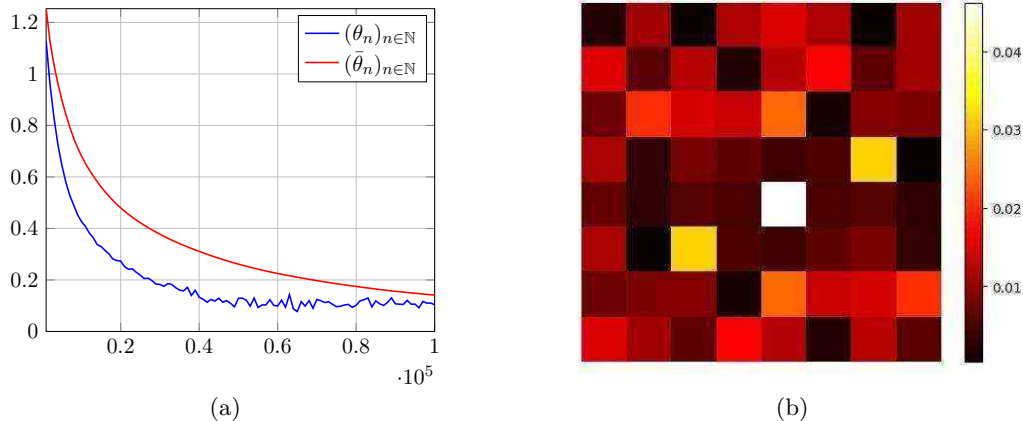


Figure 5: **Convergence of the parameters.** We recall that the parameters are initialized with $\theta_0 = 0$ and that $\gamma_n = 10^{-4}$, $\delta_n = 10^{-1}$ and $m_n = 1$. The NRMSE error in (a) rapidly decreases before oscillating around 0.1 (blue curve). A similar smoothed behaviour can be observed for the averaged sequence $(\bar{\theta}_n)_{n \in \mathbb{N}}$ (red curve). The heatmap of the NRMSE, *i.e.* a pixel $i \in \{0, \dots, 7\}^2$ in (b) corresponds to $(\theta_n(i) - \theta^*(i))^2 / \|\theta^*\|^2$.

least for a burnin period. When the behavior of the sequence $(\theta_n)_{n \in \mathbb{N}}$ becomes oscillatory, the setting can be changed in order to obtain a better approximation of θ^* . We investigate the long-time behavior after a burnin period of $N = 5 \times 10^4$ iterations with $m_n = 1$, $\gamma_n = 10^{-4}$ and $\delta_n = 10^{-1}$. After this period we set $m_n = \lceil \bar{n}^a \rceil$, $\gamma_n = \bar{n}^{-b}$ and $\delta_n = \bar{n}^{-c}$ with $\bar{n} = n - N + 1$, $a, b, c > 0$. We observe that the NRMSE error decreases from 0.1 to 0.06 for appropriate choices of rates, see Figure 7. Nevertheless, this improvement comes at a cost since the number of Markov chain iterations is no longer equal to the number of iterations n and grows as $\lceil \bar{n}^a \rceil$.

The previous comments along with Figure 7 suggest to set fix hyperparameters with $m_n = 1$ for all $n \in \mathbb{N}$. This is a good strategy to obtain acceptable approximations of the target parameter θ^* in reasonable time. However, the sampled images move slowly between the acceptable configurations of Figure 3. Increasing the fixed batch size, *i.e.* increasing m_n , for instance setting $m_n = 10^2$ for all $n \in \mathbb{N}$ we obtain more innovation in the chain $(X_0^n)_{n \in \mathbb{N}}$. Namely, for the same number of Markov chain iterations the chain $(X_0^n)_{n \in \mathbb{N}}$ visits more different acceptable configurations for $m_n = 10^2$ than for $m_n = 1$, see Figure 8. However, if $m_n = 10^2$, the NRMSE error of the sequence $(\theta_n)_{n \in \mathbb{N}}$ has a lower decrease rate than if $m_n = 1$.

Therefore, the hyperparameters of the algorithm should be adapted for the problem at hand. If we are interested in finding θ^* then fixed settings for a burnin period followed by an eventual run of the algorithm with increasing batch size and decreasing stepsizes is recommended. However, if we are concerned with the innovation of the sequence $(X_0^n)_{n \in \mathbb{N}}$

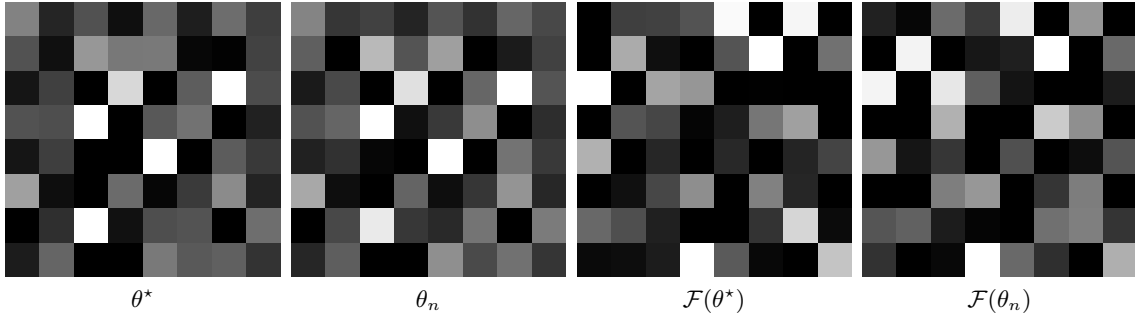


Figure 6: *Visual evaluation of parameters.* We display the target parameters θ^* and the parameters obtained after 10^6 iterations of the algorithm. Similarly we display the discrete Fourier transform of the target parameters $\mathcal{F}(\theta^*)$ and the Fourier transform of the parameters after 10^6 iterations. There is no visual difference between θ^* and θ_n .

then larger batch sizes, not necessarily increasing, are recommended. In what follows we experimentally assess some generalizations of the SOUL algorithm.

5.2 Neural network features

5.2.1 Spatially averaged CNN features

We now investigate the case where the features are given by a convolutional neural network, see Section 2.2. We briefly recall that this model is similar to the one introduced by [44] for microcanonical models but instead of considering Gram matrices for different layers of a convolutional neural network, we consider the means of different layers and channels for the same convolutional neural network to build the features. In our experiments we fix $\mathbf{K} = [-10^4, 10^4]^d$.

Model hyperparameters In the proposed model a few hyperparameters must be selected. First, a convolutional neural network architecture is to be chosen. We use the VGG19 model since it has been highlighted by [90, 44] that such an architecture is well suited for the task of texture synthesis. In [44] the neural network is pretrained on a classification task, see [85]. We first assess that this pretraining is a crucial step in our model in Figure 9. Indeed, if for each convolutional layer ℓ and channel c , the pretrained filters are replaced by filters with weights given by a Gaussian random variable which has same mean and same variance as the pretrained filters then no visually satisfying results are obtained.

Another hyperparameter of the model is the set \mathbf{j} of layers we consider to build our features, in (2). We consider three settings: (i) shallow network; (ii) deep network; (iii) full network. The structure of the network is recalled in Appendix D. In (i) we set $\mathbf{j} = \{1, 3, 6, 8, 11, 13\}$, in (ii) we set $\mathbf{j} = \{15, 24, 26, 31\}$ and in (iii) we set $\mathbf{j} =$

$\{1, 3, 6, 8, 11, 13, 15, 24, 26, 31\}$. Note that in the restricted models (i) and (ii) the dimension of the parameter space is reduced to $p = 896$ respectively $p = 1792$, whereas in the full model $p = 2688$. The influence of \mathbf{j} is visually investigated in Figure 10. In what follows we consider the full CNN model given by (iii) in order to be able to synthesize a wide variety of texture images.

It has been observed, in the case of microcanonical model, that using only CNN based features is not sufficient to describe all the textures. For instance in [64], the authors propose to add spectrum constraints in order to reimpose some spatial arrangement in the images. Similarly we can combine our neural network features with pixel-based features. In order to impose some color statistics we set $F_{\text{color}}^{\text{m}} : \mathbb{R}^d \rightarrow \mathbb{R}^3$ and $F_{\text{color}}^{\text{cov}} : \mathbb{R}^d \rightarrow \mathcal{S}_3(\mathbb{R})$ defined for any $i \in \{1, 2, 3\}$ by

$$\begin{aligned} \tilde{F}_{\text{color}}^{\text{m}}(x)(i) &= D^{-1} \sum_{k=1}^D x_i(k) , \\ \tilde{F}_{\text{color}}^{\text{cov}}(x) &= D^{-1} \begin{pmatrix} x_1 - \tilde{F}_{\text{color}}^{\text{m}}(x)(1) \\ x_2 - \tilde{F}_{\text{color}}^{\text{m}}(x)(2) \\ x_3 - \tilde{F}_{\text{color}}^{\text{m}}(x)(3) \end{pmatrix} \begin{pmatrix} x_1 - \tilde{F}_{\text{color}}^{\text{m}}(x)(1) \\ x_2 - \tilde{F}_{\text{color}}^{\text{m}}(x)(2) \\ x_3 - \tilde{F}_{\text{color}}^{\text{m}}(x)(3) \end{pmatrix}^{\text{T}} , \\ F_{\text{color}}^{\text{m}}(x) &= \tilde{F}_{\text{color}}^{\text{m}}(x) - \tilde{F}_{\text{color}}^{\text{m}}(x_0) , \quad F_{\text{color}}^{\text{cov}}(x) = \tilde{F}_{\text{color}}^{\text{cov}}(x) - \tilde{F}_{\text{color}}^{\text{cov}}(x_0) . \end{aligned} \tag{19}$$

where $d = 3D$ and $x = (x_1, x_2, x_3)$ where x_i corresponds to the i -th color channel of x . These features add 9 more parameters to the model. We refer to this model as the CNN + color features. Doing so the color statistics are imposed in expectation. It is also natural to ask that all the produced images have exactly the same color statistics as the exemplar image, *i.e.* that the equality holds almost surely. This procedure can be implemented by reimposing at each Langevin step the mean and the color covariance matrix of the images. We call this model CNN + color projection. The effect of imposing, in expectation or almost surely, the color constraints is investigated in Figure 11 and we observe that the proposed modifications do reimpose the color statistics of order 1 and 2.

Behavior of the parameter sequence We now study the behavior of the sequence $(\theta_n)_{n \in \mathbb{N}}$. In Figure 12 we present the evolution of $(\theta_n)_{n \in \mathbb{N}}$ for some layers in \mathbf{j} and three channels for each layer. The sequence $(\theta_n)_{n \in \mathbb{N}}$ does not converge, even though we observe some stabilization of the averaged sequences. The reasons for the failure of the convergence are twofold. First, in all our settings we fix the hyperparameters as follows: $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$ but run only 10^5 iterations. Considering a continuous Langevin dynamics, the images we observe correspond to a time $T = 10^5 \times \gamma_n = 1$ of the evolution. Increasing the stepsize γ_n is not an option since it yields diverging sequences of images. Second, the chain is slowly mixing and therefore it is hard to produce entirely different, yet visually coherent, samples with one run of SOUL. As a consequence, the Markov chain is prevented from efficiently exploring the image space.

It appears that the algorithm produces good visual results even though the parameter sequence is not stable.

Arbitrary size synthesis We assess in Figure 13 that contrary to the algorithm proposed in [65], our implementation can produce arbitrary large images from one input. Indeed, if for any $j \in \{1, \dots, M\}$ and $k \in \{1, \dots, c_j\}$, \tilde{A}_j^k in (1) is given by a convolutional operator, (2) can be defined for any $d \in \mathbb{N}$. The number of features does not depend on the size of the image but only on the number of layers selected in the network, since we average the neural network response in (2).

5.2.2 Comparison with existing methods

In this section we compare the proposed algorithm with several state of the art exemplar-based texture synthesis methods. We use the CNN + color version of the features and set $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$. The algorithm is run for 10^4 iterations for each image. For each comparison we systematically include the results obtained with the methodology proposed in [44], which is a microcanonical methodology using Gram matrices computed on neural network outputs as features.

First, we consider the Portilla-Simoncelli algorithm [77], see Section 5.2.2, which is a microcanonical based methodology and does not rely on neural network features, see Section 5.2.2. Our algorithm and the one from [44] provide visually satisfying results, whereas the method from [77] fails to produce realistic images.

We test our algorithm on texture images which do not exhibit salient spatial structures. Section 5.2.2 shows the results obtained using the Generative Adversarial Network approach proposed [54] in Section 5.2.2. It was already noted in [79, Figure 26] that this generative fails to produce high quality image in this case. On the other hand, our algorithm and the one from [44] yield good visual results.

We also compare our algorithm to the one of [65] in Section 5.2.2. In [65], the authors propose a similar macrocanonical methodology but do not consider more than one convolutional neural network layer to build their features.

Another experiment on highly regular textures, comparing our algorithm with the ones of [64] and [45], is presented in Appendix C.2.

5.2.3 Texture style transfer

We conclude this experimental part by considering other applications than texture synthesis and assess that the proposed algorithm can be used for the task of style transfer. Indeed given one content image x_{content} , a style image x_{style} , not necessarily of the same size, and $\mathbf{j}_{\text{content}} \subset \mathbf{j}$ we consider the same CNN feature as before but x_0 is replaced by x_{content} for $j \in \mathbf{j}_{\text{content}}$ in (2). In the rest of the neural network features, x_0 is replaced by x_{style} in (2),

i.e.

$$F(x) = \left(\overline{\mathcal{G}}_j^k(x) - \overline{\mathcal{G}}_j^k(x_0^j) \right)_{j \in \mathbf{j}, k \in \{1, \dots, c_j\}},$$

with $x_0^j = x_{\text{content}}$ if $j \in \mathbf{j}_{\text{content}}$ and x_{style} otherwise. These new features are well-suited to perform a style transfer task as illustrated in Figure 17 with $\mathbf{j}_{\text{content}} = \{1, 3, 6, 8, 11\}$ and $\mathbf{j} = \{1, 3, 6, 8, 11, 13, 15, 24, 26, 31\}$.

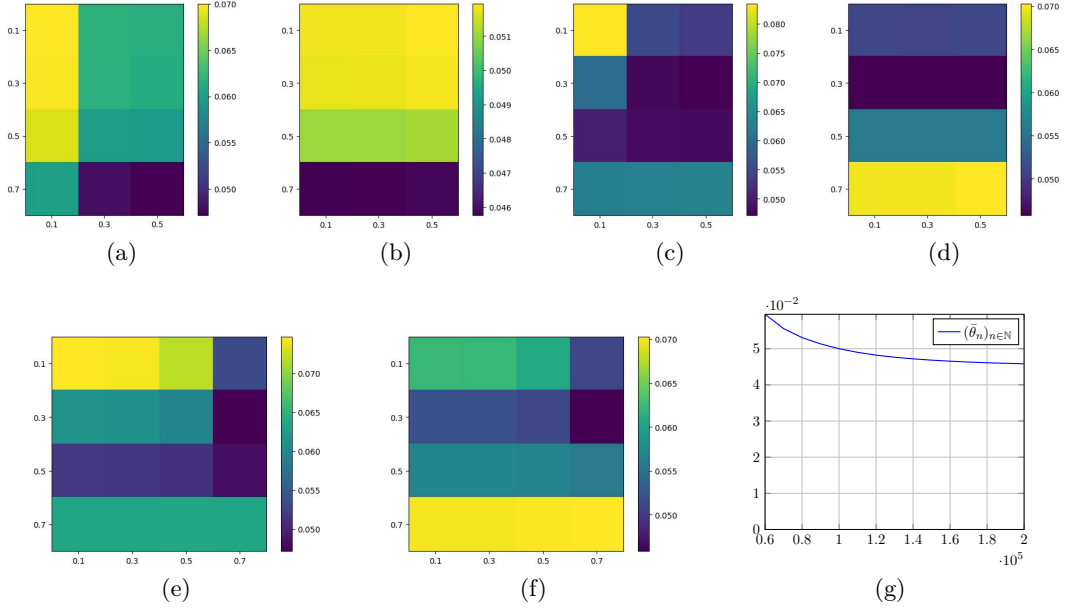


Figure 7: **Evolution of the error.** In (a) and (b) we present the heatmap of the NRMSE error between $\theta_{2 \times 10^5}$ and θ^* in (a) and $\bar{\theta}_{2 \times 10^5}$ and θ^* in (b), given different values of $b, c > 0$ where $\gamma_n = 10^{-4} \times \bar{n}^{-b}$ and $m_n = \lceil \bar{n}^c \rceil$ with $\delta_n = 10^{-1} \times \bar{n}^{-0.3}$ and $\bar{n} = n - N + 1$ with $N = 5 \times 10^4$. On the y -axis in (a) and (b) we represent the different values for parameter b and on the x -axis the different values for parameter c . Similarly, in (c) and (d) we present the heatmap of the NRMSE error between $\theta_{2 \times 10^5}$ and θ^* in (c) and $\bar{\theta}_{2 \times 10^5}$ and θ^* in (d), given different values of $a, c > 0$ where $\delta_n = 10^{-1} \times \bar{n}^{-a}$ and $m_n = \lceil \bar{n}^c \rceil$ with $\gamma_n = 10^{-4} \times \bar{n}^{-0.7}$. On the y -axis in (c) and (d) we represent the different values for parameter a and on the x -axis the different values for parameter c . In (e) and (f) we present the heatmap of the NRMSE error between $\theta_{2 \times 10^5}$ and θ^* in (e) and $\bar{\theta}_{2 \times 10^5}$ and θ^* in (f), given different values of $b, c > 0$ where $\delta_n = 10^{-1} \bar{n}^{-a}$ and $\gamma_n = \lceil \bar{n}^{-b} \rceil$ with $m_n = \lceil \bar{n}^{0.5} \rceil$. On the y -axis in (e) and (f) we represent the different values for parameter a and on the x -axis the different values for parameter b . A plot of the NRMSE for the averaged sequence is presented in (g) with $a = 0.3$ and $c = 0.7$.

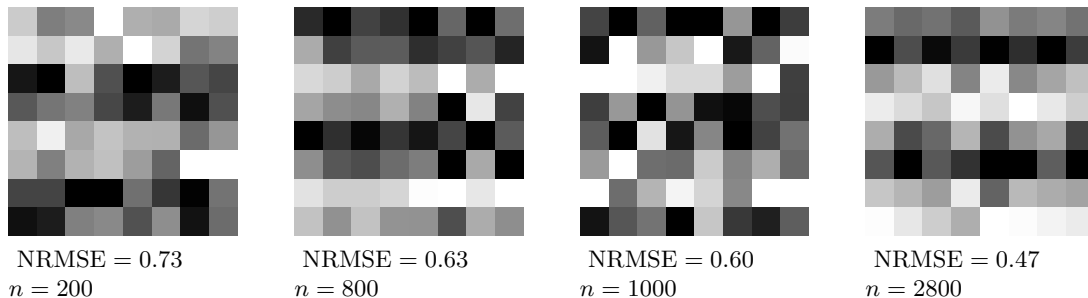


Figure 8: **Larger batch sizes improve visual quality.** The algorithm with $\delta_n = 10^{-1}$, $\gamma_n = 10^{-4}$ and $m_n = 10^2$ produces more diverse samples than the ones obtained with the same algorithm and $m_n = 1$, see Figure 3. Note that the NRMSE errors 0.73, 0.63, 0.60 and 0.47 are still high.

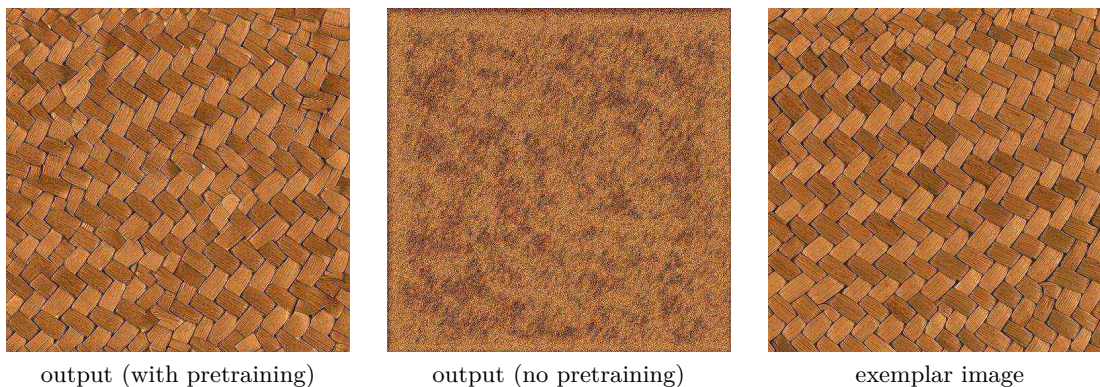


Figure 9: **Influence of the pretraining.** The exemplar image on the right is a 512×512 color image. We present the output of the SOUL algorithm on this image after 10^4 iterations. The hyperparameters are fixed as follows: $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$.

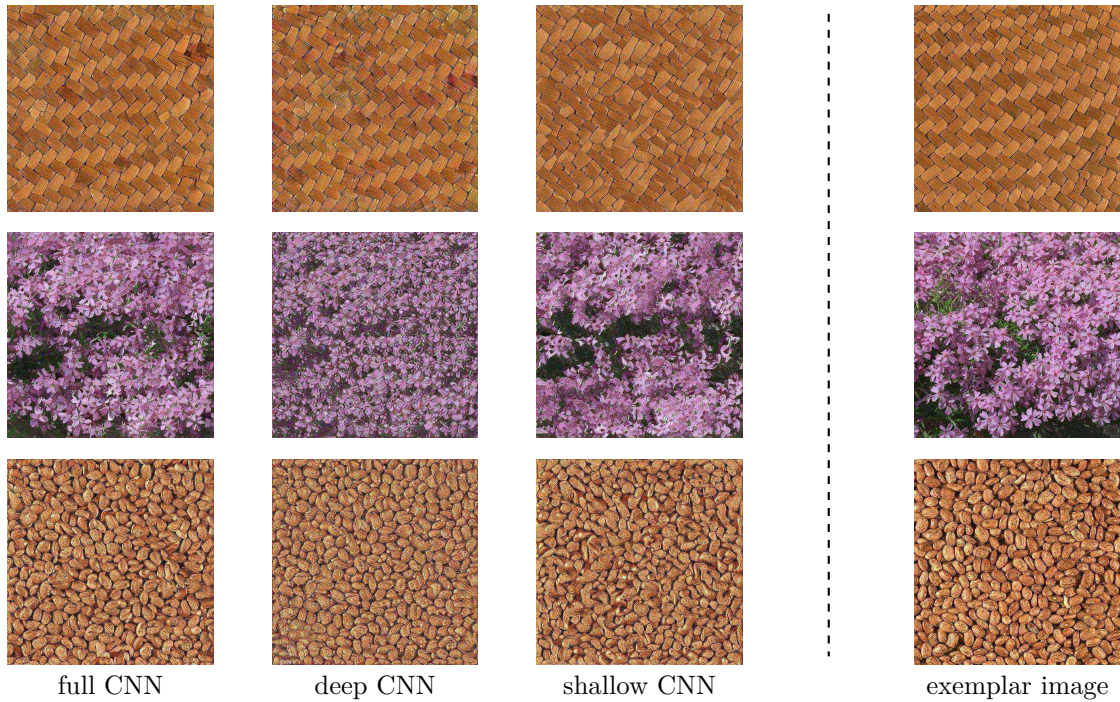


Figure 10: ***Influence of j.*** As expected the best visual results of SOUL after 10^4 iterations are obtained with the full CNN setting. The local structure and some details (the petals of the flowers, the form of the beans) are lost when using the shallow CNN setting. On the other hand, using only the deep part of the CNN is not suitable for texture with strong low frequency components. For instance in the flower image, almost no grass is retrieved when using the deep CNN setting. The hyperparameters are fixed as follows: $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$.

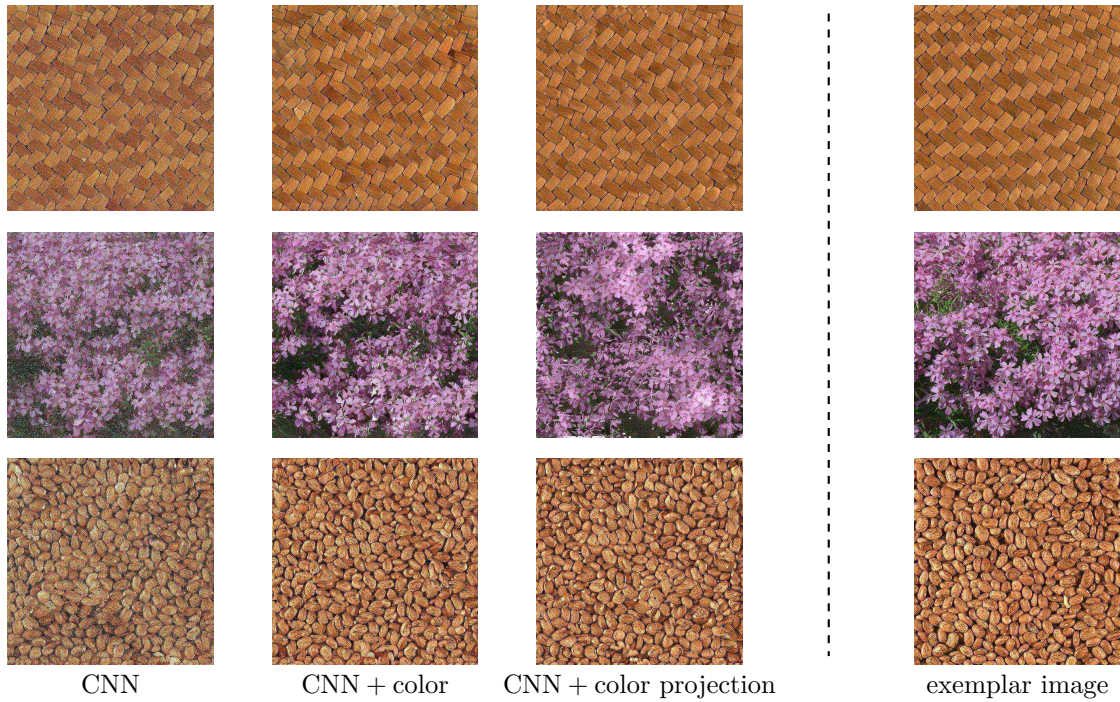


Figure 11: *Color models.* The SOUL algorithm with CNN features yields images with less contrast than the exemplar images. To address this issue we either introduce color features in the model (CNN + color), see (19), or reimpose the mean and color covariance of the image after each Langevin iteration (CNN + color projection). The results are similar for both methods. The hyperparameters are fixed as follows: $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$.

