

Collective proposal distributions for nonlinear MCMC samplers: Mean-field theory and fast implementation*

Grégoire Clarté¹, Antoine Diez^{2,3} and Jean Feydy^{4,5}

¹*Department of Computer Science, University of Helsinki, FCAI
e-mail: gregoire.clarte@helsinki.fi*

²*Department of Mathematics, Imperial College London, UK*

³*Institute for the Advanced Study of Human Biology (ASHBi), Kyoto University, Japan
e-mail: diez.antoinenicolas.4e@kyoto-u.ac.jp*

⁴*HeKA team, Inria Paris, F-75012 Paris, France, e-mail: jean.feydy@inria.fr*

⁵*Inserm, Centre de Recherche des Cordeliers, Sorbonne Université, Université Paris Cité,
F-75006 Paris, France*

Abstract: Over the last decades, various “non-linear” MCMC methods have arisen. While appealing for their convergence speed and efficiency, their practical implementation and theoretical study remain challenging. In this paper, we introduce a non-linear generalization of the Metropolis-Hastings algorithm to a proposal that depends not only on the current state, but also on its *law*. We propose to simulate this dynamics as the mean field limit of a system of interacting particles, that can in turn itself be understood as a generalisation of the Metropolis-Hastings algorithm to a population of particles. Under the double limit in number of iterations and number of particles we prove that this algorithm converges. Then, we propose an efficient GPU implementation and illustrate its performance on various examples. The method is particularly stable on multimodal examples and converges faster than the classical methods.

MSC2020 subject classifications: Primary 65C05, 65C10, 65C35; secondary 60J05, 62-08, 62-04.

Keywords and phrases: Sampling algorithm, particle method, propagation of chaos, entropy methods, GPU.

Received October 2021.

1. Introduction

1.1. Background

Monte Carlo methods are designed to estimate the integral of a function of interest φ (sometimes called *observable*) under a probability measure π . The

*This research was conducted while A.D. was supported by an EPSRC-Roth scholarship co-funded by the Engineering and Physical Sciences Research Council and the Department of Mathematics at Imperial College London.

integral $\int \varphi(x)\pi(dx)$ is approximated by a random estimator of the form

$$\frac{1}{n} \sum_{i=1}^n \varphi(X_i),$$

where the X_i are n independent and identically distributed (i.i.d.) random variables with law π . The law of large numbers ensures the convergence as $n \rightarrow +\infty$ of this estimator. For complex cases, when i.i.d. π -distributed random variables cannot be generated exactly, a now classical procedure consists in constructing a Markov chain $(X_i)_i$ with stationary distribution π . Ergodic theory results then ensure that the estimator above still converges, even though the X_i are not independent from each other. The Metropolis-Hastings algorithm [66, 65, 45] provides a simple construction for such a Markov chain that only requires to evaluate the probability density function π up to a multiplicative constant. The constructed chain is a random walk biased by an accept-reject step. Its convergence properties have been thoroughly studied, for example in [64].

This well-known procedure has become a building block for more advanced samplers, that are designed to overcome the known flaws of the Metropolis-Hastings algorithm: slow convergence, bad mixing properties for multimodal distributions *etc.* Such extensions include for instance the Wang and Landau algorithm [14], regional MCMC algorithms [22], or non Markovian (adaptive) versions [44, 43, 6, 5] where the next proposed state depends on the whole history of the process. The more recent PDMP samplers [34, 82] provide an alternative to the discrete time accept-reject scheme, replacing it by a continuous time non reversible Markov process with random jumps. Finally, more complex algorithms are based on the evaluation of the gradient of π , see for instance the Metropolis-adjusted Langevin [11] and the Hamiltonian Monte Carlo algorithms [32]. All these methods can be seen as “linear” as the next position only depends on the current position.

A non-Markovian alternative to Metropolis-Hastings like methods is given by Importance Sampling algorithms. By drawing i.i.d. samples from an auxiliary distribution q , which is usually simple and called the importance distribution, we can build an estimator using the following identity:

$$\int \varphi(x)\pi(x)dx = \int \frac{\pi(x)}{q(x)}\varphi(x)q(x)dx \simeq \frac{1}{n} \sum_{i=1}^n w_i\varphi(X_i),$$

where the X_i are i.i.d. with law q and the $w_i \propto \pi(X_i)/q(X_i)$, are called the *importance weights*. The choice of q is critical, as bad choices can lead to a degeneracy of the importance weights. Iterative methods have been developed to sequentially update the choice of the importance distribution, and to update the X_i now interpreted as particles that evolve along iterations. Among these algorithms, we can cite the Sequential Importance Sampling algorithm [39], the Population Monte Carlo (PMC) methods [30, 18, 17] or the recent Safe Adaptive Importance Sampling (SAIS) algorithm [25]. This paradigm is in particular well-suited to the study of filtering problems [39], leading to the development of

Sequential Monte Carlo (SMC) methods [24, 31]. A review of population-based algorithms and of the SMC method can be found in [49].

The SMC methodology has recently been used to design and study *nonlinear* MCMC algorithms [4]. This framework can be seen as a generalisation of some non-Markovian extensions of the Metropolis-Hastings algorithm (such as the “resampling from the past” procedure [44, 6]) but also allows the use of a wider range of algorithmic techniques. Examples are given in [4, 3] and are often based on the simulation of auxiliary chains. In the present article, we show that an alternative procedure based on the simulation of a swarm of *interacting* particles can also be used to approximate a nonlinear Markov chain. This provides a multi-particle generalisation of the Metropolis-Hastings procedure.

The duality between particle systems and non-linear Markov processes has first been underlined in statistical physics; on the mathematical side, it has been the subject of the pioneering works of [61, 62, 55, 29]. A key result is the *propagation of chaos* property formalised by [55], which implies that under an initial *chaoticity* assumption on the law of the particles, the empirical measure of the system at any further time converges towards a deterministic limit; this type of limit is called *mean-field limit*. In a continuous time framework, this limit classically satisfies a nonlinear Partial Differential Equation (PDE) [78, 63]. The original diffusion framework has been extended to jump and jump-diffusion processes in [40, 41]. We refer the interested reader to [9, 10] for recent reviews and surveys of the applications of such models. We also mention that this methodology has been used in the analysis of particle methods in filtering problems [23, 2].

1.2. Objective and methods

Let π be a target measure on $E \subset \mathbb{R}^d$, known up to a multiplicative constant and which is assumed to have a density with respect to the Lebesgue measure. We denote $\mathcal{P}(E)$ the set of probability measures on E . The goal of the present article is to build a *nonlinear* Markov chain $(\bar{X}_t)_t$ on E that samples π efficiently. Given a sample \bar{X}_t at iteration t , we draw \bar{X}_{t+1} according to

$$\bar{X}_{t+1} \sim K_{\mu_t}(\bar{X}_t, dy),$$

where the *transition kernel* is defined by:

$$K_{\mu_t}(x, dy) := \underbrace{h(\alpha_{\mu_t}(x, y))\Theta_{\mu_t}(dy|x)}_{\text{accept}} + \underbrace{\left[1 - \int_{z \in E} h(\alpha_{\mu_t}(x, z))\Theta_{\mu_t}(dz|x)\right]}_{\text{reject}} \delta_x(dy) \quad (1)$$

and where for $t \in \mathbb{N}$, $\mu_t \in \mathcal{P}(E)$ is the law of \bar{X}_t . In the discrete setting, this method is implemented by Algorithm 1, detailed below. It relies on the following quantities:

- The **proposal distribution**, a map

$$\Theta : E \times \mathcal{P}(E) \longrightarrow \mathcal{P}_0^{\text{ac}}(E),$$

where $\mathcal{P}_0^{\text{ac}}(E)$ the subset of non-vanishing absolutely continuous probability measures. For $x \in E$ and $\mu \in \mathcal{P}(E)$, its associated proposal probability density function is denoted by:

$$\Theta(x, \mu)(y)dy \equiv \Theta_\mu(y|x)dy.$$

Intuitively, the probability distribution $\Theta(x, \mu)$ can be understood as an approximation of the target π that our method uses to *propose* new samples y in a neighborhood of a point x , relying on the information that is provided by a probability measure μ . In the following μ will be the empirical distribution associated with a system of particles. To ensure a fast convergence of our method, both in terms of computation time and number of iterations, the distribution $\Theta(x, \mu)$ should be both **easy to sample** and close to the target distribution π . In practice, the choice of a good proposal Θ depends on the assumptions that can be made on the distribution π . We will present several examples in Section 3. A simple example to keep in mind and which will be detailed in Section 3.2 is $\Theta_\mu(y|x) = K \star \mu(y)$ where K is a random-walk kernel.