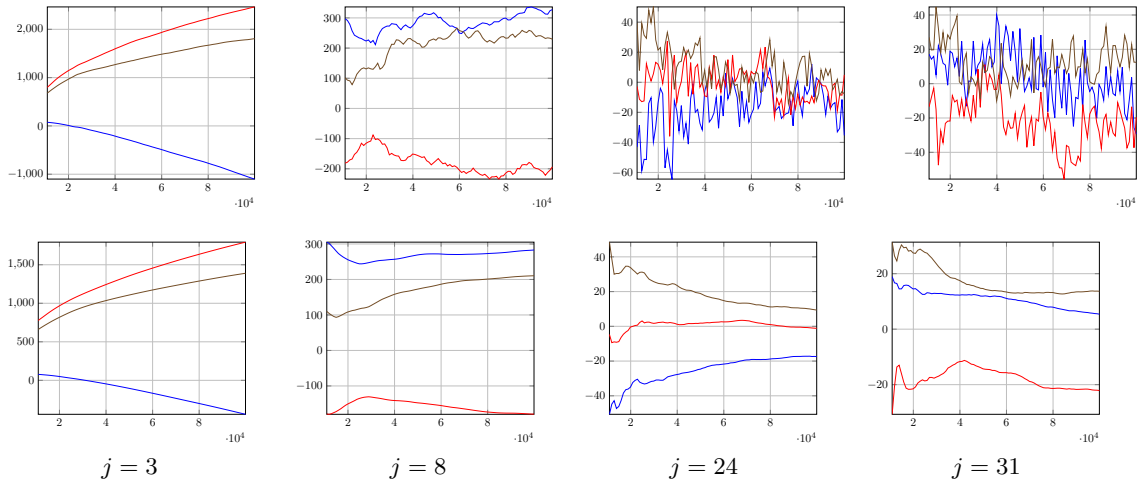


Figure 12: *Non convergence of the weights.* For the layers corresponding to $j = 3, 8, 24$ and 31 we study, on three channels ($k = 10, 20$ and 30), the behavior of the sequence $(\theta_n(i_{k,j}))_{n \in \mathbb{N}}$ (first row) and the averaged sequence $(\bar{\theta}_n(i_{k,j}))_{n \in \mathbb{N}}$ (second row), where $i_{k,j}$ is the index corresponding to layer j and channel k . These sequences have not converged yet, although the averaged sequence seems to stabilize for some layers, *i.e.* some values of j .

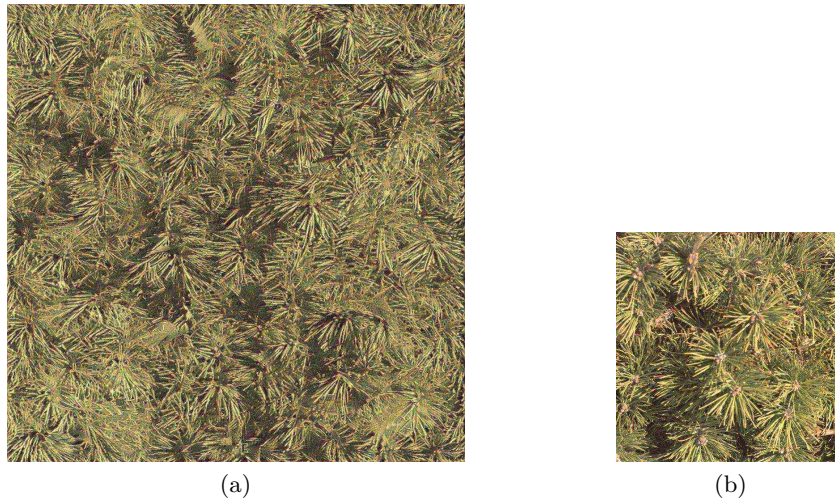


Figure 13: *Arbitrary size synthesis.* a texture of size 1024×1024 (a) is generated from an exemplar texture of size 512×512 (b). The hyperparameters are fixed as follows: $\delta_n = 10^{-3}$, $\gamma_n = 10^{-5}$ and $m_n = 1$.

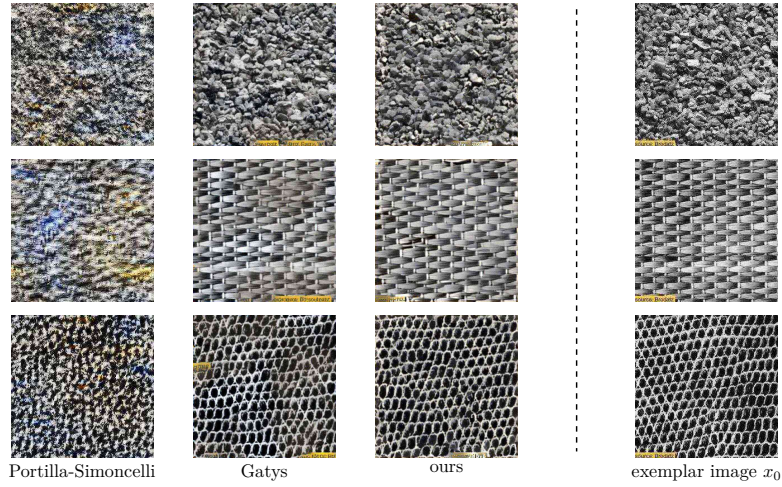


Figure 14: *Comparison with [77]*. The images presented in the column “Portilla-Simoncelli” are synthesized with the algorithm introduced in [77], the ones presented in the column “Gatys” are generated with [44] and the third column contains our results.

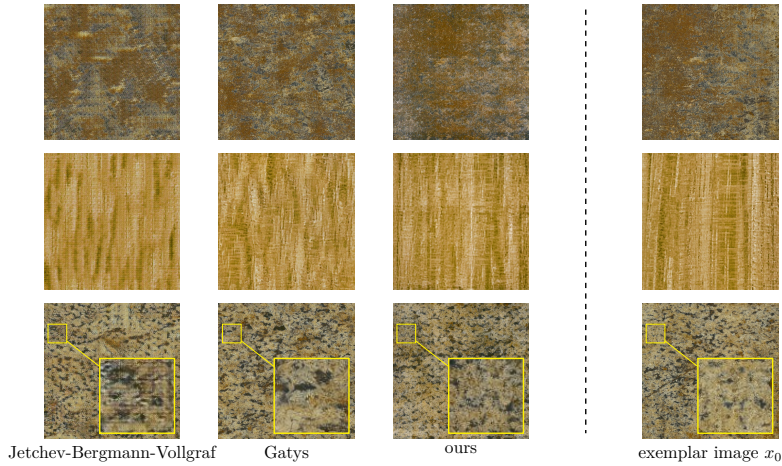


Figure 15: *Comparison with [54]*. The images presented in the column “Jetchev-Bergmann-Vollgraf” are synthesized with the algorithm introduced in [54], the ones presented in the column “Gatys” are generated with [44] and the third column contains our results.

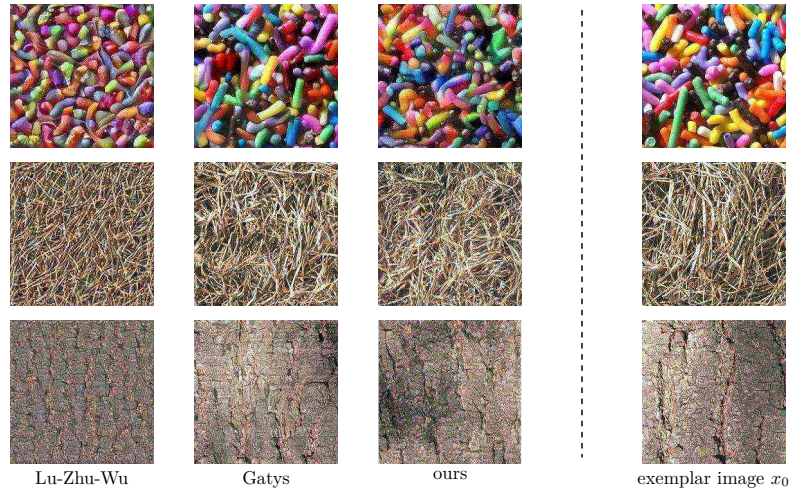


Figure 16: *Comparison with [65]*. The images presented in the column “Lu-Zhu-Wu” are synthesized with the algorithm introduced in [65], the ones presented in the column “Gatys” are generated with [44] and the third column contains our results.

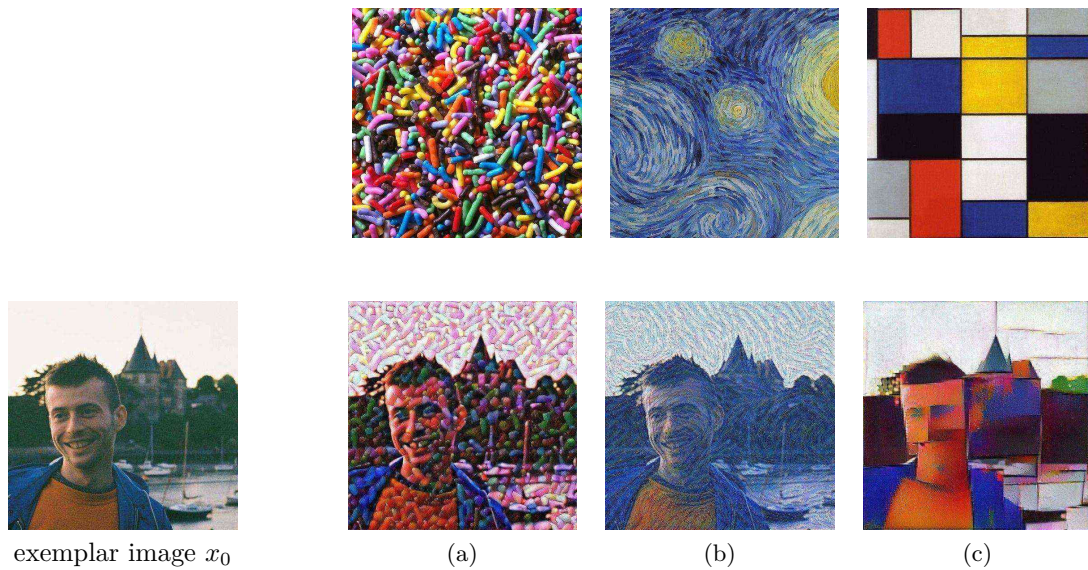


Figure 17: *Style transfer*. In (a), (b) and (c) we present the outputs of the SOUL algorithm with an exemplar content given in the leftmost column and exemplar style given by the first row. See Section 5.2.3 for more details.

References

- [1] S. AHN, A. K. BALAN, AND M. WELLING, *Bayesian posterior sampling via stochastic gradient Fisher scoring*, in Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012, 2012.
- [2] S.-I. AMARI AND H. NAGAOKA, *Methods of information geometry*, vol. 191 of Translations of Mathematical Monographs, American Mathematical Society, Providence, RI; Oxford University Press, Oxford, 2000. Translated from the 1993 Japanese original by Daishi Harada.
- [3] M. ARJOVSKY, S. CHINTALA, AND L. BOTTOU, *Wasserstein generative adversarial networks*, in Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, 2017, pp. 214–223.
- [4] Y. F. ATCHADÉ, G. FORT, AND E. MOULINES, *On perturbed proximal gradient algorithms*, J. Mach. Learn. Res., 18 (2017), pp. 10:1–10:33.
- [5] H. ATTOUCH AND Z. CHBANI, *Fast inertial dynamics and FISTA algorithms in convex optimization. perturbation aspects*, 2015, <https://arxiv.org/abs/1507.01367>.
- [6] J.-F. AUJOL, C. DOSSAL, G. FORT, AND É. MOULINES, *Rates of Convergence of Perturbed FISTA-based algorithms*. working paper or preprint, July 2019.
- [7] O. BARNDORFF-NIELSEN, *Information and exponential families in statistical theory*, Wiley Series in Probability and Statistics, John Wiley & Sons, Ltd., Chichester, 2014. Reprint of the 1978 original.
- [8] J. T. BARRON, *Continuously differentiable exponential linear units*, 2017, <https://arxiv.org/abs/1704.07483>.
- [9] U. BERGMANN, N. JETCHEV, AND R. VOLLGRAF, *Learning texture manifolds with the periodic spatial GAN*, in Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, 2017, pp. 469–477.
- [10] G. BIAU, B. CADRE, M. SANGNIER, AND U. TANIELIAN, *Some theoretical properties of GANs*, 2018, <https://arxiv.org/abs/1803.07819>.
- [11] W. M. BOOTHBY, *An introduction to differentiable manifolds and Riemannian geometry*, vol. 120 of Pure and Applied Mathematics, Academic Press, Inc., Orlando, FL, second ed., 1986.
- [12] H. BREZIS, *Functional analysis, Sobolev spaces and partial differential equations*, Universitext, Springer, New York, 2011.

- [13] J. BRUNA AND S. MALLAT, *Invariant scattering convolution networks*, IEEE transactions on pattern analysis and machine intelligence, 35 (2013), pp. 1872–1886.
- [14] J. BRUNA AND S. MALLAT, *Multiscale sparse microcanonical models*, 2018, <https://arxiv.org/abs/1801.02013>.
- [15] D. CANO AND T. MINH, *Texture synthesis using hierarchical linear transforms*, Signal Processing, (1988).
- [16] R. CHELLAPPA, S. CHATTERJEE, AND R. BAGDAZIAN, *Texture synthesis and compression using gaussian-markov random field models*, IEEE Trans. Systems, Man, and Cybernetics, 15 (1985), pp. 298–303.
- [17] H.-F. CHEN, L. GUO, AND A.-J. GAO, *Convergence and robustness of the robbins-monro algorithm truncated at randomly varying bounds*, Stochastic Processes and their Applications, 27 (1987), pp. 217–231.
- [18] I. CSISZÁR, *I-divergence geometry of probability distributions and minimization problems*, Ann. Probability, 3 (1975), pp. 146–158.
- [19] I. CSISZÁR, *Sanov property, generalized I-projection and a conditional limit theorem*, Ann. Probab., 12 (1984), pp. 768–793.
- [20] I. CSISZÁR, *Maxent, mathematics, and information theory*, in Maximum entropy and Bayesian methods (Santa Fe, NM, 1995), vol. 79 of Fund. Theories Phys., Kluwer Acad. Publ., Dordrecht, 1996, pp. 35–50.
- [21] I. CSISZÁR, F. GAMBOA, AND E. GASSIAT, *MEM pixel correlated solutions for generalized moment and interpolation problems*, IEEE Trans. Inform. Theory, 45 (1999), pp. 2253–2270.
- [22] A. S. DALALYAN, *Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent*, 2017, <https://arxiv.org/abs/1704.04752>.
- [23] A. S. DALALYAN, *Theoretical guarantees for approximate sampling from smooth and log-concave densities*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 79 (2017), pp. 651–676.
- [24] V. DE BORTOLI, A. DESOLNEUX, B. GALERNE, AND A. LECLAIRE, *Macrocanonical models for texture synthesis*, in Scale Space and Variational Methods in Computer Vision, J. Lellmann, M. Burger, and J. Modersitzki, eds., Cham, 2019, Springer International Publishing, pp. 13–24.

- [25] V. DE BORTOLI AND A. DURMUS, *Convergence of diffusions and their discretizations: from continuous to discrete processes and back*, 2019, <https://arxiv.org/abs/1904.09808>.
- [26] V. DE BORTOLI, A. DURMUS, M. PEREYRA, AND A. F. VIDAL, *Efficient stochastic optimisation by unadjusted langevin monte carlo. application to maximum marginal likelihood and empirical bayesian estimation*, 2019, <https://arxiv.org/abs/1906.12281>.
- [27] M. D. DONSKER AND S. R. S. VARADHAN, *Asymptotic evaluation of certain Markov process expectations for large time. I*, *Comm. Pure Appl. Math.*, 28 (1975), pp. 1–47; *ibid.* 28 (1975), 279–301.
- [28] M. D. DONSKER AND S. R. S. VARADHAN, *Asymptotic evaluation of certain Markov process expectations for large time. III*, *Comm. Pure Appl. Math.*, 29 (1976), pp. 389–461.
- [29] R. DOUC, E. MOULINES, P. PRIOURET, AND P. SOULIER, *Markov chains*, Springer Series in Operations Research and Financial Engineering, Springer, Cham, 2018.
- [30] R. DUITS, M. FELSBURG, G. GRANLUND, AND B. TER HAAR ROMENY, *Image analysis and reconstruction using a wavelet transform constructed from a reducible representation of the Euclidean motion group*, *International Journal of Computer Vision*, 72 (2007), pp. 79–102.
- [31] P. DUPUIS AND R. S. ELLIS, *A weak convergence approach to the theory of large deviations*, Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons, Inc., New York, 1997. A Wiley-Interscience Publication.
- [32] A. DURMUS AND E. MOULINES, *Nonasymptotic convergence analysis for the unadjusted Langevin algorithm*, *Ann. Appl. Probab.*, 27 (2017), pp. 1551–1587.
- [33] A. DURMUS, G. O. ROBERTS, G. VILMART, K. C. ZYGALAKIS, ET AL., *Fast langevin based algorithm for mcmc in high dimensions*, *The Annals of Applied Probability*, 27 (2017), pp. 2195–2237.
- [34] A. EBERLE, *Reflection couplings and contraction rates for diffusions*, *Probab. Theory Related Fields*, 166 (2016), pp. 851–886.
- [35] A. A. EFROS AND W. T. FREEMAN, *Image quilting for texture synthesis and transfer*, in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2001*, Los Angeles, California, USA, August 12-17, 2001, 2001, pp. 341–346.

- [36] A. A. EFROS AND T. K. LEUNG, *Texture synthesis by non-parametric sampling*, in ICCV, 1999, pp. 1033–1038.
- [37] G. FORT, L. RISSER, Y. ATCHADÉ, AND E. MOULINES, *Stochastic FISTA algorithms: So fast?*, in 2018 IEEE Statistical Signal Processing Workshop (SSP), IEEE, 2018, pp. 796–800.
- [38] A. FOURNIER, D. S. FUSSELL, AND L. C. CARPENTER, *Computer rendering of stochastic models*, Commun. ACM, 25 (1982), pp. 371–384.
- [39] A. GAGALOWICZ AND S. D. MA, *Model driven synthesis of natural textures for 3-d scenes*, Computers & Graphics, (1986).
- [40] B. GALERNE, Y. GOUSSEAU, AND J. MOREL, *Random phase textures: Theory and synthesis*, IEEE Trans. Image Processing, 20 (2011), pp. 257–267.
- [41] B. GALERNE AND A. LECLAIRE, *Texture inpainting using efficient gaussian conditional simulation*, SIAM Journal on Imaging Sciences, 10 (2017), pp. 1446–1474.
- [42] B. GALERNE, A. LECLAIRE, AND L. MOISAN, *A texton for fast and flexible gaussian texture synthesis*, in 22nd European Signal Processing Conference, EUSIPCO 2014, Lisbon, Portugal, September 1-5, 2014, 2014, pp. 1686–1690.
- [43] B. GALERNE, A. LECLAIRE, AND J. RABIN, *A texture synthesis model based on semi-discrete optimal transport in patch space*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 2456–2493.
- [44] L. A. GATYS, A. S. ECKER, AND M. BETHGE, *Texture synthesis using convolutional neural networks*, in Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, 2015, pp. 262–270.
- [45] N. GONTHIER, Y. GOUSSEAU, AND S. LADJAL, *High resolution neural texture synthesis with long range constraints*. working paper or preprint, 2019.
- [46] I. J. GOODFELLOW, J. POUGET-ABADIE, M. MIRZA, B. XU, D. WARDE-FARLEY, S. OZAIR, A. C. COURVILLE, AND Y. BENGIO, *Generative adversarial nets*, in Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, 2014, pp. 2672–2680.
- [47] J. HAN, K. ZHOU, L. WEI, M. GONG, H. BAO, X. ZHANG, AND B. GUO, *Fast example-based surface texture synthesis via discrete optimization*, The Visual Computer, 22 (2006), pp. 918–925.

- [48] W. K. HASTINGS, *Monte Carlo sampling methods using Markov chains and their applications*, *Biometrika*, 57 (1970), pp. 97–109.
- [49] D. J. HEEGER AND J. R. BERGEN, *Pyramid-based texture analysis/synthesis*, in *ICIP*, 1995.
- [50] C.-R. HWANG, *Laplace’s method revisited: weak convergence of probability measures*, *Ann. Probab.*, 8 (1980), pp. 1177–1182.
- [51] P. ISHWAR AND P. MOULIN, *On the existence and characterization of the maxent distribution under general moment inequality constraints*, *IEEE Trans. Inform. Theory*, 51 (2005), pp. 3322–3333.
- [52] L. ISSERLIS, *On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables*, *Biometrika*, 12 (1918), pp. 134–139.
- [53] E. T. JAYNES, *Information theory and statistical mechanics*, *Phys. Rev.*, (1957).
- [54] N. JETCHEV, U. BERGMANN, AND R. VOLLGRAF, *Texture synthesis with spatial generative adversarial networks*, 2016, <https://arxiv.org/abs/1611.08207>.
- [55] J. JOHNSON, A. ALAHI, AND L. FEI-FEI, *Perceptual losses for real-time style transfer and super-resolution*, in *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II*, 2016, pp. 694–711.
- [56] P. JUPP AND K. MARDIA, *A note on the maximum-entropy principle*, *Scandinavian Journal of Statistics*, (1983), pp. 45–47.
- [57] O. KALLENBERG, *Foundations of modern probability*, Springer Science & Business Media, 2006.
- [58] A. KASPAR, B. NEUBERT, D. LISCHINSKI, M. PAULY, AND J. KOPF, *Self tuning texture optimization*, *Comput. Graph. Forum*, 34 (2015), pp. 349–359.
- [59] H. J. KUSHNER AND G. G. YIN, *Stochastic approximation and recursive algorithms and applications*, vol. 35 of *Applications of Mathematics (New York)*, Springer-Verlag, New York, second ed., 2003. *Stochastic Modelling and Applied Probability*.
- [60] V. KWATRA, I. A. ESSA, A. F. BOBICK, AND N. KWATRA, *Texture optimization for example-based synthesis*, *ACM Trans. Graph.*, 24 (2005), pp. 795–802.
- [61] V. KWATRA, A. SCHÖDL, I. A. ESSA, G. TURK, AND A. F. BOBICK, *Graphcut textures: image and video synthesis using graph cuts*, *ACM Trans. Graph.*, 22 (2003), pp. 277–286.

- [62] E. LEVINA AND P. J. BICKEL, *Texture synthesis and nonparametric resampling of random fields*, Ann. Statist., 34 (2006), pp. 1751–1773.
- [63] C. LI AND M. WAND, *Precomputed real-time texture synthesis with markovian generative adversarial networks*, in Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III, 2016, pp. 702–716.
- [64] G. LIU, Y. GOUSSEAU, AND G. XIA, *Texture synthesis through convolutional neural networks and spectrum constraints*, in ICPR, 2016.
- [65] Y. LU, S. ZHU, AND Y. N. WU, *Learning FRAME models using CNN filters*, in Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 2016, pp. 1902–1910.
- [66] Y.-A. MA, T. CHEN, AND E. FOX, *A complete recipe for stochastic gradient mcmc*, in Advances in Neural Information Processing Systems, 2015, pp. 2917–2925.
- [67] S. MALLAT, *Group invariant scattering*, Communications on Pure and Applied Mathematics, 65 (2012), pp. 1331–1398.
- [68] S. P. MEYN AND R. L. TWEEDIE, *Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes*, Adv. in Appl. Probab., 25 (1993), pp. 518–548, <https://doi.org/10.2307/1427522>, <https://doi.org/10.2307/1427522>.
- [69] D. MUMFORD AND A. DESOLNEUX, *Pattern theory: The stochastic analysis of real-world signals*, Applying Mathematics, A K Peters, Ltd., Natick, MA, 2010.
- [70] Y. NESTEROV, *Introductory lectures on convex optimization*, vol. 87 of Applied Optimization, Kluwer Academic Publishers, Boston, MA, 2004.
- [71] Y. E. NESTEROV, *A method for solving the convex programming problem with convergence rate $O(1/k^2)$* , Dokl. Akad. Nauk SSSR, 269 (1983), pp. 543–547.
- [72] H. E. OGDEN, *A sequential reduction method for inference in generalized linear mixed models*, Electron. J. Stat., 9 (2015), pp. 135–152.
- [73] S. PATTERSON AND Y. W. TEH, *Stochastic gradient riemannian langevin dynamics on the probability simplex*, in Advances in neural information processing systems, 2013, pp. 3102–3110.
- [74] K. PERLIN, *An image synthesizer*, in Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1985, San Francisco, California, USA, July 22-26, 1985, 1985, pp. 287–296.

- [75] G. PEYRÉ, *Texture synthesis with grouplets*, IEEE Trans. Pattern Anal. Mach. Intell., (2010).
- [76] G. PÓLYA AND G. SZEGŐ, *Problems and theorems in analysis. I*, Classics in Mathematics, Springer-Verlag, Berlin, 1998. Series, integral calculus, theory of functions, Translated from the German by Dorothee Aeppli, Reprint of the 1978 English translation.
- [77] J. PORTILLA AND E. P. SIMONCELLI, *A parametric texture model based on joint statistics of complex wavelet coefficients*, IJCV, (2000).
- [78] S. R. PURKS AND W. RICHARDS, *Visual texture discrimination using random-dot patterns*, J. Opt. Soc. Am., 67 (1977), pp. 765–771.
- [79] L. RAAD, A. DAVY, A. DESOLNEUX, AND J.-M. MOREL, *A survey of exemplar-based texture synthesis*, 2017, <https://arxiv.org/abs/1707.07184>.
- [80] L. RAAD, A. DESOLNEUX, AND J. MOREL, *A conditional multiscale locally gaussian texture synthesis algorithm*, Journal of Mathematical Imaging and Vision, 56 (2016), pp. 260–279.
- [81] H. ROBBINS AND S. MONRO, *A stochastic approximation method*, The annals of mathematical statistics, (1951), pp. 400–407.
- [82] G. O. ROBERTS AND R. L. TWEEDIE, *Exponential convergence of Langevin distributions and their discrete approximations*, Bernoulli, 2 (1996), pp. 341–363.
- [83] W. RUDIN, *Real and complex analysis*, Tata McGraw-hill education, 2006.
- [84] J. SCHELLDORFER, L. MEIER, AND P. BÜHLMANN, *GLMMLasso: an algorithm for high-dimensional generalized linear mixed models using ℓ_1 -penalization*, J. Comput. Graph. Statist., 23 (2014), pp. 460–477.
- [85] K. SIMONYAN AND A. ZISSERMAN, *Very deep convolutional networks for large-scale image recognition*, 2014, <https://arxiv.org/abs/1409.1556>.
- [86] U. SIMSEKLI, R. BADEAU, T. CEMGIL, AND G. RICHARD, *Stochastic quasi-newton langevin monte carlo*, in International Conference on Machine Learning (ICML), 2016.
- [87] F. TOPSØ E, *Information-theoretical optimization techniques*, Kybernetika (Prague), 15 (1979), pp. 8–27.
- [88] D. ULYANOV, V. LEBEDEV, A. VEDALDI, AND V. S. LEMPITSKY, *Texture networks: Feed-forward synthesis of textures and stylized images*, in Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016, 2016, pp. 1349–1357.