- For $\mu \in \mathcal{P}(E)$ and $x, y \in E$, the **acceptance ratio** is defined by:

$$\alpha_\mu(x, y) := \frac{\Theta_\mu(x|y)\pi(y)}{\Theta_\mu(y|x)\pi(x)}.$$

This quantity expresses the relative appeals of the transition $x \rightarrow y$ for the “model” density $\Theta_\mu(dy|x)$ and the ground truth target $\pi(dy)$, as in classical Metropolis-Hastings methods. Noticeably the denominator is never null because of the non-vanishing assumption on Θ . Crucially, it can be computed even when the law of π is only known up to a multiplicative constant and allows our method to account for mis-matches between the proposal Θ and the distribution to sample π . Note that in practice, for the sake of numerical stability, the acceptance ratio is often manipulated through its logarithm:

$$\underbrace{\log \alpha_\mu(x, y)}_{\text{“correction”}} := \underbrace{[\log \pi(y) - \log \pi(x)]}_{\text{appeal of } x \rightarrow y \text{ for } \pi} - \underbrace{[\log \Theta_\mu(y|x) - \log \Theta_\mu(x|y)]}_{\text{appeal of } x \rightarrow y \text{ for } \Theta_\mu}.$$

In the following, we will assume that π is bounded away from zero so that the acceptance ratio is always well-defined.

- The **acceptance function** is a non-decreasing Lipschitz map of the form $h : [0, +\infty) \rightarrow [0, 1]$ which satisfies

$$\forall u \in [0, \infty), \quad uh(1/u) = h(u). \quad (2)$$

A classical example is $h(u) = \min(1, u)$. Other examples in the literature can be found in [1] and the references therein. In the following, we will occasionally use (for technical reasons) the “lazy Metropolis-Hastings” acceptance function $h(u) = \eta \min(1, u)$, $\eta < 1$ introduced in [57].

As detailed in Algorithm 1, we combine the acceptance ratio $\alpha_\mu(x, y)$ and the acceptance function h to **reject** proposed samples y that are much more appealing for $\Theta_\mu(\cdot|x)$ than they are for π .

This **necessary correction** ensures that our method samples the target π instead of the simpler proposal distribution. On the other hand, it can also slow down the method if the proposed samples y keep being rejected. Efficient proposal distributions should keep the acceptance ratio high enough to ensure a renewal of the population of samples X_t^i at every iteration and thus provide good mixing properties.

Non-linearity We say that the transition kernel is nonlinear due to its dependency on the law of the chain that it generates. When the proposal distribution does not depend on μ_t , the kernel is *linear* and we obtain the general form of the classical Metropolis-Hastings kernel.

Interest of the method The main interest of the method appears for complex distributions, that is multimodal distributions. From an MCMC point of view, the use of a non-linear sampler removes the mixing problem which is one of the main drawbacks of the Metropolis-Hastings algorithm. For nonlinear samplers, exponential convergence is ensured as soon as the distribution carries weight in every mode, without having to explore and exit each of the modes. On the contrary, the Metropolis-Hastings algorithms may remain stuck in a single mode even with long runs. Nonlinear methods allow to learn efficiently the relative weight of the distribution modes, which is unavailable even for several independent runs of Metropolis-Hastings algorithm with different initialisations, that is the most direct parallel version of Metropolis-Hastings. The numerical experiments will confirm that nonlinear methods perform better for both exploration and convergence.

Contributions We follow [4] and split our analysis into two steps:

1. We show that our non-linear kernel admits π as a stationary distribution and study its asymptotic properties.
2. We present a practical implementation based on the simulation of a system of interacting particles that enables the simulation of this kernel for different choices of the proposal distribution.

Analytical study Starting from an initial distribution $\mu_0 \in \mathcal{P}(E)$, the law μ_t of the nonlinear chain at the t -th iteration satisfies

$$\mu_{t+1} = \mathcal{T}[\mu_t]$$

where $\mathcal{T} : \mathcal{P}(E) \rightarrow \mathcal{P}(E)$ is the **transition operator** defined by duality in the space of measures by:

$$\langle \varphi, \mathcal{T}[\mu] \rangle := \int_E \varphi(x) \mathcal{T}[\mu](dx) = \iint_{E \times E} \varphi(y) K_\mu(x, dy) \mu(dx), \quad (3)$$

```

Input: An initial population of particles  $(X_0^1, \dots, X_0^N) \in E^N$ ,
a maximum time  $T \in \mathbb{N}$ , a proposal distribution  $\Theta$ 
and an acceptance function  $h$ 
Output: A sample  $(X_t^i)_{1 \leq i \leq N; 1 \leq t \leq T}$ 
for  $t = 0$  to  $T - 1$  do
  for  $i = 1$  to  $N$  do
    (Proposal) Draw  $Y_t^i \sim \Theta_{\hat{\mu}_t^N}(\cdot | X_t^i)$  a proposal for the new state of particle  $i$ ;
    (Acceptation) Compute  $\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i) = \frac{\Theta_{\hat{\mu}_t^N}(X_t^i | Y_t^i) \pi(Y_t^i)}{\Theta_{\hat{\mu}_t^N}(Y_t^i | X_t^i) \pi(X_t^i)}$ ;
    Draw  $U_t^i \sim \mathcal{U}([0, 1])$ ;
    if  $U_t^i \leq h(\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i))$  then
      | Set  $X_{t+1}^i = Y_t^i$ ; // Accept, probability  $h(\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i))$ .
    else
      | Set  $X_{t+1}^i = X_t^i$ ; // Reject, likely if  $\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i) \simeq 0$ .
    end
  end
end

```

Algorithm 1: Collective Monte Carlo (CMC)

for any continuous bounded test function $\varphi \in C_b(E)$. Thanks to the **detailed balance condition** (also called *micro-reversibility* in the context of statistical mechanics [80]), for all $x, y \in E$ and $\mu \in \mathcal{P}(E)$:

$$\pi(x)\Theta_\mu(y|x)h(\alpha_\mu(x, y)) = \pi(y)\Theta_\mu(x|y)h(\alpha_\mu(y, x)), \quad (4)$$

the transition operator can be rewritten:

$$\mathcal{T}[\mu](dx) = \mu(dx) + \int_E \pi(x)W_\mu(x \rightarrow y) \left(\frac{\mu(dy)}{\pi(y)} dx - \frac{\mu(dx)}{\pi(x)} dy \right),$$

with $W_\mu(x \rightarrow y) := \Theta_\mu(y|x)h(\alpha_\mu(x, y))$, from which it can be easily seen that $\mathcal{T}[\pi] = \pi$.

We are going to develop an analytical framework in which the convergence of the sequence of iterations of the transition operator can be analysed. Using entropy methods, we prove the exponential convergence towards π for a large class of proposal distributions. We show that in an asymptotic regime to be detailed, the rate of convergence depends only on how close from the target is the initial condition. As a byproduct, in the linear Metropolis-Hastings case, we obtain a convergence result similar to the one obtained in [27].

Efficient implementation It is not possible in general to sample directly \bar{X}_t from a nonlinear kernel because the law μ_t is not available. We therefore rely on a **mean-field** particle method to approximate such samples. Starting from a swarm of N particles $X_t^1, \dots, X_t^N \in E$ at the iteration t , we construct the next iteration by sampling independently for $i \in \{1, \dots, N\}$:

$$X_{t+1}^i \sim K_{\hat{\mu}_t^N}(X_t^i, dy),$$

where

$$\hat{\mu}_t^N := \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i} \in \mathcal{P}(E),$$

is the (random) empirical measure of the system of particles which is used as a proxy of the distribution μ_t . We show that as N goes to infinity and for each $t \in \mathbb{N}$, $\hat{\mu}_t^N$ converges towards a deterministic limit which is the t -th iterate of the nonlinear operator \mathcal{T} starting from μ_0 . Moreover we show that the N particles are asymptotically, in N , independent thus forming an approximation of a system of N independent nonlinear Markov chains with transition kernel (1).

A drawback of this approach is its high computational cost, that may scale in $\mathcal{O}(N^2)$ or $\mathcal{O}(N^3)$ for some choices of the proposal Θ . To overcome this difficulty, we propose an implementation based on GPU, more precisely on the techniques developed in the KeOps library [21] by the third author.

Outline Section 2 is devoted to the convergence analysis of Algorithm 1 for a general class of proposal distributions. The mean-field limit and the long-time asymptotic properties are studied respectively in Section 2.3 and Section 2.4. Several variants of the main algorithm are presented in Section 3. The GPU implementation of the different algorithms is detailed in Section 4. Applications to various problems are presented in Section 5. In the appendix, we present the complementary proofs (Appendix A) as well as variations on the results of the paper (Appendix B). We also add some complementary remarks on the links between our method and other classical methods (Appendix C) and other and additional numerical results (Appendix D and Appendix E).

Throughout this article, we assume that π satisfies the following assumption.