- [89] D. ULYANOV, A. VEDALDI, AND V. S. LEMPITSKY, *Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis*, in 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, 2017, pp. 4105–4113.
- [90] I. USTYUZHANINOV, W. BRENDEL, L. A. GATYS, AND M. BETHGE, *Texture synthesis using shallow convolutional networks with random filters*, 2016, <https://arxiv.org/abs/1606.00021>.
- [91] J. J. VAN WIJK, *Spot noise texture synthesis for data visualization*, in SIGGRAPH, 1991, pp. 309–318.
- [92] A. F. VIDAL, V. D. BORTOLI, M. PEREYRA, AND A. DURMUS, *Maximum likelihood estimation of regularisation parameters in high-dimensional inverse problems: an empirical bayesian approach*, 2019, <https://arxiv.org/abs/1911.11709>.
- [93] M. WELLING AND Y. W. TEH, *Bayesian learning via stochastic gradient langevin dynamics*, in Proceedings of the 28th international conference on machine learning (ICML-11), 2011, pp. 681–688.
- [94] D. WILLIAMS, *Probability with martingales*, Cambridge Mathematical Textbooks, Cambridge University Press, Cambridge, 1991.
- [95] Y. ZHOU, Z. ZHU, X. BAI, D. LISCHINSKI, D. COHEN-OR, AND H. HUANG, *Non-stationary texture synthesis by adversarial expansion*, ACM Trans. Graph., 37 (2018), pp. 49:1–49:13.
- [96] S. C. ZHU, Y. N. WU, AND D. MUMFORD, *Filters, random fields and maximum entropy (FRAME): towards a unified theory for texture modeling*, International Journal of Computer Vision, 27 (1998), pp. 107–126.

A Proofs of Section 3

We have the following variational formula which is an extension of [31, Proposition 1.4.2] to the case where F is not bounded. More precisely, allowing some growth on F , controlled by a parameter α , and restricting the set of probability measures we consider to \mathcal{P}_α we obtain the same equality. The proof is almost identical but is given for completeness.

Proposition 12. *Assume **A1**(α) with $\alpha > 0$. Then, for any $\theta \in \Theta_F$, with Θ_F defined by (3),*

$$\inf_{\pi \in \mathcal{P}_\alpha} \{ \text{KL}(\pi | \mu) + \langle \theta, \pi(F) \rangle \} = - \log \left\{ \int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) \right\} .$$

Proof. Let $\theta \in \Theta_F$ and $\pi \in \mathcal{P}_\alpha$. Note that under **A1**(α), $\pi(\|F\|) < +\infty$ and $\pi(F)$ is well defined.

If $\text{KL}(\pi|\mu) = +\infty$, then $\text{KL}(\pi|\mu) + \langle \theta, \pi(F) \rangle = +\infty$. Consider now the case $\text{KL}(\pi|\mu) < +\infty$. By definition of Θ_F , we can therefore consider π_θ , the probability measure with density with respect to μ given for any $x \in \mathbb{R}^d$ by

$$\frac{d\pi_\theta}{d\mu}(x) = \exp[-\langle \theta, F(x) \rangle] / \int_{\mathbb{R}^d} \exp[-\langle \theta, F(y) \rangle] d\mu(y) .$$

Note that since μ -almost everywhere, $(d\pi_\theta)/(d\mu)(x) > 0$, μ and π_θ are equivalent. Since $\text{KL}(\pi|\mu) < +\infty$, $\pi \ll \mu$ which implies in turn $\pi \ll \pi_\theta$ and we have

$$\begin{aligned} \text{KL}(\pi|\mu) &= \text{KL}(\pi|\pi_\theta) + \int_{\mathbb{R}^d} \log \left(\frac{d\pi_\theta}{d\mu}(x) \right) d\pi(x) \\ &= \text{KL}(\pi|\pi_\theta) - \langle \theta, \pi(F) \rangle - \log \left\{ \int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) \right\} , \end{aligned}$$

which concludes the proof, since $\text{KL}(\pi|\pi_\theta) \geq 0$. \square

A.1 Proof of Proposition 3

The proof is divided in two parts:

(a) Assume that there exists π^* , solution of (P). Let \mathcal{E} be the convex set defined by $\{\pi \in \mathcal{P}_\alpha^F : \frac{d\pi}{d\pi^*}(x) \leq 2 \text{ for } \pi^* \text{ almost every } x\}$. For any $\pi_1 \in \mathcal{E}$, consider π_2 with density with respect to π^* , $\frac{d\pi_2}{d\pi^*} = 2 - \frac{d\pi_1}{d\pi^*}$ which by definition is an element of \mathcal{E} and $\pi^* = (\pi_1 + \pi_2)/2$. Hence π^* is an algebraic inner point of \mathcal{E} . Therefore using the equality case in [18, Theorem 2.2] we obtain that for any $\pi \in \mathcal{E}$, $\text{KL}(\pi|\mu) = \text{KL}(\pi|\pi^*) + \text{KL}(\pi^*|\mu)$. Using that for any $\pi \in \mathcal{E}$, we have $\text{KL}(\pi|\pi^*)$ and $\text{KL}(\pi^*|\mu) < +\infty$, we get that

$$\begin{aligned} 0 &= \int_{\mathbb{R}^d} \log \left(\frac{d\pi}{d\mu}(x) \right) d\pi(x) - \int_{\mathbb{R}^d} \log \left(\frac{d\pi}{d\pi^*}(x) \right) d\pi(x) - \int_{\mathbb{R}^d} \log \left(\frac{d\pi^*}{d\mu}(x) \right) d\pi^*(x) \quad (20) \\ &= \int_{\mathbb{R}^d} \log \left(\frac{d\pi^*}{d\mu}(x) \right) \left[\frac{d\pi}{d\pi^*}(x) - 1 \right] d\pi^*(x) . \end{aligned}$$

Since, $\text{KL}(\pi^*|\mu) < +\infty$ we have that $\log(\frac{d\pi^*}{d\mu}) \in L^1(\mathbb{R}^d, \pi^*)$. Let $V = \text{span}\{1, \langle F, e_i \rangle : i = 1, \dots, p\}$ where $(e_i)_{i \in \{1, \dots, p\}}$ is the canonical basis of \mathbb{R}^p . V is a finite dimensional (hence closed) subspace of $L^1(\mathbb{R}^d, \pi^*)$. Hence, in order to show that $\log(\frac{d\pi^*}{d\mu}) \in V$ it suffices to show that $\log(\frac{d\pi^*}{d\mu}) \in (V^\perp)^\perp$ by [12, Proposition II.12].

We identify the topological dual space of $L^1(\mathbb{R}^d, \pi^*)$ and $L^\infty(\mathbb{R}^d, \pi^*)$, see [83, Theorem 6.16]. Let $h \in L^\infty(\mathbb{R}^d, \pi^*) \cap V^\perp$. Then by definition, $\int_{\mathbb{R}^d} F(x)h(x)d\pi^*(x) = 0$ and $\int_{\mathbb{R}^d} h(x)d\pi^*(x) = 0$. The same goes for $\tilde{h} = h/\|h\|_\infty$, and we have that π_h defined by $\frac{d\pi_h}{d\pi^*} = 1 + \tilde{h}$ is an element of \mathcal{E} . Therefore, by (20), we get that $\int_{\mathbb{R}^d} \log(\frac{d\pi^*}{d\mu}(x))h(x) = 0$

and $\log(\frac{d\pi^*}{d\mu}) \in V$. More precisely, there exists $\theta^* \in \mathbb{R}^p$, $C \in \mathbb{R}$ and $\mathbf{N} \in \mathcal{B}(\mathbb{R}^d)$ with $\pi^*(\mathbf{N}) = 0$ such that for μ almost any $x \in \mathbb{R}^d \setminus \mathbf{N}$,

$$\log\left(\frac{d\pi^*}{d\mu}(x)\right) = \langle \theta^*, F(x) \rangle + C.$$

We also have that $\pi^*(\mathbf{N}) = \int_{\mathbf{N}} \frac{d\pi^*}{d\mu}(y) d\mu(y)$ and therefore for μ almost any $x \in \mathbf{N}$, $\frac{d\pi^*}{d\mu}(x) = 0$. Using [18, Remark 2.14], for any $\pi \in \mathcal{P}_\alpha^F$ such that $\text{KL}(\pi|\mu) < +\infty$ we have $\pi \ll \pi^*$ and therefore

$$\pi(\mathbf{N}) = 0, \tag{21}$$

which concludes the proof for (7).

Finally, if there exists $\pi \in \mathcal{P}_\alpha^F$ with $\text{KL}(\pi|\mu) < +\infty$ such that $\mu \ll \pi$ then by (21), $\mu(\mathbf{N}) = \pi(\mathbf{N}) = 0$ and we get that $\frac{d\pi^*}{d\mu}(x) = \exp[-\langle \theta^*, F(x) \rangle] / \int_{\mathbb{R}^d} \exp[-\langle \theta^*, F(y) \rangle] d\mu(y)$ for μ almost every $x \in \mathbb{R}^d$. Then, using Proposition 12 and that $\theta^* \in \Theta_F$, we have by definition of (Q), see (4),

$$\begin{aligned} v(\mathbf{Q}) \leq v(\mathbf{P}) = \text{KL}(\pi^*|\mu) &= -\log\left(\int_{\mathbb{R}^d} \exp[-\langle \theta^*, F(x) \rangle] d\mu(x)\right) \\ &= \inf_{\pi \in \mathcal{P}_\alpha^F} \{\text{KL}(\pi|\mu) - \langle \theta^*, \pi(F) \rangle\} \leq v(\mathbf{Q}), \end{aligned}$$

which concludes the first part of the proof.

(b) If there exists π^* solution of (P) with $v(\mathbf{P}) < +\infty$ then $\text{KL}(\pi^*|\mu) < +\infty$ and $\pi^* \in \mathcal{P}_\alpha^F$. Now, assume that there exists $\pi \in \mathcal{P}_\alpha^F$ such that $\text{KL}(\pi|\mu) < +\infty$. Let $(\pi_n)_{n \in \mathbb{N}}$ be a sequence of probability measures such that for any $n \in \mathbb{N}$, $\pi_n \in \mathcal{P}_\alpha^F$, $\text{KL}(\pi_n|\mu) < +\infty$ and $\inf_{\mathcal{P}_\alpha^F} \text{KL}(\pi|\mu) = \lim_{n \rightarrow +\infty} \text{KL}(\pi_n|\mu)$. Using [28, Lemma 5.1] we get that $(\pi_n)_{n \in \mathbb{N}}$ is tight. Therefore we can assume, up to extraction, that $(\pi_n)_{n \in \mathbb{N}}$ converges to some probability measure π^* for the weak topology. Since $\pi \mapsto \text{KL}(\pi|\mu)$ is lower semi-continuous [31, Lemma 1.4.3 (b)] we obtain that $\text{KL}(\pi^*|\mu) = \inf_{\mathcal{P}_\alpha^F} \text{KL}(\pi|\mu)$. We recall the Donsker-Varadhan variational formula [27, Lemma 2.1] stating that for any continuous, real-valued and bounded mapping ϕ we have for any $n \in \mathbb{N}$

$$\int_{\mathbb{R}^d} \phi(x) d\pi_n(x) \leq \text{KL}(\pi_n|\mu) + \log\left(\int_{\mathbb{R}^d} e^{\phi(x)} d\mu(x)\right). \tag{22}$$

Let $\varphi_M : \mathbb{R}^d \rightarrow \mathbb{R}$ defined for any $M \geq 0$ and $x \in \mathbb{R}^d$ by $\varphi_M(x) = \eta \max(\|x\|^{\alpha'}, M)$, with η defined in **A2**(α'). Using (22), **A2**(α') and that φ_M is continuous and bounded we get that for any $n \in \mathbb{N}$ and $M \geq 0$

$$\int_{\mathbb{R}^d} \varphi_M(x) d\pi_n(x) \leq \sup_{n \in \mathbb{N}} \text{KL}(\pi_n|\mu) + \log\left(\int_{\mathbb{R}^d} \exp[\eta \|x\|^{\alpha'}] d\mu(x)\right) < +\infty.$$

Using the monotone convergence theorem we get that $\sup_{n \in \mathbb{N}} \int_{\mathbb{R}^d} \|x\|^{\alpha'} d\pi_n(x) < +\infty$. By [57, Theorem 5.16] there exist $(X_n)_{n \in \mathbb{N}}$ a sequence of \mathbb{R}^d random variables and X a \mathbb{R}^d random variable such that for any $n \in \mathbb{N}$, X_n is distributed according to π_n and X is distributed according to π^* . Since $(\pi_n)_{n \in \mathbb{N}}$ converges weakly to π^* , $(X_n)_{n \in \mathbb{N}}$ converges in distribution towards X . Therefore, we get that $(\|X_n\|^\alpha)_{n \in \mathbb{N}}$ converges in distribution to $\|X\|^\alpha$ and $\sup_{n \in \mathbb{N}} \mathbb{E}[\|X\|^{\alpha'}] < \infty$. By [57, Lemma 3.11], we get that $\mathbb{E}[\|X\|^{\alpha'}] = \int_{\mathbb{R}^d} \|x\|^{\alpha'} d\pi^*(x) < +\infty$. Hence, $\pi^* \in \mathcal{P}_\alpha$. In addition, since F is continuous by **A1**, we have that $(F(X_n))_{n \in \mathbb{N}}$ converges in distribution to $F(X)$. By [94, 13.3] and **A1**(α), we have that $(F(X_n))_{n \in \mathbb{N}}$ is uniformly integrable. Using [57, Lemma 3.11] and that for any $n \in \mathbb{N}$, $\pi_n(F) = 0$, we get $\lim_{n \rightarrow +\infty} \pi_n(F) = \pi^*(F) = 0$ and $\pi^* \in \mathcal{P}_\alpha^F$, which concludes the proof.

A.2 Proof of Proposition 4

Let $L : \text{int}(\Theta_F) \rightarrow \mathbb{R}$ be the function defined for any $\theta \in \text{int}(\Theta_F)$ by

$$L(\theta) = \log \left\{ \int_{\mathbb{R}^d} \exp[-\langle \theta, F(x) \rangle] d\mu(x) \right\} .$$

We have $L \in C^\infty(\text{int}(\Theta_F))$. The proposition is trivial if $\text{int}(\Theta_F) = \emptyset$. Therefore we suppose that $\text{int}(\Theta_F) \neq \emptyset$ and let $\theta_0 \in \text{int}(\Theta_F)$. Since $\text{int}(\Theta_F)$ is open, there exists $a_1 > 1$ such that $a_1\theta_0 \in \text{int}(\Theta_F)$. Let $a_2 > 1$ such that $1/a_1 + 1/a_2 = 1$. Let $R = \eta/(2a_2)$ with η given in **A2**(α). For any $\theta \in \mathring{B}(\theta_0, R)$, using that $t^j e^{-t} \leq j^j$ for $t \geq 0$ and $j \in \mathbb{N}$, we have for any $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$

$$\begin{aligned} \|F(x)\|^k \exp[-\langle \theta, F(x) \rangle] &\leq (k/R)^k \exp[R \|F(x)\|] \exp[-\langle \theta, F(x) \rangle] \\ &\leq (k/R)^k \exp[2R \|F(x)\|] \exp[-\langle \theta_0, F(x) \rangle] , \end{aligned}$$

The last quantity is independent of θ and μ -integrable using Hölder's inequality, since

$$\begin{aligned} &\int_{\mathbb{R}^d} \exp[2R \|F(x)\|] \exp[-\langle \theta_0, F(x) \rangle] d\mu(x) \\ &\leq \left(\int_{\mathbb{R}^d} \exp[\eta \|x\|^\alpha] d\mu(x) \right)^{1/a_2} \left(\int_{\mathbb{R}^d} \exp[-\langle a_1\theta_0, F(x) \rangle] d\mu(x) \right)^{1/a_1} < +\infty . \end{aligned}$$

This result implies that $L \in C^\infty(\text{int}(\Theta_F))$. Therefore, if θ^* is a stationary point, we have

$$\pi_{\theta^*}(F) \propto \int_{\mathbb{R}^d} F(x) \exp[-\langle \theta^*, F(x) \rangle] d\mu(x) = 0 .$$

Since $\pi_{\theta^*} \in \mathcal{P}_\alpha$ we have $\pi_{\theta^*} \in \mathcal{P}_\alpha^F$. Since $\mu \ll \pi_{\theta^*}$ we have $\pi \ll \pi_{\theta^*}$ for any $\pi \ll \mu$. Therefore for any $\pi \in \mathcal{P}_\alpha^F$ with $\pi \ll \mu$ we have

$$\text{KL}(\pi|\mu) = \int_{\mathbb{R}^d} \log \left(\frac{d\pi}{d\mu}(x) \right) d\pi(x) = \text{KL}(\pi|\pi_{\theta^*}) - L(\theta^*) = \text{KL}(\pi|\pi_{\theta^*}) + \text{KL}(\pi_{\theta^*}|\mu) \geq \text{KL}(\pi_{\theta^*}|\mu) .$$

If π is not absolutely continuous with respect to μ then $\text{KL}(\pi|\mu) = +\infty$. Therefore we have that π_{θ^*} solves (P).

A.3 Proof of Proposition 5

We preface the proof with the following lemma

Lemma 13. *Let $h : \mathbb{R}^p \rightarrow \mathbb{R}$ be a convex function such that h is ray-coercive, i.e. for any $\theta \in \mathbb{R}^p$, with $\|\theta\| = 1$ we have $\lim_{t \rightarrow +\infty} h(t\theta) = +\infty$. Then h is coercive, i.e. $\lim_{\|\theta\| \rightarrow +\infty} h(\theta) = +\infty$.*

Proof. Assume that h is not coercive. Then there exists a sequence $(\theta_n)_{n \in \mathbb{N}} \in (\mathbb{R}^d \setminus \{0\})^{\mathbb{N}}$ such that $\lim_{n \rightarrow +\infty} \|\theta_n\| = +\infty$ and the sequence $(h(\theta_n))_{n \in \mathbb{N}}$ is bounded. Upon extraction we can assume that $\lim_{n \rightarrow +\infty} \theta_n / \|\theta_n\| = \tilde{\theta}$. We have for any $t \in \mathbb{R}$

$$h(t\tilde{\theta}) = h(t\tilde{\theta}) - h(t\theta_n / \|\theta_n\|) + h(t\theta_n / \|\theta_n\|). \quad (23)$$

Let $t \geq 0$ and $\varepsilon > 0$. Since h is convex, h is continuous and there exists $n_0 \in \mathbb{N}$ such that for any $n \in \mathbb{N}$ with $n \geq n_0$,

$$\left\| h(t\tilde{\theta}) - h(t\theta_n / \|\theta_n\|) \right\| \leq \varepsilon, \quad h(t\theta_n / \|\theta_n\|) \leq t / \|\theta_n\| h(\theta_n) + (1 - t / \|\theta_n\|) h(0). \quad (24)$$

Combining (23) and (24) we obtain that for any $t \geq 0$,

$$h(t\tilde{\theta}) \leq \varepsilon + \|h(0)\| + \sup_{n \in \mathbb{N}} \|h(\theta_n)\| < +\infty,$$

which is absurd. Hence, h is coercive. □

We now turn to the proof of Proposition 5. We divide the proof in two parts.

(a) Using that $\Theta_F = \mathbb{R}^p$ and the first part of Proposition 4 we have that $L : \Theta_F \rightarrow \mathbb{R}$ is continuously differentiable over \mathbb{R}^p . In addition, for any $\theta_1, \theta_2 \in \mathbb{R}^p$ and $\theta \in (0, 1)$ we have

$$\begin{aligned} \int_{\mathbb{R}^d} \exp[-\langle \theta\theta_1 + (1 - \theta)\theta_2, F(x) \rangle] d\mu(x) \\ = \int_{\mathbb{R}^d} \exp[-\theta\langle \theta_1, F(x) \rangle] \exp[(1 - \theta)\langle \theta_2, F(x) \rangle] d\mu(x). \end{aligned}$$

Applying Hölder's inequality we get that

$$\begin{aligned} L(\theta\theta_1 + (1 - \theta)\theta_2) &= \log \left\{ \int_{\mathbb{R}^d} \exp[-\langle \theta\theta_1 + (1 - \theta)\theta_2, F(x) \rangle] d\mu(x) \right\} \\ &\leq \theta \log \left\{ \int_{\mathbb{R}^d} \exp[-\langle \theta_1, F(x) \rangle] d\mu(x) \right\} + (1 - \theta) \log \left\{ \int_{\mathbb{R}^d} \exp[\langle \theta_2, F(x) \rangle] d\mu(x) \right\}, \end{aligned}$$

hence L is convex. Using the monotone convergence theorem and (8) we have that for any $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$,

$$\begin{aligned} \lim_{t \rightarrow +\infty} L(t\theta) &= \lim_{t \rightarrow +\infty} \log \left\{ \int_{\mathbb{R}^d} \exp[-t\langle \theta, F(x) \rangle] d\mu(x) \right\} \\ &\geq \lim_{t \rightarrow +\infty} \log \left\{ \int_{\{x \in \mathbb{R}^d : \langle \theta, F(x) \rangle < 0\}} \exp[-t\langle \theta, F(x) \rangle] d\mu(x) \right\} = \infty, \end{aligned}$$

which implies that L is ray-coercive. Combining this result, the fact that L is convex and Lemma 13 we get that L is coercive. Since L is continuous and coercive it admits a minimizer θ^* and therefore $\nabla L(\theta^*) = 0$. Applying the second part of Proposition 4 concludes the first part of the proof.

(b) Let $x \in F^{-1}(\{0\})$ such that $\det(DF(x)DF(x)^T) > 0$. We obtain that $\text{Ker}(DF(x)^T) = \text{Ker}(DF(x)DF(x)^T) = \{0\}$. Hence, $\text{rank}(dF(x)) = \text{rank}(dF(x)^T) = p$ and $dF(x)$ is surjective. Using the submersion theorem, there exists $G : U \rightarrow \mathbb{R}^d$ with U an open neighborhood of $0 \in \mathbb{R}^p$ such that for any $\zeta \in U$, $F(G(\zeta)) = \zeta$. Therefore, for any $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$ there exists ζ_θ such that $\langle \theta, F(G(\zeta_\theta)) \rangle = \langle \theta, \zeta_\theta \rangle < 0$. Hence, since F is continuous, there exists an open set V in \mathbb{R}^d such that for any $y \in V$, $\langle F(y), \theta \rangle < 0$. Combining this result with the fact that for any A open and $A \neq \emptyset$, $\mu(A) > 0$ we conclude the proof.

A.4 Proof of Proposition 6

Let $\theta \in \mathbb{R}^p$ such that $\|\theta\| = 1$ and such that $\mu(\{x \in \mathbb{R}^d : \langle F(x), \theta \rangle \leq 0\}) = 0$. Then, we have using the dominated convergence theorem

$$\lim_{t \rightarrow +\infty} \int_{\mathbb{R}^d} \exp[-t\langle \theta, F(x) \rangle] d\mu(x) = \lim_{t \rightarrow +\infty} \int_{\{x \in \mathbb{R}^d : \langle F(x), \theta \rangle > 0\}} \exp[-t\langle \theta, F(x) \rangle] d\mu(x) = 0. \quad (25)$$

Recall that **A1**(α) and **A2**(α') imply that $\Theta_F = \mathbb{R}^p$. Therefore, using (25), we have $v(Q) = -\inf_{\theta \in \mathbb{R}^p} L(\theta) = +\infty$. Since $v(P) \geq v(Q)$, see (4), we have $v(P) = +\infty$. Hence, for any $\pi \in \mathcal{P}_\alpha^F$, $\text{KL}(\pi|\mu) = +\infty$ and any $\pi \in \mathcal{P}_\alpha^F$ solves (P). We conclude upon remarking that $\delta_{x_0} \in \mathcal{P}_\alpha^F$.

A.5 Proof of Proposition 7

We start to show that for any $j \in \{1, \dots, M\}$ there exists $\varepsilon_j > 0$ such that for any $x \in \bar{B}(x_0, \varepsilon_j)$

$$\mathcal{G}_j(x) = A_{j,+} \dots A_{1,+}(x). \quad (26)$$

Namely, for any $j \in \{1, \dots, M\}$, \mathcal{G}_j is locally affine around x_0 . We have that $x_0 \in \mathbb{R}^d \setminus \mathbf{N}_1$ with $\mathbf{N}_1 = \bigcup_{k=1}^{c_1 \times n_1} \text{Ker}(e_{k,1}^T A_1)$, where $(e_{k,1})_{k \in \{1, \dots, c_1 \times n_1\}}$ is the canonical basis of $\mathbb{R}^{c_1 \times n_1}$ by (13). For any $k \in \{1, \dots, n_1 \times c_1\}$, $A_1(x_0)(k) > 0$ or $A_1(x_0)(k) < 0$. Therefore, since A_1

is continuous, there exists $\varepsilon_1 > 0$ such that for any $x \in \bar{B}(x_0, \varepsilon_1)$, $A_1(x)(k)A_1(x_0)(k) > 0$. Therefore, for any $k \in \{1, \dots, n_1 \times c_1\}$, $\text{sign}(\mathcal{G}_1(x)(k)) = \text{sign}(\mathcal{G}_1(x_0)(k))$ and

$$\mathcal{G}_1(x) = \varphi(A_1(x)) = D_1(x)A_1(x) = D_1(x_0)A_1(x),$$

where D_1 is given in (12). Now assume that (26) is true for $j \in \{1, \dots, \ell\}$ with $\ell \in \{1, \dots, M-1\}$. There exists $\varepsilon_\ell > 0$ such that for any $x \in \bar{B}(x_0, \varepsilon_\ell)$

$$\mathcal{G}_{\ell+1}(x) = \varphi(A_{\ell+1}\mathcal{G}_\ell(x)) = \varphi(A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x)).$$

Since $x_0 \in \mathbb{R}^d \setminus (\bigcup_{j=1}^{\ell+1} \mathbf{N}_j)$ with $\mathbf{N}_j = \bigcup_{k=1}^{c_j \times n_j} \text{Ker}(e_{k,j}^\top A_j A_{j-1,+} \dots A_{1,+})$ with $(e_{k,j})_{k \in \{1, \dots, c_j \times n_j\}}$ the canonical basis of $\mathbb{R}^{c_j \times n_j}$, for any $k \in \{1, \dots, c_\ell \times n_\ell\}$, we get $A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x_0)(k) > 0$ or we have $A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x_0)(k) < 0$. Therefore, since $A_{\ell+1}A_{\ell,+} \dots A_{1,+}$ is continuous there exists $\varepsilon_\ell > \varepsilon_{\ell+1} > 0$ such that for any $x \in \bar{B}(x_0, \varepsilon_{\ell+1})$ and $k \in \{1, \dots, c_{\ell+1} \times n_{\ell+1}\}$,

$$A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x)(k)A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x_0)(k) > 0.$$

Therefore, for any $x \in \bar{B}(x_0, \varepsilon_{\ell+1})$ and $k \in \{1, \dots, c_{\ell+1} \times n_{\ell+1}\}$ we have $\text{sign}(\mathcal{G}_{\ell+1}(x)(k)) = \text{sign}(\mathcal{G}_{\ell+1}(x_0)(k))$ for any $x \in \bar{B}(x_0, \varepsilon_{\ell+1})$, and

$$\begin{aligned} \mathcal{G}_{\ell+1}(x) &= \varphi(A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x)) = D_{\ell+1}(x)A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x) \\ &= D_{\ell+1}(x_0)A_{\ell+1}A_{\ell,+} \dots A_{1,+}(x) = A_{\ell+1,+}A_{\ell,+} \dots A_{1,+}(x), \end{aligned}$$

where $D_{\ell+1}$ is given in (12), which concludes the recursion. Let $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$. We have for any $x \in \bar{B}(x_0, \varepsilon_M)$

$$\begin{aligned} \langle \theta, F(x) \rangle &= \sum_{j \in \mathbf{j}} \sum_{k=1}^{c_j} \theta_{j,k} \tilde{v}_{j,k}^\top (\mathcal{G}_j(x) - \mathcal{G}_j(x_0)) \\ &= \sum_{j \in \mathbf{j}} \sum_{k=1}^{c_j} \theta_{j,k} \tilde{v}_{j,k}^\top \{A_{j,+} \dots A_{1,+}(x) - A_{j,+} \dots A_{1,+}(x_0)\} \\ &= \sum_{j \in \mathbf{j}} \sum_{k=1}^{c_j} \theta_{j,k} \tilde{v}_{j,k}^\top \tilde{A}_{j,+} \dots \tilde{A}_{1,+}(x - x_0) = \left\langle \sum_{j \in \mathbf{j}} \sum_{k=1}^{c_j} \theta_{j,k} v_{j,k}, x - x_0 \right\rangle. \end{aligned}$$

Since, $(v_{j,c})_{j \in \mathbf{j}, c \in \{1, \dots, c_j\}}$ is assumed to be linearly independent we have that $v = \sum_{j \in \mathbf{j}} \sum_{c=1}^{c_j} \theta_{j,c} v_{j,c}$ is non zero and therefore setting $x = x_0 - \varepsilon_M v / \|v\|$ we get that $\langle \theta, F(x) \rangle < 0$. Since F is continuous and $\mu(\mathbf{A}) > 0$ for every non-empty open set \mathbf{A} we have that for $\theta \in \mathbb{R}^p$ with $\|\theta\| = 1$, $\mu(\{x \in \mathbb{R}^d : \langle F(x), \theta \rangle < 0\}) > 0$, which concludes the proof upon using Proposition 5-(a).