Assumption 1.1. *The support of π , denoted by E , is a compact subset of \mathbb{R}^d . The target distribution π is Lipschitz continuous and π does not vanish on E :*

$$m_0 := \inf_E \pi > 0 \quad \text{and} \quad M_0 := \sup_E \pi < +\infty.$$

Notations. *The following notations will be used throughout the article.*

- $\mathcal{P}(E)$ denotes the set of probability measures on E .
- $\mathcal{P}^{\text{ac}}(E)$ denotes the set of probability measures on E which are absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^d . A probability measure in $\mathcal{P}^{\text{ac}}(E)$ is identified with its associated probability density function: when $f \in \mathcal{P}^{\text{ac}}(E)$ we write indifferently,

$$f(\text{d}x) \equiv f(x)\text{d}x.$$

- $\mathcal{P}_0^{\text{ac}}(E) \subset \mathcal{P}^{\text{ac}}(E)$ denotes the subset of continuous probability density functions which do not vanish on E (recall that E is compact).
- A test function $\varphi \in C_b(E)$ is a continuous (bounded) function on E . For $\mu \in \mathcal{P}(E)$ we write indifferently

$$\langle \varphi, \mu \rangle \equiv \int_E \varphi(x)\mu(\text{d}x).$$

- $X \sim \mu$ means that the law of the random variable $X \in E$ is $\mu \in \mathcal{P}(E)$.
- $W^1(\mu, \nu)$ denotes the Wasserstein-1 distance between the two probability measures $\mu, \nu \in \mathcal{P}(E)$, defined by $W^1(\mu, \nu) := \inf_{X \sim \mu, Y \sim \nu} \mathbb{E}|X - Y|$ (see [85, Chapter 6] for additional details on Wasserstein distances).

2. Convergence analysis

Algorithm 1 gives a trajectorial approximation of the nonlinear Markov chain $(\overline{X}_t)_t$ with law $(\mu_t)_t$ defined by the transition kernel (1). In this section, we prove the convergence of this algorithm under general assumptions on the proposal distribution Θ described below. The proof of our main result is split into two steps each summarised in a theorem, first the mean-field limit when $N \rightarrow +\infty$ (Section 2.3) and then the long-time convergence towards π (Section 2.4).

2.1. Assumptions

For our theoretical results, we will need the following assumptions. The first three following assumptions are needed to prove the many-particle limit in Section 2.3.

Assumption 2.1 (Boundedness). *There exist two constants $\kappa^-, \kappa^+ > 0$ such that for all $\mu \in \mathcal{P}(E)$ and for all $x, y \in E$:*

$$\kappa_- \leq \Theta_\mu(y|x) \leq \kappa_+.$$

Assumption 2.2 (L^∞ Lipschitz). *The map $\Theta : E \times \mathcal{P}(E) \rightarrow \mathcal{P}_0^{\text{ac}}(E)$ is globally Lipschitz for the L^∞ -norm on E : there exists a constant $L > 0$ such that for all $x, y, x', y' \in E$ and for all $(\mu, \nu) \in \mathcal{P}(E)^2$:*

$$|\Theta_\mu(y|x) - \Theta_\nu(y'|x')| \leq L \left(W^1(\mu, \nu) + |x - x'| + |y - y'| \right).$$

Assumption 2.3 (W^1 non-expansive). *The map $\Theta : E \times \mathcal{P}(E) \rightarrow \mathcal{P}_0^{\text{ac}}(E)$ is non-expansive for the Wasserstein-1 distance: for all $x, x' \in E$ and for all $(\mu, \nu) \in \mathcal{P}(E)^2$,*

$$W^1 \left(\Theta_\mu(dy|x), \Theta_\nu(dy|x') \right) \leq W^1(\mu, \nu) + |x - x'|.$$

All the proposal distributions presented in Section 3 are based on a convolution product with one or many fixed kernels. The smoothness and boundedness properties of the proposal distribution (Assumptions 2.1, 2.2, 2.3) are thus inherited from the properties of these kernels. Moreover, we note that these assumptions on the proposal distribution and the compactness Assumption 1.1 on the target distribution could be replaced by assumptions on the ratio between the two distributions as explained at the end of Appendix A (see in particular Remark A.4).

The next two assumptions are needed to prove the long-time convergence property in Section 2.4.

Assumption 2.4. *There exists $\eta \in (0, 1)$ such that*

$$\forall u \in [0, +\infty), \quad h(u) \leq \eta.$$

Remark 2.5. Assumption 2.4 is satisfied for instance for $h(u) = \eta \min(1, u)$ (which is referred as the “lazy Metropolis-Hastings” acceptance function in [57]). We make this assumption mostly for technical reasons in order to obtain in an easy manner an explicit convergence rate in Theorem 2.15. However, using compactness arguments, we can prove that the convergence of $(\mu_t)_t$ towards π still holds without this assumption (see Corollary B.6).

The next assumption ensures that the proposal distribution is not too “far” from π .

Assumption 2.6 (Monotonicity). *The proposal distribution Θ satisfies the following monotonicity property: there exists a non decreasing function*

$$c^- : [0, 1] \rightarrow (0, 1],$$

such that for all $\mu \in \mathcal{P}_0^{\text{ac}}(E)$,

$$\inf_{(x,y) \in E^2} \frac{\Theta_\mu(y|x)}{\pi(y)} \geq c^- \left(\inf_{x \in E} \frac{\mu(x)}{\pi(x)} \right).$$

Remark 2.7. Note that under Assumptions 1.1 and 2.1, Assumption 2.6 is always satisfied with a constant function $c^- \equiv \kappa_-/M_0$. Sharper results can be obtained in specific cases. Moreover, note that Assumptions 2.1, 2.2 and 2.3 are not necessary to prove Theorem 2.15.

Remark 2.8. The monotonicity Assumption 2.6 is satisfied for all the “convolution based” methods such as Algorithms 2 and 3, which will be introduced later, since for $m > 0$, it holds that:

$$[\forall y \in E, \quad m\pi(y) \leq \mu(y)] \quad \implies \quad [\forall y \in E, \quad mK \star \pi(y) \leq K \star \mu(y)],$$

and therefore, if the left-hand side condition holds, dividing by $\pi(y)$ yields:

$$\forall y \in E, \quad \frac{K \star \mu(y)}{\pi(y)} \geq m \frac{K \star \pi(y)}{\pi(y)}.$$

On the right-hand side of the last inequality the ratio $K \star \pi(y)/\pi(y)$ depends only on π and is bounded from below, at least for small interaction kernels K , since $K \star \pi$ converges uniformly towards π as $K \rightarrow \delta_0$. In the degenerate case $K = \delta_0$, we obtain $c^-(m) = m$ for all $m > 0$ (see Remark 2.17).

2.2. Main result

The following theorem is our main convergence result.

Theorem 2.9. Let $\hat{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}$ be the random empirical distribution of the particle system constructed at the t -th iteration of Algorithm 1 with an i.i.d. $\mu_0 \in \mathcal{P}_0^{\text{ac}}(E)$ distributed initial condition. Under Assumptions 2.1, 2.2, 2.3, 2.4 and 2.6, there exist $C_1, C_2, C_3 > 0$ and $\lambda \in (0, 1)$ which depend only on μ_0, π and E such that for all $t \in \mathbb{N}$,

$$\mathbb{E}W^1(\hat{\mu}_t^N, \pi) \leq C_1 \beta(N) e^{tC_2} + C_3 (1 - \lambda)^{t/2},$$

where

$$\beta(N) := \begin{cases} CN^{-1/2} & \text{if } d = 1 \\ CN^{-1/2} \log(N) & \text{if } d = 2 \\ CN^{-1/d} & \text{if } d > 2 \end{cases}, \quad (5)$$

and $C > 0$ is a constant which depends only on E and π . In particular $\beta(N) \rightarrow 0$ as $N \rightarrow +\infty$.

Remark 2.10. The convergence speed in t of the second term on the right-hand side is geometric. It corresponds to the convergence speed of the limit nonlinear Markov chain stated in Theorem B.2. Note that it is also the convergence speed of the classical Metropolis-Hastings algorithm. This result does not account for practical mixing issues that will be described later.

Proof. This result is deduced from Theorem 2.11 and 2.15, as it is a direct consequence of the triangle inequality

$$\mathbb{E}W^1(\hat{\mu}_t^N, \pi) \leq \mathbb{E}W^1(\hat{\mu}_t^N, \mu_t) + \mathbb{E}W^1(\mu_t, \pi),$$

and the convergence results (7) and (10). In order to bound the second term on the right-hand side, we recall that on the compact set E , the total variation norm controls the Wasserstein-1 distance [85, Theorem 6.15]. \square

By the Kantorovich characterisation of the Wasserstein distance [85, Remark 6.5], this result ensures the convergence in expectation of any Lipschitz observable (and by density of any continuous observable) in the double limit $N \rightarrow +\infty$ and $t \rightarrow +\infty$ provided that $\beta(N)e^{C_2 t} \rightarrow 0$. More precisely, for any $\varepsilon > 0$, it is possible to find a time $t(\varepsilon)$ such that the second term on the right-hand side of the bound in Theorem 2.9 is smaller than $\varepsilon/2$. Then since $\beta(N) \rightarrow 0$, we can find $N(\varepsilon)$ which makes the first term at $t(\varepsilon)$ also smaller than $\varepsilon/2$. This shows that the expectation of the distance between the empirical measure of the particle system and the target distribution can be made arbitrarily small. Note however that these results and in particular the link between N and t are mostly of theoretical nature and often suboptimal; in practice higher convergence rates may be obtained (see Section 5).

Moreover, although Theorem 2.9 states a geometric convergence result for the nonlinear samplers which is similar to classical convergence results for the classical (linear) Metropolis-Hastings algorithm, the behaviours of nonlinear and linear samplers can be very different in practise. Nonlinear samplers are typically much more efficient in the sense that they need less iterations to converge and although each iteration is more costly, the total computation time remains much

shorter. In addition to the experiments shown in Section 5, Remark 2.17 gives a theoretical result obtained as a consequence of Theorem 2.9 which illustrates this difference.

2.3. Mean field approximation

In this section, we show that the system of particles defined by Algorithm 1 satisfies the propagation of chaos property and that the limiting law at each iteration is the law of the nonlinear Markov chain with transition kernel satisfying (1). From now on, we assume that the proposal distribution Θ satisfies Assumptions 2.1, 2.2 and 2.3 (see also the discussion in Remark A.4).

The main result of this section is the following theorem.

Theorem 2.11 (Coupling bound). *Let Θ be a proposal distribution which satisfies Assumptions 2.1, 2.2 and 2.3. Let $t \in \mathbb{N}$. There is a probability space on which are defined a system of N i.i.d. nonlinear Markov chains $(\bar{X}_t^i)_t$, $i \in \{1, \dots, N\}$, defined by the transition kernel (1) and a system of N particles $(X_t^i)_t$, $i \in \{1, \dots, N\}$, which is equal in law to the N -particle system constructed by Algorithm 1, such that*

$$\forall i \in \{1, \dots, N\}, \quad \mathbb{E}|\bar{X}_t^i - X_t^i| \leq \beta(N)e^{tC_\Theta}, \quad (6)$$

where $C_\Theta > 0$ is a constant which depends only on π and Θ and where $\beta(N)$ is defined by (5)

The proof of Theorem 2.11 is based on coupling arguments inspired by [78] and adapted from [28]. It can be found in Appendix A. This so-called coupling estimate classically implies the following properties (see [78]).

Corollary 2.12 (Mean-field limit and propagation of chaos). *Let Θ be a proposal distribution which satisfies Assumptions 2.1, 2.2 and 2.3. Let $(X_0^i)_{i \in \{1, \dots, N\}}$ be N i.i.d. random variables with common law $\mu_0 \in \mathcal{P}(E)$ (chaoticity assumption). Let $t \in \mathbb{N}$ and let $(X_t^i)_{i \in \{1, \dots, N\}}$ be the N particles constructed at the t -th iteration of Algorithm 1. Let $\mu_t = \mathcal{T}^{(t)}[\mu_0]$ be the t -th iterate of the transition operator (3) starting from μ_0 , that is μ_t is the law of the nonlinear Markov chain defined by the transition kernel (1) at iteration t . Then the following properties hold true.*

1. The (random) empirical measure $\hat{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}$ satisfies

$$\mathbb{E}W^1(\hat{\mu}_t^N, \mu_t) \leq C_1 \beta(N) e^{tC_2}, \quad (7)$$

where $C_1, C_2 > 0$ are two absolute constants and $\beta(N)$ is defined by (5).

2. The sequence of random empirical measures $\hat{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}$, seen a sequence of $\mathcal{P}(E)$ -valued random variables, converges in law towards the deterministic limit:

$$\hat{\mu}_t^N \xrightarrow{N \rightarrow +\infty} \mu_t.$$

We recall that the space of probability measures is endowed with the topology of the weak convergence, i.e. convergence against bounded continuous test functions. The convergence in law of a sequence of random measures thus means the weak convergence of their laws in the space $\mathcal{P}(\mathcal{P}(E))$, i.e. since μ_t is deterministic, it holds that $\text{Law}(\hat{\mu}_t^N) \in \mathcal{P}(\mathcal{P}(E)) \rightarrow \delta_{\mu_t} \in \mathcal{P}(\mathcal{P}(E))$ where δ_{μ_t} is the Dirac mass at the point $\mu_t \in \mathcal{P}(E)$ and the convergence is the weak convergence in $\mathcal{P}(\mathcal{P}(E))$. This is also equivalent to say that for all test function $\Phi \in C_b(\mathcal{P}(E))$, $\mathbb{E}[\Phi(\hat{\mu}_t^N)] \rightarrow \Phi(\mu_t)$.

3. For every ℓ -tuple of continuous bounded functions $\varphi_1, \dots, \varphi_\ell$ on E , it holds that:

$$\int_{E^\ell} \varphi_1(x_1) \dots \varphi_\ell(x_\ell) \mu_t^{\ell, N}(dx_1, \dots, dx_\ell) \xrightarrow{N \rightarrow +\infty} \prod_{k=1}^{\ell} \langle \varphi_k, \mu_t \rangle, \tag{8}$$

where $\mu_t^{\ell, N}$ is the joint law at time t of any subset of ℓ particles constructed by Algorithm 1 at iteration t .

Proof. The first property follows from the triangle inequality

$$\mathbb{E}W^1(\hat{\mu}_t^N, \mu_t) \leq \mathbb{E}W^1(\hat{\mu}_t^N, \bar{\mu}_t^N) + \mathbb{E}W^1(\bar{\mu}_t^N, \mu_t),$$

where $\bar{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}_t^i}$ is the empirical measure of the N nonlinear Markov chains constructed in Theorem 2.11. The first term on the right-hand side is bounded by (6) by definition of the Wasserstein distance and the exchangeability of the processes. The second term on the right-hand side is bounded by $\beta(N)$ by [38, Theorem 1]. The second property is a classical consequence of (6), see [46, Section 1]. The third property is equivalent to the second property by [78, Proposition 2.2]. \square

The property (8) corresponds to the original formulation of the propagation of chaos introduced by [55]. From our perspective, it justifies the use of Algorithm 1 and ensures that as the number of particles grows to infinity and despite the interactions between the particles, we asymptotically recover an i.i.d. sample. The final MCMC approximation of the expectation of an observable $\varphi \in C_b(E)$ is thus given at the t -th iteration by:

$$\int_E \varphi(x) \pi(dx) \simeq \frac{1}{N} \sum_{i=1}^N \varphi(X_t^i).$$

Remark 2.13. In [52], the authors prove the propagation of chaos towards a continuous-time (nonlinear) diffusion process for a classical random walk Metropolis-Hastings algorithm in the product space E^N where the trajectory of each of the N dimensions is interpreted as a particle, where the target distribution $\pi^{\otimes N}$ is tensorized and under a specific scaling limit in N for the time and the size of the random-walk kernel. In this algorithm, each move is globally accepted or rejected for the N particles whereas in Algorithm 1, the acceptance step is individualized for each particle. One consequence is that, unlike the algorithm in [52], for a fixed number N of particles, $\pi^{\otimes N}$ is in general *not* a stationary distribution of the particle system defined by Algorithm 1.

2.4. Long-time asymptotics

In this section, we prove that π is the unique stationary measure of the nonlinear Markov chain defined by the transition kernel (1) and we give a quantitative long-time convergence result.

2.4.1. Main result

Let $(\mu_t)_{t \in \mathbb{N}}$ be the sequence of laws of the nonlinear Markov chain defined by the transition kernel (1). It satisfies the recurrence relation

$$\mu_{t+1} = \mathcal{T}[\mu_t], \quad (9)$$

where we recall that given $\mu \in \mathcal{P}(E)$, the transition operator \mathcal{T} is defined by:

$$\mathcal{T}[\mu](dx) = \mu(dx) + \int_E \pi(x) W_\mu(x \rightarrow y) \left(\frac{\mu(dy)}{\pi(y)} dx - \frac{\mu(dx)}{\pi(x)} dy \right),$$

and

$$W_\mu(x \rightarrow y) := \Theta_\mu(y|x) h(\alpha_\mu(x, y)).$$

Note that if the initial condition has a density with respect to the Lebesgue measure, then μ_t has also a density with respect to the Lebesgue measure. In the following we make this assumption and we write with a slight abuse of notations $\mu_t(x)dx \equiv \mu_t(dx)$ for this density. The following elementary lemma shows that the ratio μ_t/π is controlled by the initial condition.

Lemma 2.14. *For any $t \in \mathbb{N}$, let $\mu_t \in \mathcal{P}_0^{\text{ac}}(E)$ be given by the recurrence relation (9) with initial condition $\mu_0 \in \mathcal{P}_0^{\text{ac}}(E)$. Then*

$$\inf_{x \in E} \frac{\mu_t(x)}{\pi(x)} \geq \inf_{x \in E} \frac{\mu_0(x)}{\pi(x)}, \quad \sup_{x \in E} \frac{\mu_t(x)}{\pi(x)} \leq \sup_{x \in E} \frac{\mu_0(x)}{\pi(x)}.$$

Proof. Since $W_\mu(x \rightarrow y) \geq 0$ and $\int_E W_\mu(x \rightarrow y) dy \leq 1$, this comes directly from the relation

$$\frac{\mathcal{T}[\mu](x)}{\pi(x)} = \left(1 - \int_E W_\mu(x \rightarrow y) dy \right) \frac{\mu(x)}{\pi(x)} + \int_E W_\mu(x \rightarrow y) \frac{\mu(y)}{\pi(y)} dy,$$

for all $\mu \in \mathcal{P}_0^{\text{ac}}(E)$ and $x \in E$. \square

The main result of this section is the following convergence result which is a direct consequence of the results presented in the following sections and discussed below.