B Proofs of Section 4

We start by introducing some notations. Let $V : \mathbb{R}^d \rightarrow [1, +\infty)$ be a measurable function. For $f \in \mathbb{F}(\mathbb{R}^d)$, the V -norm of f is given by $\|f\|_V = \|f/V\|_\infty$. Let ξ be a finite signed measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. The V -total variation norm of ξ is defined as

$$\|\xi\|_V = \sup_{f \in \mathbb{F}(\mathbb{R}^d), \|f\|_V \leq 1} \left| \int_{\mathbb{R}^d} f(x) d\xi(x) \right|.$$

If $V \equiv 1$, then $\|\cdot\|_V$ is the total variation norm denoted by $\|\cdot\|_{\text{TV}}$.

Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow (0, +\infty]$ be defined for any $x, y \in \mathbb{R}^d$ by $c(x, y) = \mathbb{1}_{\Delta_{\mathbb{R}^d}}(x, y)W(x, y)$ where $W : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, +\infty)$ is a lower semi-continuous function such that for any $x, y, z \in \mathbb{R}^d$, $W(x, y) = W(y, x)$ and $W(x, z) \leq W(x, y) + W(y, z)$. Then for any probability measures μ and ν such that there exist $x_\mu, x_\nu \in \mathbb{R}^d$ satisfying $\mu(W(\cdot, x_\mu)) < \infty$ and $\nu(W(\cdot, x_\nu)) < \infty$, we define the Wasserstein extended distance associated with cost c between μ and ν by

$$d_{W_R}(\mu, \nu) = \sup_{g \in \mathbb{G}_{\mu, W}} \left| \int_{\mathbb{R}^d} g(x) d\mu(x) - \int_{\mathbb{R}^d} g(y) d\nu(y) \right|, \quad (27)$$

with $\mathbb{G}_{\mu, W} = \{g \in \mathbb{F}(\mathbb{R}^d) : \|g(x) - g(y)\| \leq W(x, y), \text{ for all } x, y \in \mathbb{R}^d\}$.

B.1 Proof of Theorem 8

This proof is an application of [26, Theorem 2, Theorem 4]. Therefore, we are reduced to checking that [26, H1, H2] hold. More precisely, we study the geometric ergodicity of the Langevin Markov chain under **A2**(α), **B1**, **B2**(α) and **B3** with $U_2 = 0$ and $\alpha \geq 1$ as well as its discretization error. Foster-Lyapunov conditions are derived in Lemma 15 and we check that [26, H1-(a)] holds in Lemma 16. In Theorem 17 we show that [26, H1-(b)] is satisfied. We check that [26, H1-(c)] is satisfied in Lemma 18 and Proposition 19.

We denote by $K_{\gamma, \theta}$ the Markov kernel associated with the Langevin recursion (14). This kernel is given for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$

$$K_{\gamma, \theta}(x, A) = (2\pi\gamma)^{-d/2} \int_A \exp \left[-(2\gamma)^{-1} \|y - x + \gamma \nabla_x U(\theta, x)\|^2 \right] dy, \quad (28)$$

with U given by (17). Note that (28) is well-defined under **B1** and **B2**(α) with $\alpha \geq 1$. We say that a Markov kernel K on $\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)$ satisfies a discrete Foster-Lyapunov drift condition **D_d**(V, λ, b) if there exist $\lambda \in (0, 1)$, $b \geq 0$ and a measurable function $V : \mathbb{R}^d \rightarrow [1, +\infty)$ such that for all $x \in \mathbb{R}^d$

$$KV(x) \leq \lambda V(x) + b.$$

First, we state the following technical lemma.

Lemma 14. Let $m \in \mathbb{N}^*$. Then for any $u, v > 0$ and $t \geq 0$,

$$u(1+t)^{2m-1} - vt^{2m} \leq \Upsilon_m(u, v)$$

with

$$\Upsilon_m(u, v) = 2^{(4m-2)m} \max \left\{ u, u^{2m}/v^{2m-1} \right\} .$$

Proof. Let $m \in \mathbb{N}^*$, $\tilde{u}, \tilde{v} > 0$ and $\tilde{f}(t) = \tilde{u}t^{2m-1} - \tilde{v}t^{2m}$. We have for any $t \in \mathbb{R}$, $\tilde{f}'(t) = (2m-1)\tilde{u}t^{2m-2} - 2m\tilde{v}t^{2m-1}$. Since $\lim_{|t| \rightarrow +\infty} \tilde{f}(t) = -\infty$ and \tilde{f} is continuous, the maximum is attained at some point t_0 which satisfies

$$\tilde{f}'(t_0) = (2m-1)\tilde{u}t_0^{2m-2} - 2m\tilde{v}t_0^{2m-1} = 0 ,$$

and therefore $t_0 = (2m-1)\tilde{u}/(2m\tilde{v})$. We have for any $t \geq 0$

$$\tilde{u}t^{2m-1} - \tilde{v}t^{2m} \leq \tilde{u}t_0^{2m-1} \leq \tilde{u}(\tilde{u}/\tilde{v})^{2m-1} \leq \tilde{u}^{2m}/\tilde{v}^{2m-1} . \quad (29)$$

If $t \geq 1$ then $u(1+t)^{2m-1} - vt^{2m} \leq 2^{2m-1}ut^{2m-1} - vt^{2m}$ and using (29) we have $u(1+t)^{2m-1} - vt^{2m} \leq 2^{(4m-2)m}u^{2m}/v^{2m-1}$. If $t \leq 1$ then $u(1+t)^{2m-1} - vt^{2m} \leq 2^{2m-1}u$, which concludes the proof. \square

Lemma 15. Assume **B1**, **B2**(α) and **B3** with $U_2 = 0$ and $\alpha \geq 1$. Let $m \in \mathbb{N}^*$, $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}_1/\mathbb{N}^2, 1/2)$. Then $\mathbb{K}_{\gamma, \theta}$ satisfies $\mathbf{D}_d(V, \lambda^\gamma, \tilde{b}_\gamma)$ with

$$V(x) = 1 + \|x - x^*\|^{2m} , \quad \lambda = \exp[-\mathfrak{m}_1 + \bar{\gamma}\mathbb{N}^2] , \quad \tilde{b}_m = \Upsilon_m(2^{2m+1}d^m\Gamma(m+1/2), \mathfrak{m}_1) + \mathfrak{m}_1 , \quad (30)$$

where for any $t \geq 0$, $\Gamma(t) = \int_0^{+\infty} u^{t-1}e^{-u}du$ and Υ_m is given in Lemma 14. In addition, $\mathbb{K}_{\gamma, \theta}$ satisfies $\mathbf{D}_d(V, \lambda^\gamma, b_m(1+d^{\varpi_{0,m}})\gamma)$ with λ given in (30) and $b_m, \varpi_{0,m} \geq 0$ independent of the dimension d .

Proof. Let $x \in \mathbb{R}^d$, $m \in \mathbb{N}^*$, $\theta \in \mathbb{K}$, $\gamma \in (0, \bar{\gamma}]$ and Z a d -dimensional Gaussian random variable with zero mean and identity covariance matrix. First, denoting $\mathbf{Z} = (z_1, \dots, z_d)$ we have using Holder's inequality

$$\mathbb{E} \left[\|\mathbf{Z}\|^{2m} \right] = \sum_{i_1=1}^d \cdots \sum_{i_m=1}^d \mathbb{E} \left[\prod_{j=1}^m z_{i_j}^2 \right] \leq \sum_{i_1=1}^d \cdots \sum_{i_m=1}^d \mathbb{E} \left[z_1^{2m} \right] \leq (2d)^m \Gamma(m+1/2) . \quad (31)$$

Let $T_{\gamma, \theta}(x) = x - x^* - \gamma \nabla_x U(\theta, x)$. Using **B3** we get

$$\begin{aligned} \|T_{\gamma, \theta}(x)\|^2 &\leq \|x - x^*\|^2 - 2\gamma \langle \nabla_x U(\theta, x) - \nabla_x U(\theta, x^*), x - x^* \rangle + \|\nabla_x U(\theta, x) - \nabla_x U(\theta, x^*)\|^2 \\ &\leq (1 - 2\gamma\mathfrak{m}_1 + \gamma^2\mathbb{N}^2) \|x - x^*\|^2 . \end{aligned}$$

Hence, we obtain

$$\begin{aligned} \mathbb{E} \left[\|X - x^*\|^{2m} \right] &= \mathbb{E} \left[\sum_{k=0}^m \sum_{j=0}^k \binom{m}{k} \binom{k}{j} \|T_{\gamma, \theta}(x)\|^{2(m-k)} (2\gamma)^{j/2} \langle T_{\gamma, \theta}(x), Z \rangle^j (2\gamma)^{k-j} \|Z\|^{2(k-j)} \right] \\ &\leq (1 - \gamma(\mathfrak{m}_1 - 2N^2\bar{\gamma})) \|x - x^*\|^{2m} + \gamma C_m(x - x^*), \end{aligned} \quad (32)$$

where we have, using that $\|T_{\gamma, \theta}(x)\| \leq \|x - x^*\|$, (31), $2\gamma \leq 1$, the Isserlis' formula [52] and the Cauchy-Schwarz inequality

$$\begin{aligned} \gamma C_m(x - x^*) &= \sum_{k=1}^m \sum_{j=0}^k \binom{m}{k} \binom{k}{j} \|x - x^*\|^{2(m-k)} (2\gamma)^{k-j/2} \mathbb{E} \left[\langle T_{\gamma, \theta}(x), Z \rangle^j \|Z\|^{2(k-j)} \right] \\ &= \sum_{k=1}^m \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{m}{k} \binom{k}{2j} \|x - x^*\|^{2(m-k)} (2\gamma)^{k-j/2} \mathbb{E} \left[\langle T_{\gamma, \theta}(x), Z \rangle^{2j} \|Z\|^{2(k-2j)} \right] \\ &\leq 2\gamma \sum_{k=1}^m \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{m}{k} \binom{k}{2j} \|x - x^*\|^{2(m-k+j)} \mathbb{E} \left[\|Z\|^{2(k-j)} \right] \\ &\leq 2\gamma(1 + \|x - x^*\|)^{2m-1} (2d)^m \mathbf{\Gamma}(m + 1/2) \sum_{k=1}^m \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{m}{k} \binom{k}{2j} \\ &\leq 2^{2m+1} \gamma d^m \mathbf{\Gamma}(m + 1/2) (1 + \|x - x^*\|)^{2m-1}. \end{aligned} \quad (33)$$

Combining (32) and (33) we get that

$$\begin{aligned} \mathbf{K}_{\gamma, \theta}(\|x - x^*\|^{2m}) &\leq (1 - \gamma(\mathfrak{m}_1 - N^2\bar{\gamma})) \|x - x^*\|^{2m} \\ &\quad + 2^{2m+1} \gamma d^m \mathbf{\Gamma}(m + 1/2) (1 + \|x - x^*\|)^{2m-1} - \gamma \mathfrak{m}_1 \|x - x^*\|^{2m}. \end{aligned} \quad (34)$$

Using Lemma 14, we have

$$2^{2m+1} d^m \mathbf{\Gamma}(m + 1/2) (1 + \|x - x^*\|)^{2m-1} - \mathfrak{m}_1 \|x - x^*\|^{2m} \leq \Upsilon_m(2^{2m+1} d^m \mathbf{\Gamma}(m + 1/2), \mathfrak{m}_1).$$

Combining this result with (34) we get,

$$\mathbf{K}_{\gamma, \theta}(\|x - x^*\|^{2m}) \leq (1 - \gamma(\mathfrak{m}_1 - \bar{\gamma}N^2)) \|x - x^*\|^{2m} + \gamma \Upsilon_m(2^{2m+1} d^m \mathbf{\Gamma}(m + 1/2), \mathfrak{m}_1).$$

Therefore we obtain

$$\begin{aligned} \mathbf{K}_{\gamma, \theta}(1 + \|x - x^*\|^{2m}) &\leq (1 - \gamma(\mathfrak{m}_1 - \bar{\gamma}N^2))(1 + \|x - x^*\|^{2m}) \\ &\quad + \gamma \left\{ \Upsilon_m(2^{2m+1} d^m \mathbf{\Gamma}(m + 1/2), \mathfrak{m}_1) + \mathfrak{m}_1 \right\}, \end{aligned}$$

which concludes the proof upon noting that \tilde{b}_m is a polynomial in the dimension d . \square

Lemma 16. Assume **B1**, **B2**(α) and **B3** with $U_2 = 0$, $\alpha \geq 1$ and let $(X_k^n)_{n \in \mathbb{N}, k \in \{0, \dots, m_n\}}$ be given by (15) with $\bar{\gamma} < \min(m_1/N^2, 1/2)$. Let $m \in \mathbb{N}^*$, then there exist $A_{1,m} \geq 1$ and $\varpi_{1,m} \geq 0$ such that for any $n, p \in \mathbb{N}$ and $k \in \{0, \dots, m_n\}$

$$\mathbb{E} \left[K_{\gamma_n, \theta_n}^p V(X_k^n) \middle| X_0^0 \right] \leq A_{1,m} (1 + d^{\varpi_{1,m}}) V(X_0^0), \quad \mathbb{E} \left[V(X_0^0) \right] < +\infty,$$

with $V(x) = 1 + \|x - x^*\|^{2m}$ and $A_{1,m}, \varpi_{1,m} \geq 0$ which do not depend on the dimension d .

Proof. Combining [26, Lemma S15] and Lemma 15 conclude the proof. \square

Theorem 17. Assume **B1**, **B2**(α) and **B3** with $U_2 = 0$ and $\alpha \geq 1$. Then for any $m \in \mathbb{N}^*$ there exist $A_{2,m}, \varpi_{2,m} \geq 0$ and $\rho_m \in (0, 1)$ such that for any $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(m_1/N^2, 1/2)$, $K_{\gamma, \theta}$ admits an invariant probability measure $\pi_{\gamma, \theta}$ and for any $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\begin{aligned} \|\delta_x K_{\gamma, \theta}^n - \pi_{\gamma, \theta}\|_V &\leq A_{2,m} (1 + d^{\varpi_{2,m}}) \exp[-n\kappa_m \gamma / \log^2(1 + d^{\varpi_{2,m}})] V(x), \\ \|\delta_x K_{\gamma, \theta}^n - \delta_y K_{\gamma, \theta}^n\|_V &\leq A_{2,m} (1 + d^{\varpi_{2,m}}) \exp[-n\kappa_m \gamma / \log^2(1 + d^{\varpi_{2,m}})] \{V(x) + V(y)\}, \end{aligned}$$

with $V(x) = 1 + \|x - x^*\|^{2m}$ and $A_{2,m}, \varpi_{2,m} \geq 0$ and $\kappa_m > 0$ which do not depend on the dimension d .

Proof. For any $\gamma \in (0, \bar{\gamma}]$ and $\theta \in \mathbb{K}$, $K_{\gamma, \theta}$ has the Feller property and satisfies $\mathbf{D}_d(V, \lambda^\gamma, b\gamma)$ then [29, Theorem 12.3.3] applies and $K_{\gamma, \theta}$ admits an invariant probability measure $\pi_{\gamma, \theta}$.

Let $m \in \mathbb{N}^*$, $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$. Using [25, Proposition 3] with $\mathbf{A} \leftarrow \mathbb{R}^d$, we have for any $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\delta_{(x,y)} \tilde{K}_{\gamma, \theta}^{n[1/\gamma]}(\Delta_{\mathbb{R}^d}^c) \leq 1 - 2\Phi \left\{ -\alpha^{-1/2}(n) \|x - y\| / (2\sqrt{2}) \right\},$$

where for any $x, y \in \mathbb{R}^d$, $\tilde{K}_{\gamma, \theta}((x, y), \cdot)$ is the reflexive coupling between $K_{\gamma, \theta}(x, \cdot)$ and $K_{\gamma, \theta}(y, \cdot)$, see [25, Equation 5]. In addition, we have for any $n \in \mathbb{N}$

$$\alpha^{-1}(n) = (2m_1 - N^2 \bar{\gamma}) / \left\{ \exp((2m_1 - N^2 \bar{\gamma})n) - 1 \right\} \leq 2m_1 / (\lambda^{-2n} - 1) \leq 2m_1 \lambda^{2n} / (1 - \lambda).$$

and λ is given by (30). Therefore, we get that for any $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\delta_{(x,y)} \tilde{K}_{\gamma, \theta}^{n[1/\gamma]}(\Delta_{\mathbb{R}^d}^c) \leq 1 - 2\Phi \left\{ -\lambda^{n m_1^{1/2}} \|x - y\| / (1 - \lambda)^{1/2} \right\}.$$

For any $x, y \in \mathbb{R}^d$ let

$$\begin{aligned} W(x, y) &= 1 + (\|x - x^*\|^{2m} + \|y - x^*\|^{2m}) / 2, \\ K_d &= 2b_m (1 + d^{\varpi_{0,m}}) (1 + \bar{\gamma}) (1 + \log^{-1}(1/\lambda)), \quad M_d = 2K_d^{1/m}. \end{aligned}$$

Note that for any $x, y \in \mathbb{R}^d$ such that $\|x - y\| \geq M_d$, $W(x, y) \geq K_d$. In addition, let $n_0 = \max\{[-\log(m_1^{1/2} M_d / (1 - \lambda)^{1/2}) \log^{-1}(1/\lambda)], 0\}$. We have for any $x, y \in \mathbb{R}^d$ with $\|x - y\| \leq M_d$,

$$\delta_{(x,y)} \tilde{K}_{\gamma,\theta}^{n \lceil 1/\gamma \rceil} (\Delta_{\mathbb{R}^d}^c) \leq 1 - 2\Phi(-1).$$

Then, applying [25, Theorem 6b], with $\tilde{K} \leftarrow \tilde{K}_{\gamma,\theta}$, $m \leftarrow n_0$, $\varepsilon_{1,d} \leftarrow 2\Phi(-1)$, $\bar{C}_1 \leftarrow C$, $\bar{\rho}_1 \leftarrow \rho$, $\bar{A}_1 \leftarrow A$ and $\bar{c}_1 \leftarrow c$, we obtain that for any $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\|\delta_x K_{\gamma,\theta}^n - \delta_y K_{\gamma,\theta}^n\|_V \leq C \rho^{\lfloor k/(n_0 \lceil 1/\gamma \rceil) \rfloor} \{V(x) + V(y)\} / 2,$$

where

$$\begin{aligned} C &= 2 [1 + A] [1 + (A + K_d) / \{\Phi(-1)(1 - \lambda)\}] , \\ A &= b_m (1 + d^{\varpi_{0,m}}) (1 + \bar{\gamma}) (1 + \log^{-1}(1/\lambda)) , \\ \log^{-1}(1/\rho) &= \log^{-1}(1/(2\Phi(-1))) + \log^{-1}(2/(1 + \lambda)) \\ &\quad + \log(A + K_d) \log^{-1}(1/(2\Phi(-1))) \log^{-1}(2/(1 + \lambda)) . \end{aligned}$$

Since $\lfloor n/(n_0 \lceil 1/\gamma \rceil) \rfloor \geq n\gamma/(n_0(1 + \bar{\gamma})) - 1$, setting $\tilde{A}_{2,m} = C\rho^{-1}/2$ and $\rho_m = \rho^{1/(n_0(1 + \bar{\gamma}))}$, we get that for any $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\|\delta_x K_{\gamma,\theta}^n - \delta_y K_{\gamma,\theta}^n\|_V \leq \tilde{A}_{2,m} \rho_m^{\gamma n} \{V(x) + V(y)\} . \quad (35)$$

Using Lemma 15 we have that $\pi_{\gamma,\theta}(V) \leq b_m (1 + d^{\varpi_{0,m}}) (1 + \log^{-1}(1/\lambda)) V(x)$. Combining this result with (35) we get that for any $x \in \mathbb{R}^d$ and $n \in \mathbb{N}$

$$\|\delta_x K_{\gamma,\theta}^n - \pi_{\gamma,\theta}\|_V \leq \tilde{A}_{2,m} \left\{ 1 + b_m (1 + d^{\varpi_{0,m}}) (1 + \log^{-1}(1/\lambda)) \right\} \tilde{\rho}_m^{\gamma n} V(x) .$$

Since in Lemma 15, λ and b_m do not depend on the dimension d we get that K_d is upper-bounded by a polynomial in the dimension d . Hence, there exists $\varpi_{2,m}^{(a)} > 0$ which does not depend on the dimension such that $\tilde{A}_{2,m} \{1 + b_m (1 + d^{\varpi_{0,m}}) (1 + \log^{-1}(1/\lambda))\} \leq A_{2,m} (1 + d^{\varpi_{2,m}^{(a)}})$ with $A_{2,m} \geq 0$ which does not depend on the dimension d . Similarly, there exists $\varpi_{2,m}^{(b)} > 0$ independent of d such that $\sup_{d \in \mathbb{N}} \{[\log^{-1}(\rho) + n_0] / \log(1 + d^{\varpi_{2,m}^{(b)}})^{-1}\} < +\infty$ which implies that $\log^{-1}(1/\rho_m) \leq \kappa_m^{-1} \log^2(1 + d^{\varpi_{2,m}^{(b)}})$ with $\kappa_m > 0$ which do not depend on the dimension d . We conclude the proof upon setting $\varpi_{2,m} = \max(\varpi_{2,m}^{(a)}, \varpi_{2,m}^{(b)})$. \square

Similarly to the discrete setting, we say that a Markov semi-group $(P_t)_{t \geq 0}$ on $\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)$ with extended infinitesimal generator $(\mathcal{A}, D(\mathcal{A}))$ (see e.g. [68] for the definition of $(\mathcal{A}, D(\mathcal{A}))$) satisfies a continuous drift condition $\mathbf{D}_c(V, \zeta, \beta)$ if there exist $\zeta > 0$, $\beta \geq 0$ and a measurable function $V : \mathbb{R}^d \rightarrow [1, +\infty)$ with $V \in D(\mathcal{A})$ such that for all $x \in \mathbb{R}^d$

$$\mathcal{A}V(x) \leq -\zeta V(x) + \beta .$$

Let $\theta \in \mathbf{K}$ and $(P_{t,\theta})_{t \geq 0}$ be the Markov semi-group associated with the Langevin diffusion

$$d\mathbf{X}_t = - \left(\sum_{i=1}^p \theta(i) \nabla F_i(\mathbf{X}_t) + \nabla r(\mathbf{X}_t) \right) + d\mathbf{B}_t ,$$

where $(\mathbf{B}_t)_{t \geq 0}$ is a d -dimensional Brownian motion. Consider now the generator \mathcal{A}_θ of $(P_{t,\theta})_{t \geq 0}$ for any $\theta \in \mathbf{K}$, defined for any $f \in C^2(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$ by

$$\mathcal{A}_\theta f(x) = - \left\langle \nabla f(x), \sum_{i=1}^p \theta(i) \nabla F_i(x) + \nabla r(x) \right\rangle + \Delta f(x) .$$

Using [26, Lemma S14] we have that π_θ is an invariant probability measure for $(P_{t,\theta})_{t \geq 0}$.

Lemma 18. *Assume **B1**, **B2**(α) and **B3** with $U_2 = 0$ and $\alpha \geq 1$. Then for any $m \in \mathbb{N}^*$ there exist $\zeta > 0$ and $\beta \geq 0$ such that for any $\theta \in \mathbf{K}$, \mathcal{A}_θ satisfies $\mathbf{D}_c(V, \zeta, \tilde{\beta}_m)$ with*

$$V(x) = 1 + \|x - x^*\|^{2m} , \quad \zeta = -\mathfrak{m}_1 m , \quad \tilde{\beta}_m = 2m \Upsilon_m(2(m-1) + d, \mathfrak{m}_1/2) + 2\mathfrak{m}_1 m , \quad (36)$$

with Υ_m given in Lemma 14. In addition, \mathcal{A}_θ satisfies $\mathbf{D}_c(V, \zeta, \beta_m(1 + d^{\varpi'_{0,m}})\gamma)$ with ζ given in (36) and $\beta_m, \varpi'_{0,m} \geq 0$ independent of the dimension d .