Theorem 2.15. *Let $\Theta \in \mathcal{P}_0^{\text{ac}}(E)$ be a proposal distribution which satisfies Assumption 2.6 and h be an acceptance function which satisfies Assumption 2.4. Then there exist two (explicit) constants $C > 0$ and $\lambda \in (0, 1)$ which depend only on μ_0 , π and E such that*

$$\|\mu_t - \pi\|_{\text{TV}} \leq C(1 - \lambda)^{t/2} \quad (10)$$

Proof. Using the fact that the Total Variation norm is equal to the L^1 norm of the probability density functions and the Cauchy-Schwarz inequality, it holds that

$$\begin{aligned} \|\mu_t - \pi\|_{\text{TV}} &= \int_E |\mu_t(x) - \pi(x)| dx = \int_E \sqrt{\pi(x)} \sqrt{\pi(x)} \left| \frac{\mu_t(x)}{\pi(x)} - 1 \right| dx \\ &\leq \sqrt{2\mathcal{H}[\mu_t|\pi]}, \end{aligned}$$

where $\mathcal{H}[\mu_t|\pi] := \int_E \pi(x) |\mu_t(x)/\pi(x) - 1|^2 dx$ is the relative entropy which will be introduced later, see (13) with the function $\phi : s \mapsto \frac{1}{2}(s - 1)^2$. The conclusion follows from Proposition 2.20 which proves that the relative entropy decays geometrically in t . \square

Remark 2.16. The last inequality between the TV norm and the square root of the relative entropy is a simple form of a Csiszár-Kullback-Pinsker inequality, see [54, Appendix A] and [13].

Remark 2.17 (Convergence rate of nonlinear samplers). In this proof, the convergence rate λ is obtained by a crude estimate of the infimum of the jump rate $W_{f_t}(x \rightarrow y)$. We do not claim that this rate is optimal. In the degenerate case $\Theta_\mu(y|x) = \mu(y)$, the best rate obtained is equal to $h(1)$ by taking an initial condition arbitrarily close to π (see Remark 2.8). In Remark 3.1, it is shown that these proposal distributions can be obtained as the limit when $\sigma \rightarrow 0$ of the proposal distributions $\Theta_\mu(y|x) = K_\sigma \star \mu(y)$ where K_σ is a random-walk kernel with variance σ^2 . These proposal distributions are the simplest nonlinear analog of the random-walk Metropolis-Hastings algorithm with kernel K_σ (the proposal distribution in this case is $K_\sigma \star \delta_x(y)$). However in the nonlinear case, the convergence rate tends to a constant nonnegative value when $\sigma \rightarrow 0$ while in the random-walk Metropolis-Hastings case the convergence rate tends to 0 (see Theorem B.7).

In the linear case, that is when \mathcal{T} is the linear transition operator of the classical random walk Metropolis-Hastings algorithm, the convergence of the sequence of iterates of \mathcal{T} is studied in particular in [64] and more recently in [27] using analytical spectral methods. It is not possible to follow this strategy in the nonlinear case. In order to motivate our strategy, let us notice that the recurrence relation (9) can be interpreted as the explicit Euler discretization scheme

$$\frac{\mu_{t+1} - \mu_t}{\Delta t} = \mathcal{T}[\mu_t] - \mu_t, \quad (11)$$

of the nonlinear Partial Differential Equation

$$\partial_t f_t = \mathcal{T}[f_t] - f_t \quad (12)$$

with a constant time-step $\Delta t = 1$. The PDE (12) has a remarkable entropic structure which is detailed in Appendix B and which allows to prove that $f_t \rightarrow \pi$ as $t \rightarrow +\infty$. Entropy methods are by now a classical tool to study the long-time properties of both linear and nonlinear PDEs, see [54, 75]. The proof of

Theorem 2.15 is based on the adaptation of these ideas to the present discrete-time setting. A more detailed discussion of the links between (11) and (12) can be found in Appendix B.3. Entropy methods have been used previously in a similar context in [51] to prove the long-time convergence of a process obtained as the scaling limit of a particle-based Metropolis-Hastings algorithm [52] which, unlike the present case, is a continuous-time (nonlinear) diffusion process (see Remark 2.13).

2.4.2. Entropy and dissipation

For a given convex function $\phi : [0, +\infty) \rightarrow [0, +\infty)$ such that $\phi(1) = 0$, the relative entropy $\mathcal{H}[\mu|\pi]$ and dissipation $\mathcal{D}[\mu|\pi]$ of a probability density $\mu \in \mathcal{P}^{\text{ac}}(E)$ with respect to π are defined respectively by

$$\mathcal{H}[\mu|\pi] := \int_E \pi(x) \phi\left(\frac{\mu(x)}{\pi(x)}\right) dx, \quad \mathcal{D}[\mu|\pi] := - \int_E \phi'\left(\frac{\mu(x)}{\pi(x)}\right) (\mathcal{T}[\mu](x) - \mu(x)) dx. \quad (13)$$

Using the detailed balance property

$$\pi(x)W_\mu(x \rightarrow y) = \pi(y)W_\mu(y \rightarrow x),$$

for all $x, y \in E$ and $\mu \in \mathcal{P}^{\text{ac}}(E)$, it holds that

$$\begin{aligned} \mathcal{D}[\mu|\pi] &= \\ & \frac{1}{2} \iint_{E \times E} \pi(x)W_\mu(x \rightarrow y) \left(\frac{\mu(y)}{\pi(y)} - \frac{\mu(x)}{\pi(x)}\right) \left(\phi'\left(\frac{\mu(y)}{\pi(y)}\right) - \phi'\left(\frac{\mu(x)}{\pi(x)}\right)\right) dx dy, \end{aligned}$$

and thus $\mathcal{D}[\mu|\pi] \geq 0$ by convexity of ϕ .

In the following, we focus on the case $\phi(s) = \frac{1}{2}(s-1)^2$ for which we can prove the following crucial lemma. Note also that in this case the relative entropy is equal to a weighted L^2 norm and thus dominates the Total Variation norm between probability density functions.

Lemma 2.18. *Let $\phi(s) = \frac{1}{2}(s-1)^2$, then the sequence $(\mathcal{H}[\mu_t|\pi])_{t \in \mathbb{N}}$ is non-increasing.*

Proof. In this case,

$$\mathcal{H}[\mu|\pi] = \frac{1}{2} \int_E \pi(x) \left| \frac{\mu(x)}{\pi(x)} - 1 \right|^2 dx = \frac{1}{4} \iint_{E \times E} \pi(x)\pi(y) \left| \frac{\mu(x)}{\pi(x)} - \frac{\mu(y)}{\pi(y)} \right|^2 dx dy, \quad (14)$$

and

$$\mathcal{D}[\mu|\pi] = \frac{1}{2} \iint_{E \times E} \pi(x)W_\mu(x \rightarrow y) \left| \frac{\mu(y)}{\pi(y)} - \frac{\mu(x)}{\pi(x)} \right|^2 dx dy. \quad (15)$$

Moreover, since $\phi(s) = \frac{1}{2}(s-1)^2$ is a polynomial of order 2, the exact Taylor expansion

$$\phi(v) = \phi(u) + \phi'(u)(v-u) + \frac{1}{2}\phi''(u)(v-u)^2,$$

yields

$$\mathcal{H}[\mu_{t+1}|\pi] - \mathcal{H}[\mu_t|\pi] = -\mathcal{D}[\mu_t|\pi] + \frac{1}{2} \int_E \pi(x) \left| \frac{\mu_{t+1}(x)}{\pi(x)} - \frac{\mu_t(x)}{\pi(x)} \right|^2 dx. \quad (16)$$

Using that $\mu_{t+1} = \mathcal{T}[\mu_t]$ and the definition of \mathcal{T} , it holds that

$$\begin{aligned} \int_E \pi(x) \left| \frac{\mu_{t+1}(x)}{\pi(x)} - \frac{\mu_t(x)}{\pi(x)} \right|^2 dx \\ = \int_E \pi(x) \left| \int_E W_{\mu_t}(x \rightarrow y) \left(\frac{\mu_t(y)}{\pi(y)} - \frac{\mu_t(x)}{\pi(x)} \right) dy \right|^2 dx. \end{aligned}$$

By the Cauchy-Schwarz inequality, we get

$$\begin{aligned} \int_E \pi(x) \left| \frac{\mu_{t+1}(x)}{\pi(x)} - \frac{\mu_t(x)}{\pi(x)} \right|^2 dx \\ \leq \iint_{E \times E} \pi(x) W_{\mu_t}(x \rightarrow y) \left| \frac{\mu_t(y)}{\pi(y)} - \frac{\mu_t(x)}{\pi(x)} \right|^2 \int_E W_{\mu_t}(x \rightarrow z) dz dx dy. \end{aligned}$$

Reporting into (16) and using (15) yields

$$\begin{aligned} & \mathcal{H}[\mu_{t+1}|\pi] - \mathcal{H}[\mu_t|\pi] \\ & \leq -\frac{1}{2} \iint_{E \times E} \pi(x) W_{\mu_t}(x \rightarrow y) \left(1 - \int_E W_{\mu_t}(x \rightarrow z) dz \right) \left| \frac{\mu_t(y)}{\pi(y)} - \frac{\mu_t(x)}{\pi(x)} \right|^2 dx dy. \end{aligned} \quad (17)$$

Since $W_{\mu_t}(x \rightarrow y) \geq 0$ and $\int_E W_{\mu_t}(x \rightarrow z) dz \leq 1$, the right-hand side is non negative which concludes the proof. \square

2.4.3. Exponential decay of the entropy

Under the Assumptions 2.4 and 2.6, it is possible to improve the result of Lemma 2.18 and to prove a quantitative exponential decay result. Since the entropy is given by (14), owing to (17), the goal is to bound from below by a multiple of $\pi(y)$ the quantity

$$W_{\mu_t}(x \rightarrow y) \left(1 - \int_E W_{\mu_t}(x \rightarrow z) dz \right),$$

for all $x, y \in E$ and uniformly in t .

The main consequence of Assumption 2.6 is the following lower bound.

Lemma 2.19. *Let $h : [0, +\infty) \rightarrow [0, 1]$ be a continuous non-decreasing acceptance function which satisfies the relation (2). Let Θ be a proposal distribution which satisfies Assumption 2.6. Then for all $\mu \in \mathcal{P}_0^{\text{ac}}(E)$ and all $x, y \in E$,*

$$W_{\mu}(x \rightarrow y) \geq \lambda(\mu) \pi(y), \quad \lambda(\mu) := c^- \left(\inf_{x \in E} \frac{\mu(x)}{\pi(x)} \right) h(1) \in (0, 1].$$

As a consequence, using the fact that c^- is non decreasing and Lemma 2.14, for any $t \in \mathbb{N}$, it holds that

$$W_{\mu_t}(x \rightarrow y) \geq \lambda_0 \pi(y), \quad \lambda_0 := c^- \left(\inf_{x \in E} \frac{\mu_0(x)}{\pi(x)} \right) h(1).$$

Proof. Let us first prove that for any bounded interval $[a, b]$ with $a > 0$, it holds that

$$\inf_{(x,y) \in [a,b]^2} yh \left(\frac{x}{y} \right) = ah(1).$$

Since h is continuous non-decreasing, for each $y \in [a, b]$, the function $x \in [a, b] \mapsto yh(x/y)$ has a minimum in $x = a$. This shows that the minimum on the function $(x, y) \in [a, b]^2 \mapsto yh(x/y)$ is attained on the segment $\{(a, y), y \in [a, b]\}$. Since $yh(a/y) = ah(y/a)$, the same reasoning shows that this minimum is attained when $y = a$. The conclusion follows. Then, using Assumption 2.6 and applying this result with $a = c^- \left(\inf_{x \in E} \frac{\mu(x)}{\pi(x)} \right)$ yields

$$\begin{aligned} W_{\mu}(x \rightarrow y) &= \Theta_{\mu}(y|x)h(\alpha_{\mu}(x, y)) = \frac{\Theta_{\mu}(y|x)}{\pi(y)} h \left(\frac{\Theta_{\mu}(x|y)\pi(y)}{\pi(x)\Theta_{\mu}(y|x)} \right) \pi(y) \\ &\geq c^- \left(\inf_{x \in E} \frac{\mu(x)}{\pi(x)} \right) h(1)\pi(y). \square \end{aligned}$$

We are now ready to prove that in this case, the relative entropy converges exponentially fast towards zero.

Proposition 2.20. *Under Assumption 2.4 and Assumption 2.6, there exists a constant $\lambda \in (0, 1)$ such that for all $t \in \mathbb{N}$,*

$$\mathcal{H}[\mu_t|\pi] \leq (1 - \lambda)^t \mathcal{H}[\mu_0|\pi].$$

Proof. From the relation (17) and using the assumptions on h and Θ , it holds that

$$\mathcal{H}[\mu_{t+1}|\pi] - \mathcal{H}[\mu_t|\pi] \leq -2\lambda_0(1 - \eta)\mathcal{H}[\mu_t|\pi].$$

Moreover by definition of λ_0 , it holds that $\lambda_0 \leq \eta$ and thus $\lambda := 2\lambda_0(1 - \eta) < 1$. We deduce that

$$\mathcal{H}[\mu_{t+1}|\pi] \leq (1 - \lambda)\mathcal{H}[\mu_t|\pi],$$

and the conclusion follows. \square

3. Some collective proposal distributions

The proposal distribution can be fairly general and so far, we have not detailed how to choose it. Several choices of proposal distributions are gathered in this section.

This proposal should use the maximum of information coming from the value of the target π , in order to increase the fitness of Θ_{μ} to the true distribution.

However, the proposal must also allow for some exploration of the parameter space, this problem is addressed for some of the proposals.

Here, we only intend to present some of the proposals possible, each one having pros and cons, on the theoretical and practical side.

In view of Theorem 2.12, we can see each of the proposal distributions presented in this section either as a specific choice of nonlinear kernel (1) with its associated nonlinear process or as its particle approximation given by Algorithm 1. From this second perspective the proposal distribution can be seen as a specific interaction mechanism between the particles. More specifically, it depicts a specific procedure which can be interpreted as “information sharing” between the particles: given the positions of *all* the N particles at a given time, we aim at constructing the best interaction mechanism which will favour a specific aspect such as acceptance, convergence speed, exploration etc. By analogy with systems of swarming particles which exchange local information (here, the local value of the target distribution) to produce global patterns (here, a globally well distributed sample), we call this class of proposal distributions **collective**. The class of methods introduced will be referred as Collective Monte Carlo methods (CMC). On the contrary, the nonlinear kernels introduced in [3] do not belong to this class as explained in Appendix C.1. For each proposal we give an implementation which, starting from population of particles $(X^1, \dots, X^N) \in E^N$, returns a proposal Y .

Although our theoretical results (Theorem 2.12 and Theorem 2.15) are general enough to encompass almost all of the proposal distributions described here (see Section 2.1), the validity and numerical efficiency of each of them will be assessed in Section 5 on various examples of target distributions.

3.1. Metropolis-Hastings proposal (PMH)

Proposal distribution The classical Metropolis-Hastings algorithm fits into our formalism, with

$$\Theta_\mu(dy|x) = q(y|x)dy,$$

where q is a fixed random walk kernel which does not depend on μ .

Particle implementation In this case, Algorithm 1 reduces to the simulation of N independent Metropolis-Hastings chains in parallel.

3.2. Convolution kernel proposal (Vanilla CMC)

Proposal distribution Let us then consider the following proposal distribution given by the convolution:

$$\Theta_\mu(dy|x) = K \star \mu(y)dy := \left(\int_E K(y-z)\mu(dz) \right) dy, \quad (18)$$

where K is a fixed *interaction kernel*, that is a (smooth) radial function which tends to zero at infinity. Typical examples are $K(x) = \mathcal{N}(0, \sigma^2 Id)$ and $K(x) =$

$\frac{1}{|B_\sigma(0)|} \mathbf{1}_{B_\sigma(0)}(x)$, where $\sigma > 0$ is fixed and $B_\sigma(0)$ denotes the ball of radius $\sigma > 0$ centred at 0 in \mathbb{R}^d . Note that the proposal distribution does not depend on the starting point x . It may happen that the proposed state falls outside E . In this case it will be rejected since π is equal to zero outside E . One can therefore take equivalently $\Theta_\mu(dy|x) \propto K \star \mu(y) \mathbf{1}_E(y) dy$ (the same remark holds for the other collective proposal distributions).

Particle implementation At each time step t , each particle i samples uniformly another particle j and then draw a proposal $Y_t^i \sim K(X_t^j, dy)$ which can be seen as a “mutation” of X_t^j . This “resampling with mutation” procedure is somehow similar to a (genetic) Wright-Fisher model (see for instance [33] for a review of genetic models). Since a “mutation” may or may not be accepted, it can be described as a *biased Wright-Fisher model*.

The collective aspect is twofold: first the proposal distribution allows large scale move on all the domain filled by the N particles; then during the acceptance step, for a particle in X and a proposal in Y the acceptance ratio can be understood as a measure of discrepancy between the target ratio $\frac{\pi(Y)}{\pi(X)}$ and the observed ratio $\frac{K \star \hat{\mu}^N(Y)}{K \star \hat{\mu}^N(X)}$ between the (average) number of particles around Y and the (average) number of particles around X . In the linear Metropolis-Hastings case with a symmetric random-walk kernel, the acceptance ratio only takes into account the target ratio. As a consequence, the acceptance probability of a proposal state depends not only on how “good” it is when looking at the values of π but also on how many particles are (or are not) already around this proposal state compared to the present state (and therefore on how accepting the proposal would improve the current global state of the system).