Proof. Let $\theta \in \mathbf{K}$ and $m \in \mathbb{N}^*$. Then, we have for any $x \in \mathbb{R}^d$

$$\begin{aligned} V(x) &= 1 + \|x - x^*\|^{2m} , \\ \nabla V(x) &= 2m \|x - x^*\|^{2(m-1)} (x - x^*) , \\ \nabla^2 V(x) &= 4m(m-1) \|x - x^*\|^{2(m-2)} (x - x^*)(x - x^*)^\top + 2m \|x - x^*\|^{2(m-1)} \text{Id} . \end{aligned} \quad (37)$$

Hence, for any $x \in \mathbb{R}^d$, $\Delta V(x) = 4m(m-1) \|x - x^*\|^{2(m-1)} + 2md \|x - x^*\|^{2(m-1)}$. Using **B3**, (37) and Lemma 14 we get that for any $x \in \mathbb{R}^d$

$$\begin{aligned} \mathcal{A}_\theta V(x) &= -2m \|x - x^*\|^{2(m-1)} \langle \nabla_x U(\theta, x), x - x^* \rangle + 2m(2(m-1) + d) \|x - x^*\|^{2(m-1)} \\ &= -2m \|x - x^*\|^{2(m-1)} \langle \nabla_x U(\theta, x) - \nabla_x U(\theta, x^*), x - x^* \rangle + 2m(2(m-1) + d) \|x - x^*\|^{2(m-1)} \\ &\leq -2\mathfrak{m}_1 m V(x) + 2m(2(m-1) + d) \|x - x^*\|^{2(m-1)} + 2\mathfrak{m}_1 m \\ &\leq -\mathfrak{m}_1 m V(x) + 2m \left(2(m-1) + d - \mathfrak{m}_1 \|x - x^*\|^2 / 2 \right) \|x - x^*\|^{2(m-1)} + 2\mathfrak{m}_1 m \\ &\leq -\mathfrak{m}_1 m V(x) + 2m \Upsilon_m(2(m-1) + d, \mathfrak{m}_1/2) + 2\mathfrak{m}_1 m , \end{aligned}$$

which concludes the proof. \square

Proposition 19. *Assume **B1**, **B2**(α) and **B3** with $U_2 = 0$ and $\alpha \geq 1$. Then for any $m \in \mathbb{N}^*$, there exist $A_{3,m}, \varpi_{3,m} \geq 0$ such that for any $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}_1/N^2, 1/2)$,*

$$\|\pi_{\gamma,\theta} - \pi_\theta\|_{V^{1/2}} \leq A_{3,m}(1 + d^{\varpi_{3,m}})\gamma^{1/2} ,$$

with $V(x) = 1 + \|x - x^*\|^{2m}$ and $A_{3,m}, \varpi_{3,m} \geq 0$ which do not depend on the dimension d .

The proof is similar to the one of [26, Proposition S17] except that in this presentation we explicit the constants appearing in the proof and track the dependency of the constants with respect to the dimension d .

Proof. Let $m \in \mathbb{N}^*$, $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$. Since π_θ is an invariant probability measure for $(\mathbf{P}_{t,\theta})_{t \geq 0}$ we have using Theorem 17 that

$$\lim_{k \rightarrow +\infty} \|\pi_\theta \mathbf{K}_{\gamma,\theta}^k - \pi_\theta \mathbf{P}_{\gamma k,\theta}\|_{V^{1/2}} = \|\pi_{\gamma,\theta} - \pi_\theta\|_{V^{1/2}} .$$

We now give an upper bound on $\|\pi_\theta \mathbf{K}_{\gamma,\theta}^k - \pi_\theta \mathbf{P}_{\gamma k,\theta}\|_{V^{1/2}}$ for $k = q_\gamma m_\gamma$ with $m_\gamma = \lceil 1/\gamma \rceil$ and $q_\gamma \in \mathbb{N}$. We obtain using Theorem 17

$$\begin{aligned} & \|\pi_\theta \mathbf{K}_{\gamma,\theta}^k - \pi_\theta \mathbf{P}_{\gamma k,\theta}\|_{V^{1/2}} \\ & \leq \sum_{j=0}^{q_\gamma-1} \|\pi_\theta \mathbf{P}_{\gamma(j+1)m_\gamma,\theta} \mathbf{K}_{\gamma,\theta}^{(q_\gamma-(j+1))m_\gamma} - \pi_\theta \mathbf{P}_{\gamma j m_\gamma,\theta} \mathbf{K}_{\gamma,\theta}^{(q_\gamma-j)m_\gamma}\|_{V^{1/2}} \\ & \leq \sum_{j=0}^{q_\gamma-1} \left\{ A_{2,m}(1 + d^{\varpi_{2,m}}) \exp[-\kappa_m \gamma m_\gamma (q_\gamma - (j+1)) / \log^2(1 + d^{\varpi_{2,m}})] \right. \\ & \quad \left. \times \|\pi_\theta \mathbf{P}_{\gamma j m_\gamma,\theta} \mathbf{P}_{m_\gamma \gamma,\theta} - \pi_\theta \mathbf{P}_{\gamma j m_\gamma,\theta} \mathbf{K}_{\gamma,\theta}^{m_\gamma}\|_{V^{1/2}} \right\} \\ & \leq \|\pi_\theta \mathbf{P}_{m_\gamma \gamma,\theta} - \pi_\theta \mathbf{K}_{\gamma,\theta}^{m_\gamma}\|_{V^{1/2}} \sum_{j=1}^{q_\gamma} A_{2,m}(1 + d^{\varpi_{2,m}}) \exp[-\kappa_m \gamma j m_\gamma / \log^2(1 + d^{\varpi_{2,m}})] \\ & \leq A_{2,m}(1 + d^{\varpi_{2,m}})(1 + \log^2(1 + d^{\varpi_{2,m}}) / \kappa_m) \|\pi_\theta \mathbf{P}_{m_\gamma \gamma,\theta} - \pi_\theta \mathbf{K}_{\gamma,\theta}^{m_\gamma}\|_{V^{1/2}} , \end{aligned} \quad (38)$$

We now give an upper bound on $\|\pi_\theta \mathbf{P}_{m_\gamma \gamma,\theta} - \pi_\theta \mathbf{K}_{\gamma,\theta}^{m_\gamma}\|_{V^{1/2}}$. Indeed, since \mathcal{A}_θ satisfies $\mathbf{D}_c(V, \zeta, \beta)$ by Lemma 18 and $\mathbf{K}_{\gamma,\theta}$ satisfies $\mathbf{D}_d(V, \lambda^\gamma, b\gamma)$ for any $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$ by Lemma 15 and, we obtain that

$$\begin{aligned} \pi_\theta \mathbf{P}_{\gamma m_\gamma,\theta}(V) & \leq D_0 , & \pi_\theta \mathbf{K}_{\gamma,\theta}^{m_\gamma}(V) & \leq D_1 , \\ D_0 & = \beta_m(1 + d^{\varpi_{0,m}}) / \zeta , & D_1 & = D_0 + b_m(1 + d^{\varpi_{0,m}})(\bar{\gamma} + \log^{-1}(1/\lambda)) . \end{aligned} \quad (39)$$

Combining this result and [26, Lemma S16] we have for any $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$

$$\|\pi_\theta \mathbf{P}_{\gamma m_\gamma,\theta} - \pi_\theta \mathbf{K}_{\gamma,\theta}^{m_\gamma}\|_{V^{1/2}} \leq D_2 \gamma^{1/2} , \quad (40)$$

with

$$D_2 = 2D_1^{1/2}(1 + \bar{\gamma})^{1/2} \left\{ d + 2\bar{\gamma}(2N^2 + \sup_{\theta \in \mathbb{K}} \|\nabla_x U_1(\theta, 0)\|^2) D_1 \right\}^{1/2} \mathbb{N} . \quad (41)$$

Combining (38) and (40) we get

$$\|\pi_\theta \mathbf{K}_{\gamma,\theta}^k - \pi_\theta \mathbf{P}_{\gamma k,\theta}\|_{V^{1/2}} \leq D_2 A_{2,m}(1 + d^{\varpi_{2,m}})(1 + \log^2(1 + d^{\varpi_{2,m}}) / \kappa_m) \gamma^{1/2} .$$

Combining (41), (39) and Lemma 15, we get that D_0 , D_1 , D_2 are upper-bounded by polynomials in the dimension d , which concludes the proof. \square

We now turn to the proof of Theorem 8.

Proof of Theorem 8. Let $m = \lceil 2\alpha \rceil$ and $V(x) = 1 + \|x - x^*\|^{2m}$. Lemma 16 implies [26, H1a] with $A_{1,m} \leftarrow A_{1,m}(1 + d^{\varpi_{1,m}})$, Theorem 17 implies [26, H1b] with $A_{2,m} \leftarrow A_{2,m}(1 + d^{\varpi_{2,m}})$ and $\rho \leftarrow \exp[-\kappa_m / \log^2(1 + d^{\varpi_{2,m}})]$. In addition, Proposition 19 implies [26, H1c] with $\Psi(\gamma) = A_{3,m}(1 + d^{\varpi_{3,m}})\gamma^{1/2}$. Using Proposition 4 we have that [26, A1, A2, A3] hold. Since $H_\theta \leftarrow F$ in (15) we get that for any $\theta \in \mathbb{K}$ and $x \in \mathbb{R}^d$, $\|H_\theta(x)\| \leq V^{1/2}(x)$. We can apply [26, Theorem 2] and we get that for any $n \in \mathbb{N}$

$$\mathbb{E} \left[\left\{ \frac{\sum_{k=1}^n \delta_k L(\theta_k)}{\sum_{k=1}^n \delta_k} \right\} - \min_{\mathbb{K}} L \right] \leq E_n / \left(\sum_{k=1}^n \delta_k \right),$$

with,

$$\begin{aligned} E_n &= 2R_{\mathbb{K}}^2 + 2B_{1,m}R_{\mathbb{K}}\mathbb{E} \left[V^{1/2}(X_0^0) \right] \sum_{k=0}^{n-1} \delta_{k+1} / (m_k \gamma_k) \\ &\quad + 2R_{\mathbb{K}}A_{3,m}(1 + d^{\varpi_{3,m}}) \sum_{k=0}^{n-1} \delta_{k+1} \gamma_k^{1/2} + 4B_{1,m}^2 \mathbb{E} \left[V(X_0^0) \right] \sum_{k=0}^{n-1} \delta_{k+1}^2 / (m_k \gamma_k)^2 \\ &\quad + 4A_{3,m}^2 (1 + d^{\varpi_{3,m}})^2 \sum_{k=0}^{n-1} \delta_{k+1}^2 \gamma_k + B_{2,m} \sum_{k=0}^{n-1} \delta_{k+1}^2 / (m_k \gamma_k)^2, \end{aligned} \quad (42)$$

and

$$\begin{aligned} B_{1,m} &= 2(1 + d^{\varpi_m})^2 \log^2(1 + d^{\varpi_m}) A_{1,m} A_{2,m} \exp[-\bar{\gamma} \kappa_m / \log^2(1 + d^{\varpi_m})] / \kappa_m; \\ B_{2,m} &= 2(1 + \bar{\gamma})^2 \max(B_{2,m}^{(a)}, B_{2,m}^{(b)}); \\ B_{2,m}^{(a)} &= 24(1 + d^{\varpi_m})^3 A_{2,m}^2 (1 - \exp[-\kappa_m / (2 \log^2(1 + d^{\varpi_m}))])^{-2} A_{3,m}; \\ B_{2,m}^{(b)} &= 4(1 + d^{\varpi_m})^3 A_{1,m} \left[1 + 6A_{2,m}^2 (1 - \exp[-\kappa_m / (2 \log^2(1 + d^{\varpi_m}))])^{-2} \right. \\ &\quad \left. \times \left\{ A_{2,m} (1 - \exp[-\kappa_m / \log^2(1 + d^{\varpi_m})])^{-1} + 2 \right\} + A_{2,m}^2 \log^4(1 + d^{\varpi_m}) / \kappa_m^2 + A_{3,m}^2 \right], \end{aligned}$$

which concludes the proof of Theorem 8-(b), upon setting $\varpi_m = \max_{i \in \{1,2,3\}} (\varpi_{i,m})$.

We have that for any $x \in \mathbb{R}^d$, $\|H_\theta(x)\| = \|F(x)\| \leq V^{1/4}(x)$. Since H_θ does not depend on θ we get that [26, A4] is satisfied. In addition [26, Proposition 24] implies that [26, H2] is satisfied with

$$\mathbf{\Lambda}_1(\gamma_1, \gamma_2) = A_{4,m}(1 + d^{\varpi_{4,m}})\gamma_2^{-1/2} |\gamma_1 - \gamma_2|, \quad \mathbf{\Lambda}_2(\gamma_1, \gamma_2) = A_{4,m}(1 + d^{\varpi_{4,m}})\gamma_2^{1/2},$$

with $A_{4,m} \geq 0$ which does not depend on the dimension d . As a consequence we can apply [26, Theorem 4] and we get that for any $n \in \mathbb{N}^*$

$$\mathbb{E} \left[\left\{ \frac{\sum_{k=1}^n \delta_k L(\theta_k)}{\sum_{k=1}^n \delta_k} \right\} - \min_{\mathbb{K}} L \right] \leq \tilde{E}_n / \left(\sum_{k=1}^n \delta_k \right),$$

with,

$$\begin{aligned}
\tilde{E}_n &= 2R_K + 2R_K \sum_{k=0}^n \delta_{k+1} \Psi(\gamma_k) + C_{3,m} \sum_{k=0}^n |\delta_{k+1} - \delta_k| \gamma_k^{-1} \\
&+ 2R_K C_{2,m} \sum_{k=0}^n \delta_{k+1} \gamma_{k+1}^{-2} [\mathbf{A}_1(\gamma_k, \gamma_{k+1}) + \mathbf{A}_2(\gamma_k, \gamma_{k+1}) \delta_{k+1} + \delta_{k+1} \gamma_{k+1}] \\
&+ C_{3,m} \sum_{k=0}^n \delta_{k+1}^2 \gamma_{k+1}^{-1} + C_{3,m} (\delta_{n+1}/\gamma_n - \delta_0/\gamma_0) + C_{1,m} \sum_{k=0}^n \delta_{k+1}^2,
\end{aligned}$$

□

with

$$\begin{aligned}
C_{1,m} &= 2A_{1,m}(1 + d^{\varpi_m}) \mathbb{E} [V(X_0^0)] + 2 \sup_{\theta \in \mathbf{K}} \|\nabla L(\theta)\|^2, \\
C_{2,m} &= 8(1 + d^{\varpi_m})^4 \log^4(1 + d^{\varpi_m}) A_{1,m} A_{2,m}^2 \\
&\quad \times \exp[-2\bar{\gamma} \kappa_m / \log^2(1 + d^{\varpi_m})] (1 + 2A_{1,m} \mathbb{E} [V(X_0^0)]) / \kappa_m, \\
C_{3,m} &= (1 + d^{\varpi_m}) A_{1,m} C_H (4R_K + \sup_{\theta \in \mathbf{K}} \|\nabla L(\theta)\| + 1 + \delta_1 \mathbf{B}) \mathbb{E} [V(X_0^0)^{1/4}], \\
C_H &= 8(1 + d^{\varpi_m}) \log^2(1 + d^{\varpi_m}) A_{2,m} \exp[-\bar{\gamma} \kappa_m / (4 \log^2(1 + d^{\varpi_m}))] / \kappa_m.
\end{aligned}$$

Similarly to Theorem 8-(b) since $A_{1,m}$, $A_{2,m}$, $A_{3,m}$, $A_{4,m}$ and κ_m are independent of the dimension d . Setting $\varpi_m = \max_{i \in \{1,2,3,4\}} (\varpi_{i,m})$ concludes the proof of Theorem 8-(a)

B.2 Proof of Theorem 9

In this section, we give alternative results to [26, Theorem 2, Theorem 4]. The main results of [26] are stated in V -norm or total variation. However, our particular framework allows us to use a Wasserstein distance with an appropriate cost function which implies that the constants appearing in our results scale polynomially in the dimension d even if the potential is non convex. The increasing batch size case is considered in Appendix B.2.1 and the fixed batch size case in Appendix B.2.2. We check that the main assumptions **H1** and **H2** below are satisfied in the setting of Theorem 9 in Appendix B.2.3 and conclude in Appendix B.2.3.

B.2.1 Increasing batch size

In this section, we give an alternative result of [26, Theorem 2] in the case where (a) the controls on the family of Markov kernels $\{K_{\gamma,\theta} : \gamma \in (0, \bar{\gamma}], \theta \in \mathbf{K}\}$ are obtained with respect to an appropriate Wasserstein distance (b) the stochastic gradient does not depend on θ . First, we show that under **B2**(1), $\mu \mapsto \mu(F)$ is Lipschitz with respect to the considered

Wasserstein distance in Lemma 21. Then, we control the error in the perturbed gradient scheme in Lemma 22 and Lemma 23. Our main result is stated in Theorem 20.

Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, +\infty)$, defined for any $x, y \in \mathbb{R}^d$ by $c(x, y) = \mathbb{1}_{\Delta_{\mathbb{R}^d}}(x, y)(1 + \|x - y\|/R)$ where $R \geq 0$. Consider also the function $W_R : \mathbb{R}^{2d} \rightarrow \mathbb{R}_+$, and $V_R : \mathbb{R}^d \rightarrow \mathbb{R}_+$ given for $x, y \in \mathbb{R}^d$ by

$$W_R(x, y) = 1 + \|x - y\|/R, \quad V_R(x) = 1 + \|x\|/R. \quad (43)$$

We also define for any $m \in \mathbb{N}$, $\mathcal{V}_m : \mathbb{R}^d \rightarrow [1, +\infty)$ given for any $x \in \mathbb{R}^d$ by

$$\mathcal{V}_m(x) = 1 + \|x\|^{2m}. \quad (44)$$

We recall that $K_{\gamma, \theta}$ is the Markov kernel associated with the Langevin recursion (14) and expression given by (28). This kernel is well-defined under **B1** and **B2**(α) with $\alpha \geq 1$. Consider the following assumption.

H1. (i) *There exists $A_1 \geq 1$ such that for any $a \in [1, 3]$, $n, p \in \mathbb{N}$, $k \in \{0, \dots, m_n\}$*

$$\mathbb{E} \left[K_{\gamma_n, \theta_n}^p V_R^a(X_k^n) \middle| X_0^0 \right] \leq A_1 V_R^a(X_0^0), \quad \mathbb{E} \left[V_R^a(X_0^0) \right] < +\infty.$$

with $\{(X_k^\ell)_{k \in \{0, \dots, m_\ell\}} : \ell \in \{0, \dots, n\}\}$ given by (15).

(ii) *There exist $A_2, A_3 \geq 1$, $\rho \in [0, 1)$ such that for any $\gamma \in (0, \bar{\gamma}]$, $\theta \in \mathbb{K}$, $x, y \in \mathbb{R}^d$ and $n \in \mathbb{N}$, $K_{\gamma, \theta}$ has a stationary distribution $\pi_{\gamma, \theta}$ and*

$$d_{W_R}(\delta_x K_{\gamma, \theta}^n, \delta_y K_{\gamma, \theta}^n) \leq A_2 \rho^{\gamma n} W_R(x, y), \quad d_{W_R}(\delta_x K_{\gamma, \theta}^n, \pi_{\gamma, \theta}) \leq A_2 \rho^{\gamma n} V_R(x), \quad \pi_{\gamma, \theta}(V_R) \leq A_3.$$

(iii) *There exists $\Psi : \mathbb{R}_+^* \rightarrow \mathbb{R}_+$ such that for any $\gamma \in (0, \bar{\gamma}]$ and $\theta \in \mathbb{K}$, $d_{W_R}(\pi_{\gamma, \theta}, \pi_\theta) \leq \Psi(\gamma)$.*

Theorem 20. *Assume **A2**(1), **B1**, **B2**(1) and **H1**. Let $(\gamma_n)_{n \in \mathbb{N}}$, $(\delta_n)_{n \in \mathbb{N}}$ be sequences of non-increasing positive real numbers and $(m_n)_{n \in \mathbb{N}}$ a sequence of positive integers satisfying $\delta_n < (\sup_{\theta \in \mathbb{K}} \|\nabla^2 L(\theta)\|)^{-1}$ and $\gamma_n < \bar{\gamma}$. Then, there exists $(E_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}^*$*

$$\mathbb{E} \left[\left\{ \frac{\sum_{k=1}^n \delta_k L(\theta_k)}{\sum_{k=1}^n \delta_k} \right\} - \min_{\mathbb{K}} L \right] \leq E_n / \left(\sum_{k=1}^n \delta_k \right),$$

with $(\theta_k)_{k \in \mathbb{N}}$ and L are defined in (15) and (5) respectively, and for any $n \in \mathbb{N}^$*

$$\begin{aligned} E_n &= 2R_{\mathbb{K}}^2 + 6R_{\mathbb{K}} R M p A_1 A_2 (\rho^{-\bar{\gamma}} / \log(1/\rho) + 1) \sum_{k=0}^{n-1} \delta_{k+1} \{1/(m_k \gamma_k) + \Psi(\gamma_k)\} \\ &\quad + \left(2A_1 (\|F(0)\| + 3RM)^2 \mathbb{E} \left[V_R^2(X_0) \right] + 2 \sup_{\theta \in \mathbb{K}} L(\theta)^2 \right) \sum_{k=0}^{n-1} \delta_{k+1}^2. \end{aligned}$$

The proof of this result is a simple adaptation to the one of [26, Theorem 2]. However it is given for completeness.

Let $(\eta_n)_{n \in \mathbb{N}}$ be defined for any $n \in \mathbb{N}$ by

$$\eta_n = m_n^{-1} \sum_{k=1}^{m_n} \{F(X_k^n) - \pi_{\theta_n}(F)\} . \quad (45)$$

We consider the following decomposition for any $n \in \mathbb{N}$,

$$\eta_n = \eta_n^{(1)} + \eta_n^{(2)} , \quad \eta_n^{(1)} = \mathbb{E} [\eta_n | \mathcal{F}_{n-1}] , \quad \eta_n^{(2)} = \eta_n - \mathbb{E} [\eta_n | \mathcal{F}_{n-1}] , \quad (46)$$

and $(\mathcal{F}_n)_{n \in \mathbb{N} \cup \{-1\}}$ is defined for all $n \in \mathbb{N}$ by

$$\mathcal{F}_n = \sigma \left(\theta_0, \{(X_k^\ell)_{k \in \{0, \dots, m_\ell\}} : \ell \in \{0, \dots, n\}\} \right) , \quad \mathcal{F}_{-1} = \sigma(\theta_0) \quad (47)$$

We start with the following technical lemmas.

Lemma 21. *Assume **B2**(1) and **H1**. Then for any probability measures μ, ν on $\mathcal{B}(\mathbb{R}^d)$ such that $\mu(\|\cdot\|) + \nu(\|\cdot\|) < +\infty$,*

$$\|\mu(F) - \nu(F)\| \leq 3RMpd_{W_R}(\mu, \nu) ,$$

with W_R given in (43).

Proof. Using **B2**(1) we have that for any $i \in \{1, \dots, p\}$ and $x, y \in \mathbb{R}^d$, $|F_i(x) - F_i(y)| \leq 3M \|x - y\| \leq 3RMW_R(x, y)$. Let μ and ν on $\mathcal{B}(\mathbb{R}^d)$ such that $\mu(\|\cdot\|) + \nu(\|\cdot\|) < +\infty$. Using the definition of the Wasserstein distance (27), we have $\|\mu(F) - \nu(F)\| \leq \sum_{i=1}^p |\mu(F_i) - \nu(F_i)| \leq 3RMpd_{W_R}(\mu, \nu)$. \square

Lemma 22. *Assume **B1**, **B2**(1) and **H1**. Then we have for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\|\eta_n^{(1)}\| \right] \leq B_1 \left\{ \mathbb{E} \left[V_R(X_0^0) \right] / (m_n \gamma_n) + \Psi(\gamma_n) \right\} ,$$

with $B_1 = 3RMpA_1A_2(\rho^{-\bar{\gamma}} / \log(1/\rho) + 1)$.

Proof. Using the definition of $(\mathcal{F}_n)_{n \in \mathbb{N}}$, see (47), the Markov property, **H1**-(ii)-(iii), Lemma 21 and that for any $\theta \in \mathbb{K}$, we have for any $n \in \mathbb{N}^*$

$$\begin{aligned} \|\mathbb{E} [\eta_n | \mathcal{F}_{n-1}]\| &\leq m_n^{-1} \sum_{k=1}^{m_n} \left\| \mathbb{K}_{\gamma_n, \theta_n}^k F(X_0^n) - \pi_{\theta_n}(F) \right\| \\ &\leq 3RMpm_n^{-1} \sum_{k=1}^{m_n} \left\{ d_{W_R}(\delta_{X_0^n} \mathbb{K}_{\gamma_n, \theta_n}^k, \pi_{\gamma_n, \theta_n}) \right\} + 3RMpd_{W_R}(\pi_{\gamma_n, \theta_n}, \pi_{\theta_n}) \\ &\leq 3RMpm_n^{-1} \sum_{k=1}^{m_n} \left\{ A_2 \rho^{\gamma_n k} V_R(X_{m_n}^n) \right\} + 3RMp\Psi(\gamma_n) \leq \frac{3RMpA_2\rho^{-\bar{\gamma}}V_R(X_{m_n}^n)}{\log(1/\rho)\gamma_n m_n} + 3RMp\Psi(\gamma_n) . \end{aligned}$$

In a similar manner, we have

$$\left\| \mathbb{E} \left[\eta_0 \middle| X_0^0 \right] \right\| \leq \frac{3RMpA_2\rho^{-\bar{\gamma}}V_R(X_0^0)}{\log(1/\rho)\gamma_0m_0} + 3RMp\Psi(\gamma_0).$$

We conclude using **H1-(i)**. \square

Lemma 23. *Assume **A2(1)**, **B1**, **B2(1)** and **H1**. Then we have for any $n \in \mathbb{N}$, $\mathbb{E} [\|\eta_n\|^2] \leq B_2$, with*

$$B_2 = 2A_1(\|F(0)\| + 3RM)^2\mathbb{E} \left[V_R^2(X_0^0) \right] + 2 \sup_{\theta \in \mathbb{K}} \|\nabla L(\theta)\|^2.$$

Proof. Using that $\|x + y\|^2 \leq 2(\|x\|^2 + \|y\|^2)$ for any $x, y \in \mathbb{R}^d$, the Cauchy-Schwarz inequality **H1-(i)** and Proposition 4, we get for any $n \in \mathbb{N}$,

$$\mathbb{E} \left[\|\eta_n\|^2 \right] \leq 2m_n^{-1} \sum_{k=1}^{m_n} \|F(X_k^n)\|^2 + 2\|\nabla L(\theta_n)\|^2.$$

We conclude using that for any $x \in \mathbb{R}^d$, $\|F(x)\| \leq (\|F(0)\| + 3RM)V_R(x)$ and the fact that $\sup_{\theta \in \mathbb{K}} \|\nabla L(\theta)\| < +\infty$. \square

Proof of Theorem 20. Taking the expectation in [4, Theorem 3, Equation (8)], using the Cauchy-Schwarz inequality and the fact that $(\eta_n^{(2)})_{n \in \mathbb{N}}$ is a martingale increment with respect to $(\mathcal{F}_n)_{n \in \mathbb{N}}$, we get that for every $n \in \mathbb{N}$

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^n \delta_k \left\{ L(\theta_k) - \min_{\mathbb{K}} L \right\} \right] \\ & \leq \mathbb{E} \left[2R_{\mathbb{K}}^2 - \sum_{k=0}^{n-1} \delta_{k+1} \langle \Pi_{\mathbb{K}}(\theta_k - \delta_{k+1} \nabla L(\theta_k)) - \theta^*, \eta_k \rangle + \sum_{k=0}^{n-1} \delta_{k+1}^2 \|\eta_k\|^2 \right] \\ & \leq 2R_{\mathbb{K}}^2 + 2R_{\mathbb{K}} \sum_{k=0}^{n-1} \delta_{k+1} \mathbb{E} \left[\|\eta_k^{(1)}\| \right] + 2 \sum_{k=0}^{n-1} \delta_{k+1}^2 \mathbb{E} \left[\|\eta_k\|^2 \right]. \end{aligned}$$

Combining this result, Lemma 22 and Lemma 23 completes the proof. \square

B.2.2 Fixed batch size

In this section, we give an alternative result of [26, Theorem 4], in the case where $m_n = 1$ and $\gamma_n = \gamma_0$ for all $n \in \mathbb{N}$. We consider the following additional assumption on the family of kernels $\{\mathbb{K}_{\gamma, \theta} : \theta \in \mathbb{K}, \gamma \in (0, \bar{\gamma}]\}$.