Remark 3.1 (Moderate interaction, part 1). When $\sigma \rightarrow 0$ we obtain the degenerate proposal distribution $\Theta_\mu(dy|x) = \mu(dy)$ (which does not satisfy the assumption $\Theta_\mu(dy|x) \in \mathcal{P}_0^{\text{ac}}(E)$ in general). It would not make sense to take this proposal distribution at the particle level in Theorem 2.11. However, it makes sense to consider the case $\Theta_\mu(y|x) = \mu(y) dy$ in the nonlinear kernel (1) where $\mu \in \mathcal{P}_0^{\text{ac}}(E)$. This degenerate proposal distribution still satisfies the assumptions of Theorem 2.15 and could lead to a better rate of convergence (see Remark 2.17). It is thus worth mentioning that this degenerate proposal distribution can also be obtained as the many-particle limit of a system of particles under an additional *moderate interaction* assumption [68, 53, 28]. See Remark A.3 for additional details.

Draw uniformly $j \in \{1, \dots, N\}$;
 Draw $e \sim K$;
 Set $Y = X^j + e$;

Algorithm 2: Vanilla CMC proposal generation

3.3. Markovian Mixture of Kernels Proposal (MoKA and MoKA-Markov)

Proposal distribution A limitation of the Convolution Kernel Algorithm 2 is the fixed size of the interaction kernel. A remedy is given by the following collective proposal distribution which is a convolution with a mixture of kernels (with different sizes) with (potentially) nonlinear mixture weights:

$$\Theta_\mu(dy|x) = \sum_{p=1}^P \alpha_p(\mu) K_p \star \mu(y) dy,$$

A possible choice for the weights is to take a solution of the following minimisation problem:

$$\min_{\alpha \in S^p} \int_E \phi \left(\frac{\frac{\sum_p \alpha_p K_p \star \mu(x)}{\int_E \sum_p \alpha_p K_p \star \mu(x') \mu(dx')}}{\frac{\pi(x)}{\int_E \pi(x') \mu(dx')}} \right) \frac{\pi(x)}{\int_E \pi(x') \mu(dx')} \mu(dx), \quad (19)$$

where S_p denotes the p -simplex and where ϕ is convex non-negative such that $\phi(1) = 0$. Typically $\phi(s) = s \log s - s + 1$. In this case, it corresponds to minimising the Kullback-Leibler divergence between the probability distributions $\frac{\sum_p \alpha_p K_p \star \mu(x) \mu(dx)}{\int_E \sum_p \alpha_p K_p \star \mu(x') \mu(dx')}$ and $\frac{\pi(x) \mu(dx)}{\int_E \pi(x') \mu(dx')}$. Note that if μ is interpreted as the current knowledge (i.e. the distribution of the particles or of the associated nonlinear chain), then this minimisation problem can be understood as choosing the proposal distribution which is the closest to π relatively to μ . In our experiments, we found that choosing $\phi(s) = |1 - s|$, leads to similar, slightly better results compared to $\phi(s) = s \log(s) - s + 1$. Moreover, this choice is also numerically more stable so we chose to implement this version, that we call Markovian Mixture of Kernels (MoKA-Markov).

Another choice for the weights, which is **non-Markovian**, is to take α_p proportional to the geometric mean of the acceptance ratio of the particles which have chosen the kernel p at the previous iteration. This method will be referred as Mixture of Kernels Adaptive CMC (MoKA). It shares similarities with the D -kernel algorithm of [30] and the arguments developed by the authors suggest that the two versions (Markovian and MoKA) may be asymptotically equivalent. The proof is left for future work.

Particle implementation Same as in Algorithm 2 but with an additional step to choose a “mutation kernel” at each proposal step. The computation of the weights of the mixture can be done in a fully Markovian way at the beginning of each iteration before the proposal step or, in MoKA, are computed at the end of the iteration and used at the next iteration.

In Section 5 we will show that this algorithm can favour initial exploration

if initially the particles are in an area of low potential.

Compute the weights $\alpha_1, \dots, \alpha_P$ using (19);
 Draw $p \in \{1, \dots, P\}$ with probability $(\alpha_1, \dots, \alpha_P)$;
 Draw uniformly $j \in \{1, \dots, N\}$;
 Draw $e \sim K_p$;
 Set $Y = X^j + e$;

Algorithm 3: Markovian Mixture of Kernels proposal generation

3.4. Kernelised importance-by-deconvolution sampling (KIDS)

Proposal distribution Algorithms based on a simple convolution operator (Algorithms 2 and 3) keep a “blind” resampling step. In order to improve the convergence speed of such algorithms one may want to favour the selection of “good” states. For a fixed interaction kernel K , one can choose a proposal distribution of the form:

$$\Theta_\mu(y|x) = K \star \nu_\mu(y),$$

where ν_μ solves the following deconvolution problem with an absolute continuity constraint:

$$\min_{\nu \ll \mu} \int_E \log \left(\frac{K \star \nu(x)}{\pi(x)} \right) K \star \nu(x) dx. \quad (20)$$

That is, we are looking for a weight function $w \geq 0$ which satisfies the constraint

$$\int_E w(x) \mu(dx) = 1$$

and such that the measure defined by $\nu_\mu(A) = \int_A w(x) \mu(dx)$ minimises the KL divergence above. The function w is the Radon-Nikodym derivative of ν with respect to μ . In other words the proposal distribution focuses on the parts of the support of μ which are “closer” to π . Note also that this optimization procedure does not depend on the normalisation constant of π . Similarly to Algorithm 3, the idea is to minimize a KL divergence between a set of proposal distributions and the distribution π . However, in the minimization problem 19 the goal was to optimize the choice of the convolution kernel K among a given family while in the optimization problem 20, the goal is to optimize the support of the distribution μ to focus on the parts which are “closer” to π . Finally, it should be noted that the two minimization procedures are not excluding and that both can be conducted at the same time.

Remark 3.2. Although the proposal distributions KIDS and MoKA-Markov perform well in practice, we did not manage to prove the regularity Assumption 2.3 because we do not know the regularity of the solutions of the minimization problems (19) and (20). Since the weights α_p are bounded in (19), MoKA-Markov satisfies Assumption 2.2 but we did not manage to prove this

assumption for KIDS for the same reason as before. However, Vanilla CMC satisfies all the assumptions since the proposal distribution is a simple convolution with a smooth kernel.

Particle implementation In the case of an empirical measure $\hat{\mu}^N = \frac{1}{N}\delta_{X^i}$, the Radon-Nikodym weight function w can simply be seen as a vector of N weights $(w^1, \dots, w^N) \in [0, 1]^N$ such that $\sum_i w^i = 1$ and the measure $\nu_{\hat{\mu}^N}$ is thus a weighted empirical measure:

$$\nu_{\hat{\mu}^N} := \sum_{i=1}^N w^i \delta_{X^i}. \quad (21)$$

The deconvolution procedure gives more weight to the set of particles that are already “well distributed” according to π . These particles are thus more often chosen in the Wright-Fisher resampling step (see Algorithm 2).

Note that although the weighted empirical measure proposal (21) is very reminiscent of an Importance Sampling procedure, the computation of the weights here follows from a completely different idea.

In practice we solve the deconvolution problem using the Richardson-Lucy algorithm [71, 59] (also known as the Expectation Maximisation algorithm). See for instance [67, Section 5.3.2] where it is proved that the iterative algorithm below converges towards a minimiser of the Kullback-Leibler divergence (20) in the case of an empirical measure μ . Note that the computation of the weights (Richardson-Lucy loop) can be done before the resampling step.

```

Set  $w_i^{(0)} = 1$  for all  $i \in \{1, \dots, N\}$ ;
for  $s = 0$  to  $S - 1$  do
  For all  $i \in \{1, \dots, N\}$ , update the weight by:
   $w_i^{(s+1)} = w_i^{(s)} \frac{\sum_{j=1}^N \frac{\pi(X_i^j)K(X_i^j - X_i^i)}{\sum_{k=1}^N w_k^{(s)}K(X_i^j - X_k^i)}$ ;
end
Set  $w_i = w_i^{(S)}$  for all  $i \in \{1, \dots, N\}$ ;
Normalize the weights  $(w_1, \dots, w_N)$ ;
Draw  $j \in \{1, \dots, N\}$  with probability  $(w_1, \dots, w_N)$ ;
Draw  $e \sim K$ ;
Set  $Y = X^j + e$ ;

```

Algorithm 4: Kernel Importance-by-Deconvolution Sampling proposal generation

3.5. Bhatnagar-Gross-Krook sampling (BGK)

Proposal distribution In Algorithms 2, 3 and 4, the proposal distribution is based on a (mixture of) symmetric kernels: this symmetry property is reflected in the proposal distribution and might not well represent the local properties of the target distribution. In dimension $d \geq 2$, we can adopt a different strategy by

sampling proposals from a multivariate Gaussian distribution with a covariance matrix that is computed locally. An example is given by the following proposal distribution:

$$\Theta_\mu(dy|x) = \left(\int_E G_{\widehat{\Sigma}_\mu(z)}(\widehat{m}_\mu(z) - y)\mu(dz) \right) dy,$$

where

$$\widehat{m}_\mu(z) = \frac{1}{\int_E K(z - z')\mu(dz')} \int_E K(z - z')z'\mu(dz'), \quad (22)$$

and

$$\widehat{\Sigma}_\mu(z) = \frac{1}{\int_E K(z - z')\mu(dz')} \int_E K(z - z')(z' - \widehat{m}_\mu(z))(z' - \widehat{m}_\mu(z))^T \mu(dz'), \quad (23)$$

and where K is a fixed interaction kernel. This proposal distribution and the associated transition operator are reminiscent of a Bhatnagar-Gross-Krook (BGK) type operator [12].