H2. *There exists $\mathbf{\Lambda} : \mathbb{R}_+^* \rightarrow \mathbb{R}_+$ such that for any $\gamma \in (0, \bar{\gamma}]$, $\theta_1, \theta_2 \in \mathbb{K}$, $x \in \mathbb{R}^d$*

$$\|\delta_x \mathbb{K}_{\gamma, \theta_1} - \delta_x \mathbb{K}_{\gamma, \theta_2}\|_{V_R} \leq \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x).$$

Theorem 24. Assume **A2**(1), **B1**, **B2**(1), **B3**, **B4**, **H1** and **H2**. Let $(\gamma_n)_{n \in \mathbb{N}}$, $(\delta_n)_{n \in \mathbb{N}}$ be sequences of non-increasing positive real numbers and $(m_n)_{n \in \mathbb{N}}$ a sequence of positive integers satisfying $\delta_n < 1/(\sup_{\theta \in \mathbb{K}} \|\nabla^2 L(\theta)\|)$ and for any $n \in \mathbb{N}$, $\gamma_n = \gamma < \bar{\gamma}$, $m_n = m_0$ and $\sup_{n \in \mathbb{N}} |\delta_{n+1} - \delta_n| \delta_n^{-2} < +\infty$. Then, there exists $(\tilde{E}_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}^*$

$$\mathbb{E} \left[\left\{ \frac{\sum_{k=1}^n \delta_k L(\theta_k)}{\sum_{k=1}^n \delta_k} \right\} - \min_{\mathbb{K}} L \right] \leq \tilde{E}_n / \left(\sum_{k=1}^n \delta_k \right),$$

with $(\theta_k)_{k \in \mathbb{N}}$ and L are defined in (15) and (5) respectively, and

$$E_n = D \left\{ 1 + \sum_{k=0}^{n-1} \delta_{k+1}^2 / \gamma + \sum_{k=0}^{n-1} \delta_{k+1}^2 \Lambda(\gamma) / \gamma^2 + \sum_{k=0}^{n-1} \delta_{k+1} \Psi(\gamma) + \delta_{n+1} / \gamma \right\},$$

and $D = 2R_{\mathbb{K}}^2 + 6R_{\mathbb{K}} R M p + 2R_{\mathbb{K}} \tilde{B}_2 + 3\tilde{B}_1 + B_2$ where \tilde{B}_2 is given in Lemma 23, \tilde{B}_1 in Lemma 25 and \tilde{B}_2 in Lemma 28.

The proof of this result is an adaptation to the one of [26, Theorem 4]. The main difference in the proof consists in a refinement of [26, Lemma 16] which can be established in our setting and is given Lemma 28.

Similarly to the proof of Theorem 20, we need to analyze the error η_n for $n \in \mathbb{N}$ defined by (45), but the decomposition (46) has to be improved. For that purpose, we introduce *Poisson solutions* associated with F . Under **H1** for any $\theta \in \mathbb{K}$ and $\gamma \in (0, \bar{\gamma}]$, consider $\hat{F}_{\gamma, \theta} : \mathbb{R}^d \rightarrow \mathbb{R}^p$ solution of the *Poisson equation*,

$$(\text{Id} - \mathbb{K}_{\gamma, \theta}) \hat{F}_{\gamma, \theta} = F - \pi_{\gamma, \theta}(F). \quad (48)$$

Note that by **H1**-(ii), $\hat{F}_{\gamma, \theta}$ is well defined and is given for any $x \in \mathbb{R}^d$ by

$$\hat{F}_{\gamma, \theta}(x) = \sum_{j \in \mathbb{N}} \{ \mathbb{K}_{\gamma, \theta}^j F(x) - \pi_{\gamma, \theta}(F) \}. \quad (49)$$

In addition, by Lemma 21 and **H1**-(ii), we have for any $\theta \in \mathbb{K}$ and $x \in \mathbb{R}^d$

$$\begin{aligned} \left\| \hat{F}_{\gamma, \theta}(x) \right\| &\leq \left\| \sum_{j \in \mathbb{N}} \mathbb{K}_{\gamma, \theta}^j F(x) - \pi_{\gamma, \theta}(F) \right\| \leq 3R M p \sum_{j \in \mathbb{N}} d_{W_R}(\mathbb{K}_{\gamma, \theta}^j, \pi_{\gamma, \theta}) \\ &\leq 3R M p A_2 \sum_{j \in \mathbb{N}} \rho^{\gamma j} V_R(x) \leq 3R M p A_2 \log^{-1}(1/\rho) \rho^{-\bar{\gamma}} \gamma^{-1} V_R(x) \leq C_F \gamma^{-1} V_R(x), \end{aligned} \quad (50)$$

with $C_F = 3R M p A_2 \log^{-1}(1/\rho) \rho^{-\bar{\gamma}}$. We now denote for any $n \in \mathbb{N}$, $\tilde{X}_{n+1} = X_1^n$ and therefore η_n defined by (45) is given for any $n \in \mathbb{N}$ by $\eta_n = F(\tilde{X}_{n+1}) - \pi_{\theta_n}(F)$. Using (48) an alternative expression of $(\eta_n)_{n \in \mathbb{N}}$ is given for any $n \in \mathbb{N}$ by

$$\eta_n = \hat{F}_{\gamma, \theta_n}(\tilde{X}_{n+1}) - \mathbb{K}_{\gamma, \theta_n} \hat{F}_{\gamma, \theta_n}(\tilde{X}_{n+1}) + \pi_{\gamma, \theta_n}(F) - \pi_{\theta_n}(F) = \eta_n^{(a)} + \eta_n^{(b)} + \eta_n^{(c)} + \eta_n^{(d)},$$

where

$$\begin{aligned}
\eta_n^{(a)} &= \hat{F}_{\gamma, \theta_n}(\tilde{X}_{n+1}) - \mathbf{K}_{\gamma, \theta_n} \hat{F}_{\gamma, \theta_n}(\tilde{X}_n), \\
\eta_n^{(b)} &= \mathbf{K}_{\gamma, \theta_n} \hat{F}_{\gamma, \theta_n}(\tilde{X}_n) - \mathbf{K}_{\gamma, \theta_{n+1}} \hat{F}_{\gamma, \theta_{n+1}}(\tilde{X}_{n+1}), \\
\eta_n^{(c)} &= \mathbf{K}_{\gamma, \theta_{n+1}} \hat{F}_{\gamma, \theta_{n+1}}(\tilde{X}_{n+1}) - \mathbf{K}_{\gamma, \theta_n} \hat{F}_{\gamma, \theta_n}(\tilde{X}_{n+1}), \\
\eta_n^{(d)} &= \pi_{\gamma, \theta_n}(F) - \pi_{\theta_n}(F).
\end{aligned} \tag{51}$$

In the next results, we analyze each term in this decomposition separately, except for $(\eta_n^{(a)})_{n \in \mathbb{N}}$ which is a sequence of martingale increments with respect to $(\mathcal{F}_n)_{n \in \mathbb{N}}$.

Lemma 25. *Assume **A2**(1), **B1**, **B2**(1) and **H1**. Then, for any $n \in \mathbb{N}$*

$$\begin{aligned}
\mathbb{E} \left[\left\| \sum_{k=0}^n \delta_{k+1} \langle a_{k+1}, \eta_k^{(b)} \rangle \right\| \right] \\
\leq \tilde{B}_1 \left[\sum_{k=0}^n |\delta_{k+1} - \delta_k| \gamma^{-1} + \sum_{k=0}^n \delta_k^2 \gamma^{-1} + (\delta_{n+1}/\gamma - \delta_1/\gamma) \right].
\end{aligned}$$

with $(\eta_n^{(b)})_{n \in \mathbb{N}}$ defined in (51), $a_{k+1} = \Pi_{\mathbf{K}}[\theta_k - \delta_{k+1} \nabla L(\theta_k)] - \theta^*$, $\theta^* \in \operatorname{argmin}_{\mathbf{K}} L$ and

$$\begin{aligned}
\tilde{B}_1 &= A_1 C_F (2R_{\mathbf{K}} + \sup_{\theta \in \mathbf{K}} \|\nabla L(\theta)\|) \mathbb{E} [V_R(\tilde{X}_0)] \\
&\quad + A_1 C_F (1 + \delta_1 \sup_{\theta \in \mathbf{K}} \|\nabla^2 L(\theta)\|) (\|F(0)\| + RM) \mathbb{E} [V_R^2(\tilde{X}_0)] + 4A_1 C_F R_{\mathbf{K}} \mathbb{E} [V_R(\tilde{X}_0)].
\end{aligned}$$

Proof. By (51) we have for any $n \in \mathbb{N}$ and $\theta \in \mathbf{K}$

$$\begin{aligned}
\sum_{k=0}^n \delta_{k+1} \langle a_{k+1}, \eta_k^{(b)} \rangle &= \sum_{k=1}^n \langle \delta_{k+1} a_{k+1} - \delta_k a_k, \mathbf{K}_{\gamma, \theta_k} \hat{F}_{\gamma, \theta_k}(\tilde{X}_k) \rangle \\
&\quad - \langle \delta_{n+1} a_{n+1}, \mathbf{K}_{\gamma, \theta_{n+1}} \hat{F}_{\gamma, \theta_{n+1}}(\tilde{X}_{n+1}) \rangle + \langle \delta_1 a_1, \mathbf{K}_{\gamma, \theta_0} \hat{F}_{\gamma, \theta_0}(\tilde{X}_0) \rangle.
\end{aligned} \tag{52}$$

In addition, using Proposition 4, we have for any $n \in \mathbb{N}$ and $\theta \in \mathbf{K}$,

$$\begin{aligned}
\|\delta_{n+1} a_{n+1} - \delta_n a_n\| &\leq 2R_{\mathbf{K}} |\delta_{n+1} - \delta_n| + \delta_{n+1} \|a_{n+1} - a_n\| \\
&\leq 2R_{\mathbf{K}} |\delta_{n+1} - \delta_n| + \delta_{n+1} (1 + \delta_n \sup_{\theta \in \mathbf{K}} \|\nabla^2 L(\theta)\|) \|\theta_n - \theta_{n-1}\| + \delta_{n+1} |\delta_{n+1} - \delta_n| \|\nabla L(\theta_n)\| \\
&\leq (2R_{\mathbf{K}} + \delta_1 \sup_{\theta \in \mathbf{K}} \|\nabla L(\theta)\|) |\delta_{n+1} - \delta_n| \\
&\quad + \delta_n^2 (1 + \delta_n \sup_{\theta \in \mathbf{K}} \|\nabla^2 L(\theta)\|) (\|F(0)\| + 3RM) V_R(\tilde{X}_{n+1}),
\end{aligned} \tag{53}$$

where we have used in the last inequality that $\Pi_{\mathbf{K}}$ is non-expansive, **B2**(1) and **H1**-(i) and Proposition 4 again. Combining (52), (53), (50), the Cauchy-Schwarz inequality and

H1-(i) we get that

$$\begin{aligned} \mathbb{E} \left[\left\| \sum_{k=0}^n \delta_{k+1} \langle a_k, \eta_k^{(b)} \rangle \right\|^2 \right] &\leq (2R_K + \delta_1 \sup_{\theta \in \mathbb{K}} \|\nabla L(\theta)\|) A_1 C_F \mathbb{E} \left[V_R(\tilde{X}_0) \right] \sum_{k=0}^n |\delta_{k+1} - \delta_k| \gamma^{-1} \\ &\quad + A_1 C_F (\|F(0)\| + 3RM) (1 + \delta_1 \sup_{\theta \in \mathbb{K}} \|\nabla^2 L(\theta)\|) \mathbb{E} \left[V_R^2(\tilde{X}_0) \right] \sum_{k=0}^n \delta_k^2 \gamma^{-1} \\ &\quad + 2A_1 R_K C_F \mathbb{E} \left[V_R(\tilde{X}_0) \right] \{ \delta_{n+1}/\gamma + \delta_1/\gamma \} , \end{aligned}$$

which concludes the proof of Lemma 25. \square

We now upper bound $\mathbb{E} \left[\|\eta_n^{(c)}\| \right]$ for $n \in \mathbb{N}$ using the two following lemmas.

Lemma 26. *Assume B1, B2(1), H1 and H2. Then, for any $\gamma \in (0, \bar{\gamma}]$, $\theta_1, \theta_2 \in \mathbb{K}$ and $x \in \mathbb{R}^d$*

$$d_{W_R}(\pi_{\gamma, \theta_1}, \pi_{\gamma, \theta_2}) \leq A_1 A_2 \rho^{-\bar{\gamma}} \log^{-1}(1/\rho) \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) \gamma^{-1} .$$

Proof. Let $\gamma \in (0, \bar{\gamma}]$, $\theta_1, \theta_2 \in \mathbb{K}$, $\ell \in \mathbb{N}^*$, $j \in \mathbb{N}$ with $\ell \geq j + 1$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable such that for any $y, z \in \mathbb{R}^d$, $|g(y) - g(z)| \leq W_R(y, z)$. Using **H1-(ii)** we have

$$\left| \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g(x) - \pi_{\gamma, \theta_2}(g) \right| \leq A_2 \rho^{(\ell-1-j)\gamma} V_R(x) .$$

Combining this result and **H2** we have that

$$\left| (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g(x) \right| \leq A_2 \rho^{\gamma(\ell-1-j)} \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) . \quad (54)$$

Using **H1-(i)** in (54), we get

$$\left| \mathbf{K}_{\gamma, \theta_1}^j (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g(x) \right| \leq A_1 A_2 \rho^{\gamma(\ell-1-j)} \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) .$$

Combining this result and the triangular inequality we obtain

$$\begin{aligned} \left| \mathbf{K}_{\gamma, \theta_1}^\ell g(x) - \mathbf{K}_{\gamma, \theta_2}^\ell g(x) \right| &\leq \sum_{j=0}^{\ell-1} \left| \mathbf{K}_{\gamma, \theta_1}^{j+1} \mathbf{K}_{\gamma, \theta_2}^{\ell-j-1} g(x) - \mathbf{K}_{\gamma, \theta_1}^j \mathbf{K}_{\gamma, \theta_2}^{\ell-j} g(x) \right| \\ &\leq \sum_{j=0}^{\ell-1} \left| \mathbf{K}_{\gamma, \theta_1}^j (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g(x) \right| \\ &\leq A_1 A_2 \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) \sum_{j=0}^{\ell-1} \rho^{\gamma(\ell-1-j)} \\ &\leq A_1 A_2 \rho^{-\bar{\gamma}} \log^{-1}(1/\rho) \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) \gamma^{-1} . \end{aligned}$$

Taking the limit $\ell \rightarrow +\infty$ and using **H1-(ii)** concludes the proof. \square

Lemma 27. Assume **B1**, **B2(1)**, **B4** and **H1**. Then, for any $\gamma \in (0, \bar{\gamma}]$, $\theta_1, \theta_2 \in \mathbf{K}$, $\ell \in \mathbb{N}^*$, $j \in \mathbb{N}$ with $\ell \geq j + 1$ and $x \in \mathbb{R}^d$

$$\left\| \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \pi_{\gamma, \theta_2}(F) \right\} \right\| \leq D_F \|\theta_1 - \theta_2\| V_R(x) \gamma \rho^{\gamma \ell},$$

with $D_F = A_2^2 \mathbf{B} \mathbf{M} (1 + 2R_{\mathbf{K}} \mathbf{B} \bar{\gamma}) R + 2\mathbf{M}^2$.

Proof. Let $\gamma \in (0, \bar{\gamma}]$, $\theta_1, \theta_2 \in \mathbf{K}$, $\ell \in \mathbb{N}^*$, $j \in \mathbb{N}$ with $\ell \geq j + 1$ and $x, y \in \mathbb{R}^d$. First, we have

$$\begin{aligned} & (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} \{F(x) - \pi_{\gamma, \theta_2}(F)\} - (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} \{F(x) - \pi_{\gamma, \theta_2}(F)\} \\ &= (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(y) \\ &= \mathbf{K}_{\gamma, \theta_1} \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \mathbf{K}_{\gamma, \theta_1} \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(y) + \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(y) \\ &= \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} (F(x + \Delta_\gamma(x)) - F(x) - F(y + \Delta_\gamma(y)) + F(y)) \\ &= \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} G(x) - \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} G(y), \end{aligned} \tag{55}$$

with $\Delta_\gamma(x) = \gamma(\nabla_x U(\theta_1, x) - \nabla_x U(\theta_2, x)) = \gamma \sum_{i=1}^p (\theta_1^i - \theta_2^i) \nabla F_i(x)$ and $G : \mathbb{R}^d \rightarrow \mathbb{R}^p$ defined for any $z \in \mathbb{R}^d$ by

$$G(z) = F(z + \Delta_\gamma(z)) - F(z). \tag{56}$$

Using **B2(1)** and **B4** we have that for any $x, y \in \mathbb{R}^d$,

$$\|\Delta_\gamma(x)\| \leq \mathbf{M} \gamma \|\theta_1 - \theta_2\|, \quad \|\Delta_\gamma(x) - \Delta_\gamma(y)\| \leq R_{\mathbf{K}} \mathbf{B} \gamma \|x - y\|. \tag{57}$$

Using (56), (57), we have for any $x, y \in \mathbb{R}^d$ with $x \neq y$

$$\begin{aligned} \|G(x) - G(y)\| &= \left\| \int_0^1 \{dF(x + t\Delta_\gamma(x))(\Delta_\gamma(x)) - dF(y + t\Delta_\gamma(y))(\Delta_\gamma(y))\} dt \right\| \\ &\leq \int_0^1 \|dF(x + t\Delta_\gamma(x)) - dF(y + t\Delta_\gamma(y))\| \|\Delta_\gamma(x)\| dt \\ &\quad + \int_0^1 \|dF(y + t\Delta_\gamma(y))\| (\|\Delta_\gamma(x)\| + \|\Delta_\gamma(y)\|) dt \\ &\leq \mathbf{B} (\|x - y\| + \|\Delta_\gamma(x) - \Delta_\gamma(y)\|) \|\Delta_\gamma(x)\| + 2\mathbf{M}^2 \gamma \|\theta_1 - \theta_2\| \\ &\leq \mathbf{B} \mathbf{M} (1 + 2R_{\mathbf{K}} \mathbf{B} \bar{\gamma}) \gamma \|\theta_1 - \theta_2\| \|x - y\| + 2\mathbf{M}^2 \gamma \|\theta_1 - \theta_2\| \\ &\leq \tilde{D}_F \gamma \|\theta_1 - \theta_2\| (1 + \|x - y\| / R), \end{aligned}$$

with $\tilde{D}_F = \mathbf{B} \mathbf{M} (1 + 2R_{\mathbf{K}} \mathbf{B} \bar{\gamma}) R + 2\mathbf{M}^2$. Therefore, for any $x, y \in \mathbb{R}^d$,

$$\|G(x) - G(y)\| \leq \tilde{D}_F \gamma \|\theta_1 - \theta_2\| W_R(x, y). \tag{58}$$

Combining (55), (58) and **H1-(ii)** we obtain that

$$\begin{aligned} & \left\| (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} \{F(x) - \pi_{\gamma, \theta_2}(F)\} - (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} \{F(x) - \pi_{\gamma, \theta_2}(F)\} \right\| \\ & \leq A_2 \tilde{D}_F \|\theta_1 - \theta_2\| \gamma \rho^{\gamma(\ell-j)} W_R(x, y). \end{aligned}$$

Therefore, using **H1-(ii)** we get

$$\left\| \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \pi_{\gamma, \theta_2}(F) \right\} \right\| \leq A_2^2 \tilde{D}_F \|\theta_1 - \theta_2\| V_R(x) \gamma \rho^{\gamma \ell},$$

which concludes the proof. \square

Lemma 28. *Assume **B1**, **B2(1)**, **B4**, **H1** and **H2**. Then we have for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\left\| \eta_n^{(c)} \right\| \right] \leq \tilde{B}_2 \delta_{n+1} \gamma^{-2} (\mathbf{\Lambda}(\gamma) + \gamma),$$

with

$$\tilde{B}_2 = A_1 (\|F(0)\| + 3RM) \rho^{-2\bar{\gamma}} \log^{-2}(1/\rho) \max \{D_F, E_F\}, \quad (59)$$

with $(\eta_n^{(c)})_{n \in \mathbb{N}}$ defined in (51), $E_F = 3RMpA_1A_2^2$ and D_F in Lemma 27.

Proof. We first give an upper bound on $\left\| \mathbf{K}_{\gamma, \theta_1} \hat{F}_{\gamma, \theta_1}(x) - \mathbf{K}_{\gamma, \theta_2} \hat{F}_{\gamma, \theta_2}(x) \right\|$ for any $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$. By (49) we have for any $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$,

$$\begin{aligned} & \left\| \mathbf{K}_{\gamma, \theta_1} \hat{F}_{\gamma, \theta_1}(x) - \mathbf{K}_{\gamma, \theta_2} \hat{F}_{\gamma, \theta_2}(x) \right\| \\ & = \left\| \sum_{\ell \in \mathbb{N}^*} \left\{ \mathbf{K}_{\gamma, \theta_1}^\ell F(x) - \pi_{\gamma, \theta_1}(F) \right\} - \sum_{\ell \in \mathbb{N}^*} \left\{ \mathbf{K}_{\gamma, \theta_2}^\ell F(x) - \pi_{\gamma, \theta_2}(F) \right\} \right\| \\ & \leq \sum_{\ell \in \mathbb{N}^*} \left\| \mathbf{K}_{\gamma, \theta_1}^\ell F(x) - \pi_{\gamma, \theta_1}(F) - \mathbf{K}_{\gamma, \theta_2}^\ell F(x) - \pi_{\gamma, \theta_2}(F) \right\|. \quad (60) \end{aligned}$$

We now bound each term of the series in the right hand side. For any measurable functions g_1, g_2 with $g_i : \mathbb{R}^d \rightarrow \mathbb{R}^p$ and such that $\sup_{x \in \mathbb{R}^d} \|g_i(x)\| / V_R(x) < +\infty$ with $i \in \{1, 2\}$, $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$, $x \in \mathbb{R}^d$ and $\ell \in \mathbb{N}^*$, using that π_{γ, θ_1} is invariant for $\mathbf{K}_{\gamma, \theta_1}$, it holds that

$$\begin{aligned} \mathbf{K}_{\gamma, \theta_1}^\ell g_1(x) - \mathbf{K}_{\gamma, \theta_2}^\ell g_2(x) &= \mathbf{K}_{\gamma, \theta_1}^\ell g_1(x) - \mathbf{K}_{\gamma, \theta_2}^\ell g_1(x) + \mathbf{K}_{\gamma, \theta_2}^\ell (g_1(x) - g_2(x)) \\ &= \sum_{j=0}^{\ell-1} \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g_1(x) - \pi_{\gamma, \theta_2}(g_1) \right\} \\ & \quad + \sum_{j=0}^{\ell-1} \pi_{\gamma, \theta_1} \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g_1(x) - \mathbf{K}_{\gamma, \theta_2}^{\ell-j} g_1(x) \right\} + \mathbf{K}_{\gamma, \theta_2}^\ell (g_1(x) - g_2(x)) \\ &= \sum_{j=0}^{\ell-1} \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} g_1(x) - \pi_{\gamma, \theta_2}(g_1) \right\} \\ & \quad - \pi_{\gamma, \theta_1} (\mathbf{K}_{\gamma, \theta_2}^\ell g_1(x) - g_1(x)) + \mathbf{K}_{\gamma, \theta_2}^\ell (g_1(x) - g_2(x)). \end{aligned}$$

Setting $g_1 = F - \pi_{\gamma, \theta_1}(F)$ and $g_2 = F - \pi_{\gamma, \theta_2}(F)$, we obtain that

$$\begin{aligned} & \mathbf{K}_{\gamma, \theta_1}^\ell F(x) - \pi_{\gamma, \theta_1}(F) - \mathbf{K}_{\gamma, \theta_2}^\ell F(x) - \pi_{\gamma, \theta_2}(F) \\ &= \sum_{j=0}^{\ell-1} \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \pi_{\gamma, \theta_2}(F) \right\} + \Xi_\ell, \end{aligned} \quad (61)$$

where, using that π_{γ, θ_2} is invariant for $\mathbf{K}_{\gamma, \theta_2}$, we have

$$\begin{aligned} \Xi_\ell &= -\pi_{\gamma, \theta_1}(\mathbf{K}_{\gamma, \theta_2}^\ell F(x) - F(x)) + \mathbf{K}_{\gamma, \theta_2}^\ell [\pi_{\gamma, \theta_2}(F) - \pi_{\gamma, \theta_1}(F)] \\ &= (\pi_{\gamma, \theta_2} - \pi_{\gamma, \theta_1}) \mathbf{K}_{\gamma, \theta_2}^\ell F(x). \end{aligned}$$

Using Lemma 27 we obtain for any $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$, $x \in \mathbb{R}^d$ and $\ell \in \mathbb{N}^*$

$$\begin{aligned} & \left\| \sum_{j=0}^{\ell-1} \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j} F(x) - \pi_{\gamma, \theta_2}(F) \right\} \right\| \\ & \leq \sum_{j=0}^{\ell-1} \left\| \left\{ \mathbf{K}_{\gamma, \theta_1}^j - \pi_{\gamma, \theta_1} \right\} (\mathbf{K}_{\gamma, \theta_1} - \mathbf{K}_{\gamma, \theta_2}) \left\{ \mathbf{K}_{\gamma, \theta_2}^{\ell-1-j}(F) - \pi_{\gamma, \theta_2}(F) \right\} \right\| \\ & \leq \sum_{j=0}^{\ell-1} D_F \gamma \|\theta_1 - \theta_2\| \rho^{\gamma \ell} V_R(x) \leq D_F \gamma V_R(x) \|\theta_1 - \theta_2\| \ell \rho^{\gamma \ell}. \end{aligned} \quad (62)$$

Using **H1-(ii)**, Lemma 21 and Lemma 26, we obtain for any $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$, $x \in \mathbb{R}^d$ and $\ell \in \mathbb{N}^*$

$$\begin{aligned} \left\| (\pi_{\gamma, \theta_1} - \pi_{\gamma, \theta_2}) \mathbf{K}_{\gamma, \theta_2}^\ell F(x) \right\| & \leq 3RMpA_2 \rho^{\gamma \ell} d_{W_R}(\pi_{\gamma, \theta_1}, \pi_{\gamma, \theta_2}) \\ & \leq E_F \rho^{-\bar{\gamma}} \log^{-1}(1/\rho) \mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| V_R^2(x) \gamma^{-1} \rho^{\gamma \ell}. \end{aligned} \quad (63)$$

with $E_F = 3RMpA_1A_2^2$. Combining (62) and (63) in (61), we obtain that for any $\theta_1, \theta_2 \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$ that

$$\begin{aligned} & \mathbf{K}_{\gamma, \theta_1}^\ell F(x) - \pi_{\gamma, \theta_1}(F) - \mathbf{K}_{\gamma, \theta_2}^\ell F(x) - \pi_{\gamma, \theta_2}(F) \\ & \leq D_F \|\theta_1 - \theta_2\| V_R(x) \gamma \ell \rho^{\gamma \ell} + E_F \rho^{-\bar{\gamma}} \log^{-1}(1/\rho) \|\theta_1 - \theta_2\| V_R^2(x) \mathbf{\Lambda}(\gamma) \gamma^{-1} \rho^{\gamma \ell}. \end{aligned}$$

Using this result in (60) and that for any $t \in (-1, 1)$ and $a > 0$, $\sum_{k \in \mathbb{N}} kt^{ak} = t(1 - t^a)^{-2} \leq a^{-2} t^{-a} \log^{-2}(1/t)$, we get that

$$\begin{aligned} \left\| \mathbf{K}_{\gamma, \theta_1} \hat{F}_{\gamma, \theta_1}(x) - \mathbf{K}_{\gamma, \theta_2} \hat{F}_{\gamma, \theta_2}(x) \right\| & \leq D_F \rho^{-2\bar{\gamma}} \log^{-2}(1/\rho) \|\theta_1 - \theta_2\| V_R(x) \gamma^{-1} \\ & \quad + E_F \rho^{-2\bar{\gamma}} \log^{-2}(1/\rho) \|\theta_1 - \theta_2\| V_R^2(x) \mathbf{\Lambda}(\gamma) \gamma^{-2} \\ & \leq C_c \gamma^{-2} (\mathbf{\Lambda}(\gamma) \|\theta_1 - \theta_2\| + \gamma \|\theta_1 - \theta_2\|) V_R^2(x), \end{aligned} \quad (64)$$

with $C_c = \rho^{-2\bar{\gamma}} \log^{-2}(1/\rho) \max\{D_F, E_F\}$. Note that for any $k \in \mathbb{N}$, by **B2**(1) and the fact that Π_K is non-expansive we have $\|\theta_{k+1} - \theta_k\| \leq \delta_{k+1}(\|F(0)\| + 3RM)V_R(\tilde{X}_{k+1})$. Therefore, plugging this result in (64), we get for any $k \in \mathbb{N}$,

$$\begin{aligned} & \left\| K_{\gamma, \theta_k} \hat{F}_{\gamma, \theta_k}(\tilde{X}_{k+1}) - K_{\gamma, \theta_{k+1}} \hat{F}_{\gamma, \theta_{k+1}}(\tilde{X}_{k+1}) \right\| \\ & \leq C_{c,2}(\|F(0)\| + 3RM)\delta_{k+1}\gamma^{-2}(\Lambda(\gamma) + \gamma)V_R^3(\tilde{X}_{k+1}). \end{aligned} \quad (65)$$

Therefore by definition of $(\eta_k^{(c)})_{k \in \mathbb{N}}$, see (51), and using **H1**-(i) in (65) we get that for any $k \in \mathbb{N}$

$$\mathbb{E} \left[\left\| \eta_k^{(c)} \right\| \right] \leq \tilde{B}_2 \delta_{k+1} \gamma^{-2} (\Lambda(\gamma) + \gamma),$$

with \tilde{B}_2 given by (59). □

Lemma 29. *Assume **B1**, **B2**(1) and **H1**. Then we have for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\left\| \eta_n^{(d)} \right\| \right] \leq 3RMp\Psi(\gamma),$$

with $(\eta_n^{(d)})_{n \in \mathbb{N}}$ defined in (51).

Proof. The proof is a direct consequence of Lemma 21 and **H1**-(iii). □

Proof of Theorem 24. Taking the expectation in [4, Theorem 3, Equation (8)], using the Cauchy-Schwarz inequality, the decomposition of the error (51) and the fact that $(\eta_n^{(a)})_{n \in \mathbb{N}}$ is a martingale increment with respect to $(\mathcal{F}_n)_{n \in \mathbb{N}}$, we get that for every $n \in \mathbb{N}$

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^n \delta_k \left\{ L(\theta_k) - \min_K L \right\} \right] \\ & \leq \mathbb{E} \left[2R_K^2 - \sum_{k=0}^{n-1} \delta_{k+1} \langle \Pi_K(\theta_k - \delta_{k+1} \nabla L(\theta_k)) - \theta^*, \eta_k \rangle + \sum_{k=0}^{n-1} \delta_{k+1}^2 \|\eta_k\|^2 \right] \\ & \leq 2R_K^2 + 2R_K \sum_{k=0}^{n-1} \delta_{k+1} \mathbb{E} \left[\left\| \eta_k^{(c)} \right\| + \left\| \eta_k^{(d)} \right\| \right] + \mathbb{E} \left[\sum_{k=0}^{n-1} \delta_{k+1} \langle a_{k+1}, \eta_k^{(b)} \rangle \right] + \sum_{k=0}^{n-1} \delta_{k+1}^2 \mathbb{E} \left[\|\eta_k\|^2 \right]. \end{aligned}$$

Combining this result, Lemma 25, Lemma 28, Lemma 29 and Lemma 23 completes the proof. □

B.2.3 Proof of Theorem 9

In this section, we check that **H1** and **H2** are satisfied in order to apply Theorem 24. More precisely, we study the geometric ergodicity of the Langevin Markov chain under **B1**, **B2**(1) and **B3** as well as its discretization error. We begin with the following technical lemma

Lemma 30. Assume **B1**, **B2(1)** and **B3**. Let $\mathfrak{m} = \mathfrak{m}_1/2$, $\tilde{\mathfrak{N}} = 2\mathfrak{N}$, $R = 4\mathcal{M}/\mathfrak{m}_1$ and $\mathfrak{v} = \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\|$. In addition, for any $\theta \in \mathbf{K}$, $\gamma > 0$ and $x \in \mathbb{R}^d$, let

$$T_{\gamma, \theta}(x) = \|x - \gamma \nabla_x U(\theta, x)\|^2 . \quad (66)$$

Then for any $\theta \in \mathbf{K}$ and $x, y \in \mathbb{R}^d$

- (a) $\|\nabla_x U(\theta, x) - \nabla_x U(\theta, y)\| \leq \tilde{\mathfrak{N}} \|x - y\|$;
- (b) if $\|x - y\| \geq R$, $\langle \nabla_x U(\theta, x) - \nabla_x U(\theta, y), x - y \rangle \geq \mathfrak{m} \|x - y\|^2$,
- (c) we have

$$\|T_{\gamma, \theta}(x)\| \leq (1 + \gamma \tilde{\mathfrak{N}}) \|x\| + \gamma \mathfrak{v} ,$$
- (d) if $\|x\| \geq \max(R, 2\mathfrak{v}/\mathfrak{m})$ and $\gamma \leq \mathfrak{m}/(2\tilde{\mathfrak{N}}^2)$

$$\|T_{\gamma, \theta}(x)\| \leq (1 - \gamma \mathfrak{m}/2 + \gamma^2 \tilde{\mathfrak{N}}^2/2) \|x\| .$$

Proof. Let $\theta \in \mathbf{K}$. The proof of (a) is straightforward. Let $\mathfrak{m} = \mathfrak{m}_1/2$ and $x, y \in \mathbb{R}^d$ such that $\|x - y\| \geq R$ with $R = 4\mathcal{M}/\mathfrak{m}_1$. Using **B3-(b)-(c)** we have

$$\begin{aligned} & \langle \nabla_x U(\theta, x) - \nabla_x U(\theta, y), x - y \rangle \\ &= \langle \nabla_x U_1(\theta, x) - \nabla_x U_1(\theta, y), x - y \rangle + \langle \nabla_x U_2(\theta, x) - \nabla_x U_2(\theta, y), x - y \rangle \\ &\geq \mathfrak{m}_1 \|x - y\|^2 - 2\mathcal{M} \|x - y\| \geq (\mathfrak{m}_1 - 2\mathcal{M}/R) \|x - y\|^2 \\ &\geq (\mathfrak{m}_1/2) \|x - y\|^2 , \end{aligned}$$

which concludes the proof of (b). For any $x \in \mathbb{R}^d$,

$$\|T_{\gamma, \theta}(x)\| \leq \|T_{\gamma, \theta}(x) - T_{\gamma, \theta}(0)\| + \|T_{\gamma, \theta}(0)\| \leq (1 + \gamma \tilde{\mathfrak{N}}) \|x\| + \gamma \mathfrak{v} \leq (1 + \gamma(\tilde{\mathfrak{N}} + \mathfrak{v})) (1 + \|x\|) - 1 ,$$

and therefore (c) holds. Finally, let $\|x\| \geq \max(R, 2\mathfrak{v}/\mathfrak{m})$ and $\gamma \leq \mathfrak{m}/(2\tilde{\mathfrak{N}}^2)$. Using that for any $t \geq 0$, $\sqrt{1+t} \leq 1 + t/2$,

$$\begin{aligned} \|T_{\gamma, \theta}(x)\| &\leq \|T_{\gamma, \theta}(x) - T_{\gamma, \theta}(0)\| + \gamma \|\nabla_x U(\theta, 0)\| \\ &\leq (1 - \gamma \mathfrak{m} + \gamma^2 \tilde{\mathfrak{N}}^2/2) \|x\| + \gamma \mathfrak{v} \leq (1 - \gamma \mathfrak{m}/2 + \gamma^2 \tilde{\mathfrak{N}}^2/2) \|x\| , \end{aligned}$$

which concludes the proof of (d). \square

Lemma 31. Assume **B1**, **B2(1)** and **B3**. Let $\mathfrak{m}, \tilde{\mathfrak{N}}$ and R be given by Lemma 30. Then for any $m \in \mathbb{N}^*$, $\theta \in \mathbf{K}$ and $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}/(2\tilde{\mathfrak{N}}^2), 1/2)$, $K_{\gamma, \theta}$ satisfies $\mathbf{D}_d(\mathcal{V}_m, \lambda^\gamma, \tilde{b}_m \gamma)$

with \mathcal{V}_m given in (44) and

$$\begin{aligned}
\lambda &= \exp[-\mathfrak{m}/4 + \bar{\gamma}\tilde{\mathfrak{N}}^2/2] , \\
\tilde{b}_m &= \Upsilon_m(2^{2m+1}d^m(1 + \bar{\gamma}\tilde{\mathfrak{N}})^{2m-1}\Gamma(m + 1/2), \mathfrak{m}/4) + \mathfrak{m}/4 + e^{\kappa\bar{\gamma}}(\kappa + \log(1/\lambda))\mathcal{V}_m(\tilde{R}) + C_m(\tilde{R}) , \\
\kappa &= \left\{ (1 + \bar{\gamma}\tilde{\mathfrak{N}}) \max(\tilde{R}, 1) + \mathfrak{v} \right\}^{2m} , \\
\mathfrak{v} &= \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\| , \quad \tilde{R} = \max(R, 2\mathfrak{v}/\mathfrak{m}) , \\
C_m(\tilde{R}) &= 2^{2m+1}d^m \{1 + \bar{\gamma}(\tilde{\mathfrak{N}} + \mathfrak{v})\}^{2m-1} \Gamma(m + 1/2)(1 + \tilde{R})^{2m-1} ,
\end{aligned} \tag{67}$$

where for any $t \geq 0$, $\Gamma(t) = \int_0^{+\infty} u^{t-1}e^{-u}du$ and Υ_m is given in Lemma 14. In addition, $\mathbf{K}_{\gamma,\theta}$ satisfies $\mathbf{D}_d(V, \lambda^\gamma, b_m(1 + d^{\varpi_{0,m}})\gamma)$ with λ given in (67) and $b_m, \varpi_{0,m} \geq 0$ independent of the dimension d .

Proof. First, note that using **B2(1)-(a)** we get

$$\mathfrak{v} \leq \left\| \sum_{i=1}^p \theta_i \nabla F_i(0) + \nabla r(0) \right\| \leq \|\nabla r(0)\| + pR_{\mathbf{K}} \sup_{i \in \{1, \dots, p\}} \|\nabla F_i(0)\| < +\infty .$$

Let $m \in \mathbb{N}^*$, $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$. Similarly to Lemma 15, we obtain that

$$\begin{aligned}
\int_{\mathbb{R}^d} \|y\|^{2m} \mathbf{K}_{\gamma,\theta}(x, dy) &\leq \|T_{\gamma,\theta}(x)\|^{2m} + \gamma 2^{2m+1} \gamma d^m \Gamma(m + 1/2)(1 + \|T_{\gamma,\theta}(x)\|)^{2m-1} \\
&\leq \|T_{\gamma,\theta}(x)\|^{2m} + \gamma 2^{2m+1} \gamma C_m(x) ,
\end{aligned} \tag{68}$$

with $T_{\gamma,\theta}$ defined in (66) and

$$C_m(x) = 2^{2m+1}d^m \{1 + \bar{\gamma}(\tilde{\mathfrak{N}} + \mathfrak{v})\}^{2m-1} \Gamma(m + 1/2)(1 + \|x\|)^{2m-1} , \tag{69}$$

where we have used Lemma 30-(c). Let $\tilde{R} = \max(R, 2\mathfrak{v}/\mathfrak{m})$. We divide the rest of the proof in two parts:

(a) Let $\|x\| \geq \tilde{R}$. We have using Lemma 30-(d),

$$\|T_{\gamma,\theta}(x)\| \leq (1 - \gamma\mathfrak{m}/2 + \gamma^2\tilde{\mathfrak{N}}^2/2) \|x\| .$$

Hence, $\|T_{\gamma,\theta}(x)\|^{2m} \leq (1 - \gamma\mathfrak{m}/4 + \gamma^2\tilde{\mathfrak{N}}^2/2) \|x\|^{2m} - \gamma\mathfrak{m} \|x\|^{2m}/4$ and we have using Lemma 14 and (69) in (68)

$$\begin{aligned}
\mathbf{K}_{\gamma,\theta}\mathcal{V}_m(x) &\leq (1 - \gamma\mathfrak{m}/4 + \gamma^2\tilde{\mathfrak{N}}^2/2)(1 + \|x\|^{2m}) \\
&\quad + \gamma \left[2^{2m+1}d^m \{1 + \bar{\gamma}(\tilde{\mathfrak{N}} + \mathfrak{v})\}^{2m-1} \Gamma(m + 1/2)(1 + \|x\|)^{2m-1} - \mathfrak{m} \|x\|^{2m}/4 + \mathfrak{m}/4 \right] \\
&\leq (1 - \gamma\mathfrak{m}/4 + \gamma^2\tilde{\mathfrak{N}}^2/2)(1 + \|x\|^{2m}) \\
&\quad + \gamma \left[\Upsilon_m(2^{2m+1}d^m \{1 + \bar{\gamma}(\tilde{\mathfrak{N}} + \mathfrak{v})\}^{2m-1} \Gamma(m + 1/2), \mathfrak{m}/4) + \mathfrak{m}/4 \right] .
\end{aligned}$$

(b) Now assume that $\|x\| \leq \tilde{R}$. Let $\kappa = \{(1 + \tilde{N}) \max(1, \tilde{R}) + \mathbf{v}\}^{2m}$. We have, using that $\gamma \leq 1$,

$$(1 + \gamma \tilde{N})^{2m} \leq 1 + \gamma \sum_{k=1}^{2m} \binom{2m}{k} \tilde{N}^k \leq 1 + \gamma(1 + \tilde{N})^{2m} \leq 1 + \gamma \kappa.$$

Combining this result with Lemma 30-(c) and the fact that $\bar{\gamma} \leq 1$, we get

$$\begin{aligned} 1 + \|T_{\gamma, \theta}(x)\|^{2m} &\leq 1 + [(1 + \gamma \tilde{N}) \|x\| + \gamma \mathbf{v}]^{2m} \\ &\leq 1 + (1 + \gamma \tilde{N})^{2m} \|x\|^{2m} + \gamma \sum_{k=1}^{2m} \binom{2m}{k} (1 + \gamma \tilde{N})^{2m-k} \tilde{R}^{2m-k} \mathbf{v}^k \\ &\leq 1 + (1 + \gamma \kappa) \|x\|^{2m} + \gamma \kappa \leq (1 + \gamma \kappa)(1 + \|x\|^{2m}). \end{aligned} \quad (70)$$

Let

$$C_m(\tilde{R}) = 2^{2m+1} d^m \{1 + \bar{\gamma}(\tilde{N} + \mathbf{v})\}^{2m-1} \mathbf{\Gamma}(m + 1/2)(1 + \tilde{R})^{2m-1}.$$

Using (70) in (68) and that for any $a \geq b$, $e^a - e^b \leq (a - b)e^a$ we have

$$\begin{aligned} \mathbf{K}_{\gamma, \theta}(1 + \|x\|^{2m}) &\leq 1 + \|T_{\gamma, \theta}(x)\|^{2m} + \gamma C_m(\tilde{R}) \\ &\leq \lambda^\gamma \mathcal{V}_m(x) + \gamma e^{\kappa \bar{\gamma}} (\kappa + \log(1/\lambda)) \mathcal{V}_m(\tilde{R}) + \gamma C_m(\tilde{R}), \end{aligned}$$

which concludes the proof upon noting that \tilde{b}_m is a polynomial in the dimension d . □

Lemma 32. *Assume **B1**, **B2**(1), **B3** and let $(X_k^n)_{n \in \mathbb{N}, k \in \{0, \dots, m_n\}}$ be given by (15) with $\bar{\gamma} < \min(\mathbf{m}/(2\tilde{N}^2), 1/2)$. There exist $A_1 \geq 1$ and $\varpi_1 \geq 0$ such that for any $a \in [1, 3]$, $n, p \in \mathbb{N}$ and $k \in \{0, \dots, m_n\}$*

$$\mathbb{E} \left[\mathbf{K}_{\gamma_n, \theta_n}^p V_R^a(X_k^n) \middle| X_0^0 \right] \leq A_1 V_R^a(X_0^0), \quad \mathbb{E} \left[V_R^a(X_0^0) \right] < +\infty,$$

with V_R given in (43) and A_1, ϖ_1 which do not depend on the dimension d .

Proof. Using Jensen's inequality it suffices to prove the result for $a = 3$. Using Lemma 31, there exist $\lambda \in (0, 1)$ and $b \geq 0$ such that for any $m \in \{1, 2, 3\}$, $\theta \in \mathbf{K}$ and $\gamma \in (0, \bar{\gamma}]$, $\mathbf{K}_{\gamma, \theta}$ satisfies $\mathbf{D}_d(\mathcal{V}_m, \lambda^\gamma, b_m(1 + d^{\varpi_0, m})\gamma)$ with \mathcal{V}_m given in (44). Hence, since λ and b_m do not depend on the dimension d , using [26, Lemma S15], there exists $\tilde{A}_1, \varpi_1 \geq 0$ such that for any $m \in \{1, 2, 3\}$

$$\mathbb{E} \left[\mathbf{K}_{\gamma_n, \theta_n}^p \mathcal{V}_m(X_k^n) \middle| X_0^0 \right] \leq \tilde{A}_1 (1 + d^{\varpi_1}) \mathcal{V}_m(X_0^0), \quad \mathbb{E} \left[\mathcal{V}_m(X_0^0) \right] < +\infty,$$

with \tilde{A}_1 and ϖ_1 which do not depend on the dimension d . Combining this result and Jensen's inequality we obtain that

$$\begin{aligned}
\mathbb{E} \left[\mathbf{K}_{\gamma_n, \theta_n}^p V_R^3(\mathbf{X}_k^n) \middle| \mathbf{X}_0^0 \right] &\leq R^{-m} \sum_{m=0}^3 \binom{3}{m} \mathbb{E} \left[\mathbf{K}_{\gamma_n, \theta_n}^p \mathcal{V}_m(\mathbf{X}_k^n) \middle| \mathbf{X}_0^0 \right]^{1/2} \\
&\leq R^{-m} \sum_{m=0}^3 \binom{3}{m} \tilde{A}_1^{1/2} (1 + d^{\varpi_1})^{1/2} (1 + \|\mathbf{X}_0^0\|^{2m})^{1/2} \\
&\leq R^{-m} \tilde{A}_1^{1/2} (1 + d^{\varpi_1})^{1/2} \sum_{m=0}^3 \binom{3}{m} (1 + \|\mathbf{X}_0^0\|^m) \\
&\leq 9 \tilde{A}_1^{1/2} (1 + d^{\varpi_1})^{1/2} (1 + \|\mathbf{X}_0^0\|/R)^3,
\end{aligned}$$

which concludes the proof. \square

Theorem 33. *Assume **B1**, **B2**(1) and **B3**. Then there exist $A_2, \varpi_2 \geq 0$ and $\rho \in (0, 1)$ such that for any $\theta \in \mathbf{K}$ and $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(m/(2\tilde{\mathbf{N}}^2), 1/2)$, $\mathbf{K}_{\gamma, \theta}$ admits an invariant probability measure $\pi_{\gamma, \theta}$ and for any $n \in \mathbb{N}$ and $x \in \mathbb{R}^d$*

$$d_{W_R}(\delta_x \mathbf{K}_{\gamma, \theta}^n, \pi_{\gamma, \theta}) \leq A_2 (1 + d^{\varpi_2}) \rho^{\gamma n} V_R(x), \quad d_{W_R}(\delta_x \mathbf{K}_{\gamma, \theta}^n, \delta_y \mathbf{K}_{\gamma, \theta}^n) \leq A_2 (1 + d^{\varpi_2}) \rho^{\gamma n} W_R(x, y),$$

with V_R, W_R given in (43) and $A_2, \varpi_2 \geq 0$ and $\rho \in (0, 1)$ which do not depend on the dimension d .

Proof. Let $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$, $n \in \mathbb{N}$ and $x, y \in \mathbb{R}^d$. Applying [25, Theorem 10], we obtain that there exist $\tilde{A}_2 \geq 0$ and $\rho \in (0, 1)$ independent of the dimension d such that

$$d_{W_R}(\delta_x \mathbf{K}_{\gamma, \theta}^n, \delta_y \mathbf{K}_{\gamma, \theta}^n) \leq \tilde{A}_2 \rho^{\gamma n} W_R(x, y).$$

In addition, $\mathbf{K}_{\gamma, \theta}$ admits an invariant probability measure $\pi_{\gamma, \theta}$. Hence, we have

$$d_{W_R}(\delta_x \mathbf{K}_{\gamma, \theta}^n, \pi_{\gamma, \theta}) \leq \tilde{A}_2 \rho^{\gamma n} \int_{\mathbb{R}^d} W_R(x, y) d\pi_{\gamma, \theta}(y). \quad (71)$$

By Lemma 31 we get that $\mathbf{K}_{\gamma, \theta}$ satisfies $\mathbf{D}_d(\mathcal{V}_2, \lambda^\gamma, b_2(1 + d^{\varpi_{0,2}})\gamma)$ where $\mathcal{V}_2(x)$ is given by (44) with $m = 2$, λ, b_2 and $\varpi_{0,2}$ are independent of the dimension d . Hence, using Jensen's inequality and that $\sqrt{1+t} \leq 1 + t/2$ for $t \geq 0$ we get

$$\pi_{\gamma, \theta}(\mathcal{V}_2^{1/2}) \leq \lambda^{\gamma/2} \pi_{\gamma, \theta}(\mathcal{V}_2^{1/2}) + b_2(1 + d^{\varpi_{0,2}}) \lambda^{-\bar{\gamma}/2} \gamma/2.$$

Therefore we obtain that

$$\pi_{\gamma, \theta}(\mathcal{V}_2^{1/2}) \leq b_2(1 + d^{\varpi_{0,2}}) \lambda^{-\bar{\gamma}/2} \gamma / (2 - 2\lambda^{\gamma/2}) \leq b_2(1 + d^{\varpi_{0,2}}) \lambda^{-\bar{\gamma}} \log^{-1}(1/\lambda). \quad (72)$$

In addition, for any $y \in \mathbb{R}^d$,

$$1 + \|y\|/R \leq 2^{1/2}(1 + \|y\|^2/R^2)^{1/2} \leq 2^{1/2}(1 + 1/R^2)^{1/2}\mathcal{V}_2^{1/2}(y). \quad (73)$$

Combining (71), (72), (73) and Jensen's inequality we get that

$$\begin{aligned} d_{W_R}(\delta_x K_{\gamma,\theta}^n, \pi_{\gamma,\theta}) &\leq \tilde{A}_2 \rho^{\gamma n} \int_{\mathbb{R}^d} W_R(x, y) d\pi_{\gamma,\theta}(y) \\ &\leq \tilde{A}_2 \rho^{\gamma n} \{ \|x\|/R + \pi_{\gamma,\theta}(V_R) \} \\ &\leq \tilde{A}_2 \rho^{\gamma n} \left\{ \|x\|/R + 2^{1/2}(1 + 1/R^2)^{1/2} b_2 (1 + d^{\varpi_{0,2}}) \lambda^{-\bar{\gamma}} / \log(1/\lambda) \right\} \\ &\leq \tilde{A}_2 \left\{ 1 + 2^{1/2}(1 + 1/R^2)^{1/2} b_2 (1 + d^{\varpi_{0,2}}) \lambda^{-\bar{\gamma}} / \log(1/\lambda) \right\} V_R(x), \end{aligned}$$

which concludes the proof. \square

Lemma 34. *Assume B1, B2(1) and B3. Then there exists $\tilde{A}_3, \varpi'_3 \geq 0$ such that for any $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}/(2\tilde{\mathfrak{N}}^2), 1/2)$ and $k \in \mathbb{N}$*

$$\| \pi_\theta P_{k\gamma[1/\gamma], \theta} P_{\gamma[1/\gamma], \theta} - \pi_\theta P_{k\gamma[1/\gamma], \theta} K_{\gamma,\theta}^{\lceil 1/\gamma \rceil} \|_{V_R} \leq \tilde{A}_3 (1 + d^{\varpi'_3}) \gamma^{1/2},$$

with V_R given in (43) and $\tilde{A}_3, \varpi'_3 \geq 0$ which do not depend on the dimension d .

Proof. Let $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$. First, we show that $(P_{t,\theta})_{t \geq 0}$ satisfies a drift condition $\mathbf{D}_c(\mathcal{V}_2, \zeta, \beta_m(1 + d^{\varpi'_{0,m}}))$, with $\mathcal{V}_2(x) = 1 + \|x\|^2$, $\zeta > 0$ and $\beta_m, \varpi'_{0,m} \geq 0$ independent of the dimension d . We have that for any $x \in \mathbb{R}^d$, $\nabla \mathcal{V}_2(x) = 2x$ and $\Delta \mathcal{V}_2(x) = 2d$. Hence, for any $x \in \mathbb{R}^d$

$$\mathcal{A}_\theta \mathcal{V}_2(x) = -\langle \nabla_x U(\theta, x), \nabla \mathcal{V}_2(x) \rangle + \Delta \mathcal{V}_2(x) = -2\langle \nabla_x U(\theta, x), x \rangle + 2d.$$

We now distinguish two cases.

(a) If $\|x\| \geq R$, using Lemma 30-(b) we have

$$\begin{aligned} \mathcal{A}_\theta \mathcal{V}_2(x) &\leq -2\mathfrak{m} \|x\|^2 + 2d + 2 \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\| \|x\| \\ &\leq -\mathfrak{m} \mathcal{V}_2(x) + 2\{d + \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\| \|x\| - \mathfrak{m} \|x\|^2 / 2 + \mathfrak{m}\} \\ &\leq -\mathfrak{m} \mathcal{V}_2(x) + 2\{d + \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\|^2 / (2\mathfrak{m}) + \mathfrak{m}\}. \end{aligned}$$

(b) If $\|x\| \leq R$, using Lemma 30-(a) we have

$$\begin{aligned} \mathcal{A}_\theta \mathcal{V}_2(x) &\leq 2(\tilde{\mathfrak{N}} \|x\| + \sup_{\theta} \|\nabla_x U(\theta, x)\|) \|x\| + 2d \\ &\leq -\mathfrak{m} \mathcal{V}_2(x) + 2(\tilde{\mathfrak{N}} R + \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, x)\|) R + 2d + \mathfrak{m} \mathcal{V}_2(R). \end{aligned}$$

Hence, there exists $\zeta > 0$ and $\beta_m, \varpi'_{0,m} \geq 0$ such that $(P_{t,\theta})_{t \geq 0}$ satisfies $\mathbf{D}_c(\mathcal{Y}_2, \zeta, \beta_m(1 + d^{\varpi'_{0,m}}))$, with ζ, β_m and $\varpi'_{0,m}$ independent of the dimension d . This implies by [68, Theorem 4.5]

$$\pi_\theta(\mathcal{Y}_2) \leq \beta_m(1 + d^{\varpi'_{0,m}})/\zeta. \quad (74)$$

Using a generalized Pinsker inequality [32, Lemma 24], [32, Equation 15] and that for any $y \in \mathbb{R}^d$, $V_R(y) \leq (1 + 1/R^2)^{1/2}\mathcal{Y}_2(y)$, we get that

$$\begin{aligned} & \|\pi_\theta P_{k\gamma[1/\gamma],\theta} P_{\gamma[1/\gamma],\theta} - \pi_\theta P_{k\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{[1/\gamma]}\|_{V_R} \\ & \leq 2(1 + 1/R^2)^{1/2} (\pi_\theta P_{(k+1)\gamma[1/\gamma]} \mathcal{Y}_2 + \pi_\theta P_{k\gamma[1/\gamma]} K_{\gamma,\theta}^{[1/\gamma]} \mathcal{Y}_2)^{1/2} \\ & \quad \times \text{KL} \left(\pi_\theta P_{k\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{[1/\gamma]} \mid \pi_\theta P_{k\gamma[1/\gamma],\theta} P_{\gamma[1/\gamma],\theta} \right)^{1/2} \\ & \leq (1 + 1/R^2)^{1/2} (\pi_\theta(\mathcal{Y}_2) + \pi_\theta K_{\gamma,\theta}^{[1/\gamma]} \mathcal{Y}_2)^{1/2} \\ & \quad \times \tilde{\mathbf{N}} \left(2\tilde{\mathbf{N}}\bar{\gamma} \sup_{j \in \mathbb{N}} \left\{ \pi_\theta K_{\gamma,\theta}^{[1/\gamma]j} \mathcal{Y}_2 \right\} + 2\bar{\gamma} \sup_{\theta \in \mathbf{K}} \|\nabla_x U(\theta, 0)\|^2 + d \right)^{1/2}. \end{aligned}$$

Combining this result, (74) and Lemma 31 completes the proof. \square

Proposition 35. *Assume **B1**, **B2**(1) and **B3**. Then there exist $A_3, \varpi_3 \geq 0$ such that for any $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}/(2\tilde{\mathbf{N}}^2), 1/2)$,*

$$d_{W_R}(\pi_{\gamma,\theta}, \pi_\theta) \leq A_3(1 + d^{\varpi_3})\gamma^{1/2},$$

with W_R given in (43) and $A_3, \varpi_3 \geq 0$ which do not depend on the dimension d .

Proof. Let $\theta \in \mathbf{K}$, $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$. Using [25, Theorem 10], we get that

$$d_{W_R}(\pi_{\gamma,\theta}, \pi_\theta) = \lim_{n \rightarrow +\infty} d_{W_R}(\pi_\theta K_{\gamma,\theta}^{n[1/\gamma]}, \pi_\theta).$$

By Theorem 33, Lemma 34 and that for any $\theta \in \mathbf{K}$, π_θ is an invariant probability measure for $(P_{t,\theta})_{t \geq 0}$, we get for any $n \in \mathbb{N}$

$$\begin{aligned} d_{W_R}(\pi_\theta K_{\gamma,\theta}^{n[1/\gamma]}, \pi_\theta P_{n\gamma[1/\gamma],\theta}) & \leq \sum_{k=0}^{n-1} d_{W_R}(\pi_\theta P_{(k+1)\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{(n-k-1)[1/\gamma]}, \pi_\theta P_{k\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{(n-k)[1/\gamma]}) \\ & \leq A_2(1 + d^{\varpi_2}) \sum_{k=0}^{n-1} \rho^{n-k-1} d_{W_R}(\pi_\theta P_{k\gamma[1/\gamma],\theta} P_{\gamma[1/\gamma],\theta}, \pi_\theta P_{k\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{[1/\gamma]}) \\ & \leq A_2(1 + d^{\varpi_2}) \sum_{k=0}^{n-1} \rho^{n-k-1} \|\pi_\theta P_{k\gamma[1/\gamma],\theta} P_{\gamma[1/\gamma],\theta} - \pi_\theta P_{k\gamma[1/\gamma],\theta} K_{\gamma,\theta}^{[1/\gamma]}\|_{V_R} \\ & \leq \gamma^{1/2} A_2 \tilde{A}_3 (1 + d^{\varpi_2})^2 / \log(1/\rho), \end{aligned}$$

which concludes the proof since A_2, \tilde{A}_3 and ρ do not depend on the dimension d . \square

Lemma 36. *There exist $A_4, \varpi_4 \geq 0$ such that for any $\theta \in \mathbb{K}$, $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} < \min(\mathfrak{m}/(2\bar{\mathfrak{N}}^2), 1/2)$,*

$$\|\delta_x \mathbb{K}_{\gamma_1, \theta_1} - \delta_x \mathbb{K}_{\gamma_2, \theta_2}\|_{V_R} \leq A_4(1 + d^{\varpi_4}) \left[\gamma_2^{-1/2} |\gamma_1 - \gamma_2| + \gamma_2^{1/2} \|\theta_1 - \theta_2\| \right] V_R^2(x),$$

with V_R given in (43) and $A_4, \varpi_4 \geq 0$ which do not depend on the dimension d .

Proof. The proof is similar to the one of [26, Proposition S18]. \square

Proof of Theorem 9. Combining Lemma 32, Theorem 33 and Proposition 35 we obtain that **H1** is satisfied. Lemma 36 implies that **H2** holds. Therefore Theorem 20 and Theorem 24 can be applied. \square

B.3 Proof of Proposition 11

We recall that for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ we define $f_\varepsilon(x) = \|F(x)\|^2 - \varepsilon$.

Proposition 37. *Assume **A1**(α) and **A2**(2α) with $\alpha > 0$. In addition, assume that F is continuous, $F^{-1}(\{0\}) \neq \emptyset$, $F^{-1}(\{0\}^c) \neq \emptyset$ and that for every open set $\mathbf{A} \neq \emptyset$, $\mu(\mathbf{A}) > 0$. Then there exists $\varepsilon_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$, the macrocanonical model π_ε associated with f_ε and the reference measure μ , solution of (P), exists and is given by $(d\pi_\varepsilon/d\mu)(x) \propto \exp[-\vartheta_\varepsilon f_\varepsilon(x)]$, with $\vartheta_\varepsilon > 0$. In addition $\lim_{\varepsilon \rightarrow 0} \vartheta_\varepsilon = +\infty$.*

Proof. Let $\varepsilon_0 = \mu(\|F\|^2)/2 > 0$, since $\mu(F^{-1}(\{0\}^c)) > 0$. Let $\varepsilon \in (0, \varepsilon_0]$. Since **A2**(2α) holds, for any $\vartheta > -\eta/C_\alpha^2$, with C_α given in **A1**(α) and η given in **A2**(2α), we have $\int_{\mathbb{R}^d} \exp[-\vartheta f_\varepsilon(x)] d\mu(x) < +\infty$. Let $\mathfrak{l} = (-\eta/C_\alpha^2, +\infty)$ and $L_\varepsilon : \mathfrak{l} \rightarrow \mathbb{R}$ such that for any $\vartheta \in \mathfrak{l}$, $L(\vartheta) = \log \left\{ \int_{\mathbb{R}^d} \exp[-\vartheta f_\varepsilon(x)] d\mu(x) \right\}$. By Proposition 4, we have that L is continuously differentiable on \mathfrak{l} . Since $F^{-1}(\{0\}) \neq \emptyset$ we have that there exists a non-empty open set \mathfrak{l}_ε such that for any $x \in \mathfrak{l}_\varepsilon$, $f_\varepsilon(x) < 0$. Therefore

$$\lim_{\vartheta \rightarrow +\infty} L(\vartheta) \geq \lim_{\vartheta \rightarrow +\infty} \log \left\{ \int_{\mathfrak{l}_\varepsilon} \exp[-\vartheta f_\varepsilon(x)] d\mu(x) \right\} = +\infty,$$

where we used the monotone convergence theorem in the last inequality. Since L is continuous we obtain that there exists $\vartheta_\varepsilon \in [0, +\infty)$ such that $L(\vartheta_\varepsilon) = \min_{[0, +\infty)} L(\vartheta)$. We have that $L'(0) \leq \varepsilon - \mu(\|F\|^2) < 0$, therefore $\vartheta_\varepsilon \in (0, +\infty)$ and $L'(\vartheta_\varepsilon) = 0$. Applying Proposition 4, we obtain that $\pi_{\vartheta_\varepsilon}$ is a solution of (P). We denote π_ε this solution.

Assume that there exists a sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ with $\varepsilon_n > 0$ such that $(\vartheta_{\varepsilon_n})_{n \in \mathbb{N}}$ is bounded. Then, up to extraction, there exists $\vartheta^* \geq 0$ such that $\lim_n \vartheta_{\varepsilon_n} = \vartheta^*$. Using the dominated convergence theorem we obtain that

$$0 = \lim_n \varepsilon_n = \lim_n \pi_{\varepsilon_n}(f_{\varepsilon_n}) = \pi_{\vartheta^*}(\|F\|^2) > 0,$$

which is a contradiction. Therefore, $\lim_{\varepsilon \rightarrow 0} \vartheta_\varepsilon = +\infty$. \square

We now turn to the study of the tightness of the sequence $(\pi_\varepsilon)_{\varepsilon>0}$ in the special case where F is given by (2). Under the assumptions of Proposition 38, for each sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ such that $\lim_n \varepsilon_n = 0$, up to extraction, we have that $(\pi_{\varepsilon_n})_{n \in \mathbb{N}}$ converges to a probability measure π_∞ which concentrates on $F^{-1}(\{0\})$.

Proposition 38. *Assume A2(2) and that for any non-empty open set $A \subset \mathbb{R}^d$, $\mu(A) > 0$. Let F be given by (2) assume that $1 \in \mathbf{j}$ and that there exists $k \in \{1, \dots, c_1\}$ such that for any $x \in \mathbb{R}^d$ with $x \neq 0$, there exists $\ell \in \{1, \dots, n_1\}$ with $e_\ell^\top \tilde{A}_1^k x > 0$. Then for any sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ with $\lim_n \varepsilon_n = 0$, $(\pi_{\varepsilon_n})_{n \in \mathbb{N}}$ is tight.*

Proof. Let F be given by (2). Then for any $x \in \mathbb{R}^d$, $f_\varepsilon(x) = \|F(x)\|^2 - \varepsilon$. We show that the level sets of $x \mapsto \|F(x)\|^2$ are compact. Let \mathbb{S}^{d-1} be the sphere in \mathbb{R}^d and define $f: \mathbb{S}^{d-1} \rightarrow (0, +\infty)$ for any $x \in \mathbb{S}^{d-1}$ by

$$f(x) = \max_{\ell \in \{1, \dots, c_1\}} \{e_\ell^\top A_1^k x\}.$$

f is continuous and since \mathbb{S}^{d-1} is compact, f reaches its minimum f_0 and therefore $f_0 > 0$. Let $x \in \mathbb{R}^d$, using that φ is non-increasing we have for any $k \in \{1, \dots, c_1\}$

$$\begin{aligned} n_1^{-1} \sum_{\ell=1}^{n_1} \mathcal{G}_1^k(x)(\ell) &= n_1^{-1} \sum_{\ell=1}^{n_1} \varphi(e_\ell^\top A_1^k x) \\ &\geq n_1^{-1} \sum_{\ell=1}^{n_1} \varphi(e_\ell^\top \tilde{A}_1^k x + e_\ell^\top A_1^k 0) \\ &\geq n_1^{-1} \varphi(f_0 \|x\|) + \min_{\ell \in \{1, \dots, n_1\}} e_\ell^\top b_1^k. \end{aligned}$$

This result combined with the fact that $\lim_{t \rightarrow +\infty} \varphi(t) = +\infty$ implies that

$$\lim_{\|x\| \rightarrow +\infty} \|F(x)\|^2 = +\infty.$$

Therefore, $F^{-1}(\{0\}^c) \neq \emptyset$. $F^{-1}(\{0\}) \neq \emptyset$ since $F(x_0) = 0$. F is continuous and A1(1) is satisfied. Therefore, Proposition 37 applies and $(\pi_{\varepsilon_n})_{n \in \mathbb{N}}$ is well-defined for any sequence $\lim_{n \rightarrow +\infty} \varepsilon_n = 0$. In addition, $\lim_{n \rightarrow +\infty} \vartheta_{\varepsilon_n} = +\infty$.

Since $x \mapsto \|F(x)\|^2$ is continuous and coercive, the level sets of f_ε are compact for any $\varepsilon > 0$. We conclude using [50, Proposition 2.3]. \square

proof of Proposition 11. The proof is then a direct consequence of the tightness of any sequence $(\pi_{\varepsilon_n})_{n \in \mathbb{N}}$ and that $\lim_{n \rightarrow +\infty} \vartheta_{\varepsilon_n} = +\infty$ combined with [50, Proposition 2.2, Theorem 2.1, Theorem 3.1]. \square

It should be noted that Proposition 11 is merely a strengthening of Proposition 38 under additional assumptions on the form of $F^{-1}(\{0\})$.

C Additional experiments

C.1 Accelerations and noiseless versions of SOUL

First we investigate the following discrete dynamics

$$\begin{cases} \tilde{X}_{k+1}^n = \tilde{X}_k^n - \gamma_n \left(\sum_{i=1}^p \tilde{\theta}_n^i \nabla F_i(\tilde{X}_k^n) + \nabla r(\tilde{X}_k^n) \right) & \text{and } \tilde{X}_0^n = \tilde{X}_{n-1}^{m_{n-1}} ; \\ \tilde{\theta}_{n+1} = \Pi_{\mathcal{K}} \left[\tilde{\theta}_n - \delta_{n+1} m_n^{-1} \sum_{j=1}^{m_n} F(\tilde{X}_j^n) \right] , \end{cases} \quad (75)$$

which corresponds to the one of SOUL (15) without the Gaussian noise term in the Langevin update. We refer to this algorithm as noiseless SOUL. Note that the families $\{\tilde{\theta}_n : n \in \mathbb{N}\}$ and $\{\tilde{X}_k^n : n \in \mathbb{N}, k \in \{0, \dots, m_n\}\}$ are deterministic up to their initialization. In the setting (75), the sequence $(\tilde{X}_0^n)_{n \in \mathbb{N}}$ seems to converge to one of the configurations presented in Figure 3, whereas the sequence $(\theta_n)_{n \in \mathbb{N}}$ does not converge towards the optimal parameters, see Figure 18. This experiment highlights that the use of a Markov kernel in the SOUL dynamics cannot be avoided in order to obtain the convergence of $(\theta_n)_{n \in \mathbb{N}}$ towards θ^* .

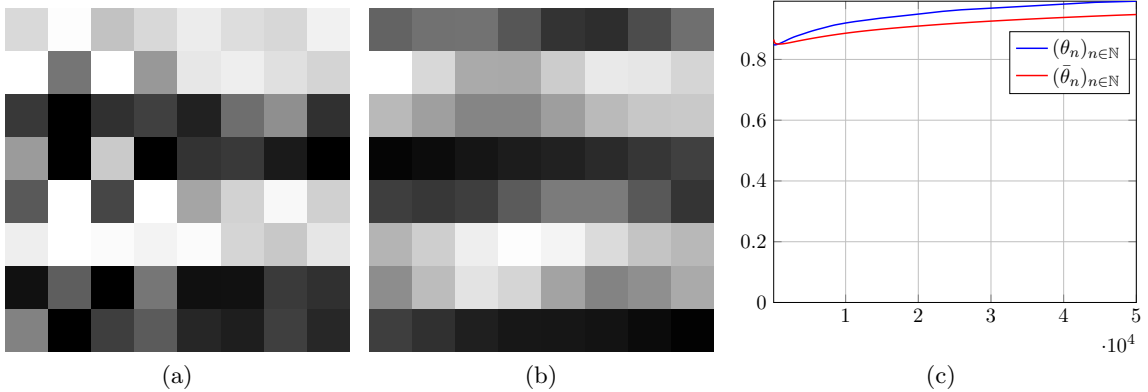


Figure 18: **Noiseless SOUL**. The original target image is recalled in (a) and the limiting configuration obtained with the noiseless SOUL algorithm (75) is given in (b), whereas the non-convergence of the error towards 0 can be observed in (c). The blue curve is the NRMSE of the sequence $(\theta_n)_{n \in \mathbb{N}}$ and the red curve is the NRMSE of the associated averaged sequence.

Another modification of the SOUL algorithm can be considered replacing the gradient descent step in (15) by another optimization methodology. Here, we focus on a popular extrapolation technique: the Nesterov acceleration. The accelerated SOUL algorithm is

then given by the following recursion

$$\begin{cases} \tilde{X}_{k+1}^n = \tilde{X}_k^n - \gamma_n \left(\sum_{i=1}^p \tilde{\theta}_n^i \nabla F_i(\tilde{X}_k^n) + \nabla r(\tilde{X}_k^n) \right) + \sqrt{2\gamma_n} Z_{k+1}^n & \text{and } \tilde{X}_0^n = \tilde{X}_{n-1}^{m_{n-1}} ; \\ \tilde{\theta}_{n+1/2} = \Pi_K \left[\tilde{\theta}_n - \delta_{n+1} m_n^{-1} \sum_{j=1}^{m_n} F(\tilde{X}_j^n) \right] ; \\ \tilde{\theta}_{n+1} = \tilde{\theta}_{n+1} + \frac{n-2}{n+1} \left\{ \tilde{\theta}_{n+1/2} - \tilde{\theta}_{n-1/2} \right\} , \end{cases} \quad (76)$$

where the sequence $(Z_k^n)_{n \in \mathbb{N}, k \in \{1, \dots, m_n\}}$ is a sequence of independent d -dimensional zero mean Gaussian random variables with covariance identity. This algorithm is not a descent algorithm but reaches the optimal convergence rate $\mathcal{O}(1/n^2)$ for convex functions in a deterministic setting, see [70, 71]. The perturbed gradient case is treated in [5] in a general framework and in [37, 6] when the perturbation is given by a Monte Carlo approximation of the gradient. Recall that for any $n \in \mathbb{N}$ we define $\eta_n = \nabla L(\theta_n) - m_n^{-1} \sum_{k=1}^{m_n} F(X_k^n)$. The assumption on the summability of the sequence of perturbations $(\eta_n)_{n \in \mathbb{N}}$ is of the form $\sum_{n \in \mathbb{N}} n \|\eta_n\| < +\infty$ in [5, Theorem 5.1]. This is a stronger requirement than $\sum_{n \in \mathbb{N}} \|\eta_n\| < +\infty$ which is a common assumption for the convergence of the perturbed gradient descent, see [59, Section 5.2.1]. In this accelerated setting (76), letting $m_n = 1$, generates oscillatory sequences $(\theta_n)_{n \in \mathbb{N}}$ which do not reduce the NRMSE. However this oscillatory effect can be counterbalanced with the use of a larger batch size, e.g. $m_n = 10$, see Figure 19.

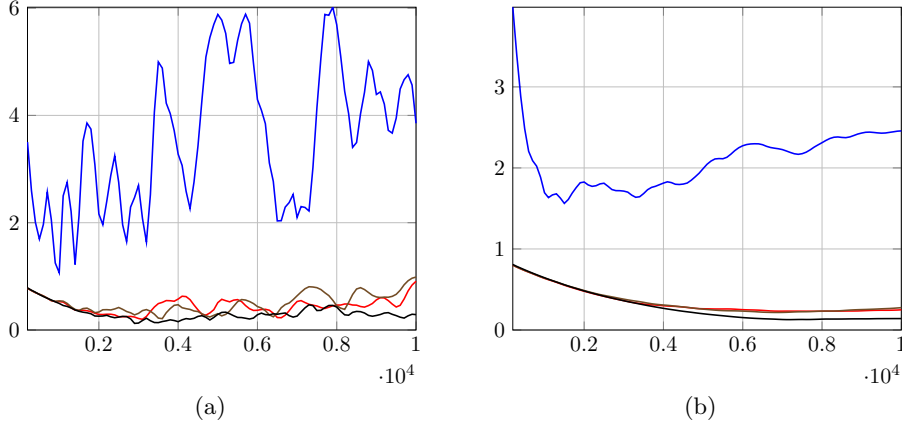


Figure 19: **Nesterov acceleration.** The Nesterov accelerated version of SOUL (76) does not yield a satisfactory sequence $(\theta_n)_{n \in \mathbb{N}}$ in terms of NRMSE (blue curve in (a)) nor a satisfactory averaged sequence $(\bar{\theta}_n)_{n \in \mathbb{N}}$ (blue curve in (b)) with parameters $\delta_n = 10^{-1}$, $\gamma_n = 10^{-4}$ and $m_n = 1$. If $\delta_n = 10^{-2}$ then the results are improved (red curves) or $\delta_n = 10^{-2} \times n^{-0.5}$ (brown curve). The best results are obtained if $\delta_n = 10^{-2} \times n^{-0.5}$ and $m_n = \lceil n^{0.5} \rceil$ (black curve).

C.2 Highly regular textures

In Appendix C.2 we perform the comparison on highly regular textures. On these textures our algorithm and the one of [44] fail at reproducing visually satisfying images (with the notable exception of the brick image for which [44] yields excellent results). Adding spectral constraints, as in [64], yields more regular images although the results are still not satisfactory. A solution is proposed in [45] where autocorrelation features are considered at each layers. This method yields the best visual results but the parameter space is much larger than the initial image space.

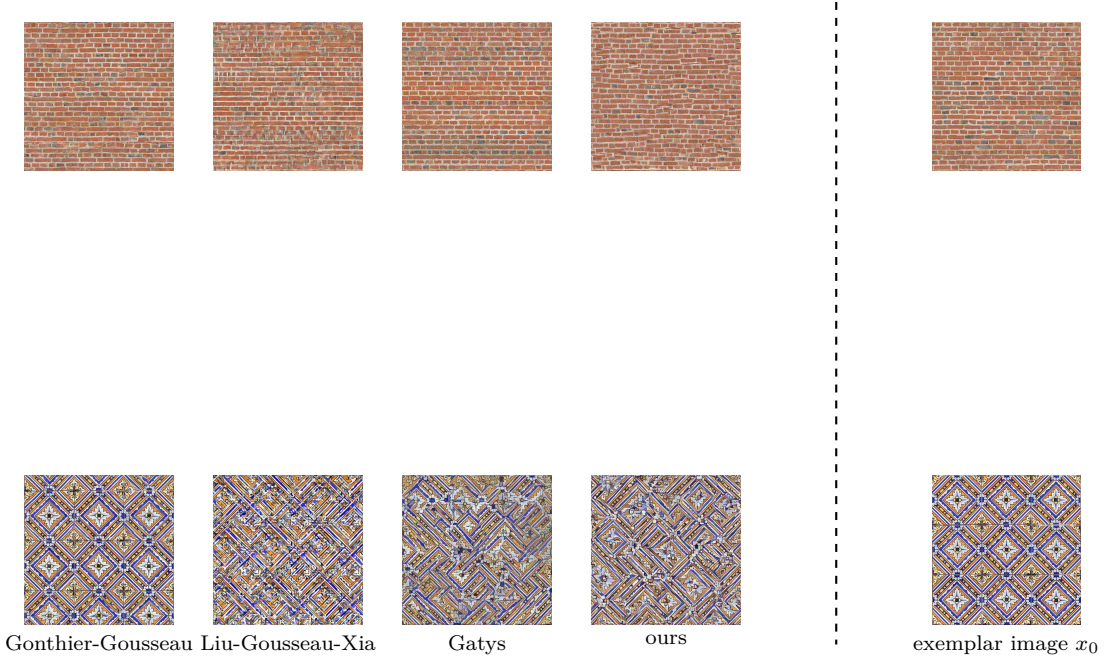


Figure 20: *Comparison with [64]*. The images presented in the column “Gonthier-Gousseau” corresponds to the features described in [45] “Liu-Gousseau-Xia” are synthesized with the features considered in [64], the ones presented in the column “Gatys” are generated with [44] and the fourth column contains our results.

D Structure of VGG19

The layers of the VGG19 network [85] are given as follows (for each convolutional layer we indicate $(c_j, n_j) \rightarrow (c_{j+1}, n_{j+1})$):

0. Convolutional layer, $(3, n_0) \rightarrow (64, n_0)$
1. ReLU layer
2. Convolutional layer, $(64, n_0) \rightarrow (64, n_0)$
3. ReLU layer
4. Max-pooling layer
5. Convolutional layer, $(64, n_0/2) \rightarrow (128, n_0/2)$
6. ReLU layer
7. Convolutional layer, $(128, n_0/2) \rightarrow (128, n_0/2)$
8. ReLU layer
9. Max-pooling layer
10. Convolutional layer, $(128, n_0/4) \rightarrow (256, n_0/4)$
11. ReLU layer
12. Convolutional layer, $(256, n_0/4) \rightarrow (256, n_0/4)$
13. ReLU layer
14. Convolutional layer, $(256, n_0/4) \rightarrow (256, n_0/4)$
15. ReLU layer
16. Convolutional layer, $(256, n_0/4) \rightarrow (256, n_0/4)$
17. ReLU layer
18. Max-pooling layer
19. Convolutional layer, $(256, n_0/8) \rightarrow (512, n_0/8)$
20. ReLU layer
21. Convolutional layer, $(512, n_0/8) \rightarrow (512, n_0/8)$
22. ReLU layer
23. Convolutional layer, $(512, n_0/8) \rightarrow (512, n_0/8)$
24. ReLU layer
25. Convolutional layer, $(512, n_0/8) \rightarrow (512, n_0/8)$
26. ReLU layer
27. Max-pooling layer
28. Convolutional layer, $(512, n_0/16) \rightarrow (512, n_0/16)$
29. ReLU layer
30. Convolutional layer, $(512, n_0/16) \rightarrow (512, n_0/16)$
31. ReLU layer
32. Convolutional layer, $(512, n_0/16) \rightarrow (512, n_0/16)$
33. ReLU layer
34. Convolutional layer, $(512, n_0/16) \rightarrow (512, n_0/16)$
35. ReLU layer
36. Max-pooling layer