In the particular case of $K(x, y) \equiv 1$, we have a more simple algorithm which can be interpreted as a Markovian version of the Adaptive Metropolis-Hastings algorithm introduced by [44]. However such algorithm does not benefit from the appealing properties of local samplers. Indeed, when the target distribution is multimodal, we can take more advantageously as interaction kernel $K(x) \propto \mathbf{1}_{|x| < \sigma}$, where the threshold σ allows the proposal to be adapted to the local mode. Note that moving across the modes is still possible thanks to the choice of another particle at the first step (see proposal algorithm below). The main issue is the choice of σ that must be higher than the size of the modes but smaller than the distance between the modes.

Particle implementation Each particle, say X_t^i samples another particle, say X_t^j . Then we compute the local mean and covariance around X_t^j and we draw a proposal Y_t^i for X_t^i according to a Normal law with the locally computed parameters. As before it may be cheaper to compute and store all the local means and covariances before the resampling loop.

Draw $j \in \{1 \dots, N\}$ uniformly;
 Compute $\kappa = \sum_i K(X_t^i, X_t^j)$;
 Compute the local mean $\widehat{m}_{X_t^j} = \frac{1}{\kappa} \sum_i K(X_t^i, X_t^j)X_t^i$;
 Compute the local covariance
 $\widehat{\Sigma}_{X_t^j} = \frac{1}{\kappa} \sum_i K(X_t^i, X_t^j)(X_t^i - \widehat{m}_{X_t^j})(X_t^i - \widehat{m}_{X_t^j})^T$;
 Draw $Y \sim \mathcal{N}(\widehat{m}_{X_t^j}, \widehat{\Sigma}_{X_t^j})$;

Algorithm 5: BGK proposal generation

4. GPU implementation

A major bottleneck for the development of collective samplers based on interacting particles is the computational cost associated to the simulation of N^2

pair-wise interactions between particles at every time step. To overcome this issue, many methods have been proposed throughout the years, often using techniques borrowed from N -body simulations or molecular dynamics and developed since the 60's. Let us cite for instance the Verlet and cell-list methods for short-ranged interactions [83, 47, 76], the super-particle and fast multipole methods [7, 74, 42] for long-ranged interactions or the more recent Random Batch Method introduced by [50]. In the MCMC literature, such methods have been applied for instance to particle smoothing in [56].

Most importantly, the last decade has seen the introduction of massively parallel GPU chips. Beyond the training of convolutional neural networks, GPUs can now be used to simulate generic particle systems at ever faster rates and, noticeably, without relying on numerical approximations, contrary to classical N -body simulation algorithms. One goal of the present paper is to discuss the consequences of this hardware revolution on Monte Carlo sampling, from the practitioner point of view. The use of GPU in this context traces back at least to [58]. Although in this article, the authors have focused on traditional SMC and MCMC methods which are often embarrassingly parallelizable (with a linear complexity in N), it should be noted that, as such, GPU programming is not a trivial task. In particular, let us quote [58] on the CUDA language, which is the development environment for the Nvidia GPU chips: “*a programmer proficient in C should be able to code effectively in CUDA within a few weeks of dedicated study*”. In the present article, our goal is to advocate the use of very recent GPU routines developed for high-level languages (in particular Python) which allow the end-user to benefit from GPU acceleration without any particular knowledge or experience in the CUDA programming language. Using these techniques, our implementation offers a transparent Python interface for massively parallelized MCMC methods, including both the traditional methods (Metropolis-Hastings algorithm, SMC) which are linear in N and the population based methods (CMC-related methods) with a quadratic footprint. For the latter methods, the main specific technical difficulties are briefly described below.

From the evaluation of kernel densities to the computation of importance weights in the Richardson-Lucy loop (Algorithm 4), the bottleneck of the Collective Proposal framework presented in Section 3 is the computation of an off-grid discrete convolution of the form:

$$a_i = \sum_{j=1}^N K(X_t^i, X_t^j) b_j \quad (24)$$

for all i between 1 and N , with arbitrary weights b_j . All CMC-related methods have $\mathcal{O}(N^2)$ time complexity, with a constant that is directly related to the efficiency of the underlying implementation of the “kernel sum” (24).

In the machine learning literature, this fundamental operation is often understood as a matrix-vector product between an N -by- N *kernel matrix* ($K(X_t^i, X_t^j)$) and a vector (b_j) of size N . Common implementations generally rely on linear algebra routines provided e.g. by the PyTorch library [69] and have a quadratic

memory footprint: even on modern hardware, this prevents them from handling populations of more than $N = 10^4$ to 10^5 samples without making approximations or relying on specific sub-sampling schemes [86].

Fortunately though, over the last few years, efficient GPU routines have been developed to tackle computations in the mould of (24) with maximum efficiency. These methods can be accessed through the KeOps extension for PyTorch [70], NumPy [81], R [79] and Matlab [60], that is developed by the third author [21, 35, 37] and freely available at <https://www.kernel-operations.io>. In practice, this library supports arbitrary kernels on the GPU with a linear memory footprint, log-domain implementations for the sake of numerical stability and outperforms baseline PyTorch GPU implementations by one to two orders of magnitude. Computing the convolution of (24) with a cloud of $N = 10^6$ points in dimension 3 takes no more than 1s on a modern chip. Pairwise interactions between populations of $N = 10^5$ samples may also be simulated in no more than 10ms, without making any approximation.

We run all the tests of Section 5 using a single gaming GPU, the Nvidia GeForce RTX 2080 Ti on the GPU cluster of the Department of Mathematics at Imperial College London. With 10^4 to 10^6 particles handled at any given time, our simulations run in a handful of seconds at most, with performances that enable the real-time sampling of large populations of *interacting* samples. Our code is available on the third author's GitHub page.

5. Numerical experiments

We now run experiments on several target distributions in low and moderately high dimension.

Our code and its documentation are available online at

<http://www.kernel-operations.io/monaco>

and on the GitHub page of the third author. All our experiments can be run on Google Colaboratory (colab.research.google.com) with a GPU within a few seconds to a few minutes, depending on the method and number of independent runs.

5.1. Methodology

5.1.1. Methods in competition

Among the variants of our methods, we show the results for

- The Vanilla CMC Algorithm 2 (named CMC in the following) with a fixed interaction kernel K which is taken equal to the indicator of the ball centered at zero with a given radius R . The quadratic cost of this algorithm is only due to the computation of discrete convolutions.

- The Markovian version of the MoKA algorithm 3 with several interaction kernels and adaptive weights (named MoKA Markov in the following). In addition to the computation of the discrete convolutions, solving the optimization problem (19) slightly increases the computational cost.
- The non-Markovian version of MoKA outline in Section 3.3 (named MoKA in the following). Thanks to the simplified computation procedure of the mixture weights compared to MoKA Markov, this algorithm has the same computation cost as CMC.
- The hybrid algorithm which features both the adaptive interaction kernels of MoKA (non-Markovian) and an adaptive weighting of the particles obtained with Algorithm 4. It will be named MoKA KIDS in the following. It is a significantly heavier method than CMC, as it relies on the Richardson-Lucy iterations to optimise the deconvolution weights. In the experiments shown below, this algorithm is roughly five times slower than CMC in terms of computation time (but needs less iterations to converge).

Please note that we do not include the BGK and KIDS samplers in these experiments: although we believe that the ideas behind these methods are interesting enough to justify their presentation, we observe that they generally do not perform as well as the other CMC variants and leave them aside for the sake of clarity.

For the competitors, we consider the three following algorithms.

- As a baseline, we include a parallel implementation of Metropolis-Hastings (named PMH in the following) with a number of parallel runs that is equal to the number of particles in CMC and its variants.
- The Safe Adaptive Importance sampling (SAIS) algorithm from [25], which is one of the state-of-the-art importance sampling based methods. It is detailed in Section C.2. We note that unlike all the other Markovian methods, the memory and computational cost of SAIS significantly increases with the number of iterations since the number of particles is not constant – to the best of our knowledge, no procedure has yet been proposed to remove particles with time. We have implemented the so-called mini-batch version as proposed in [25] but without the sub-sampling procedure since the computation time is not a problem with our fast GPU implementation. The parameters (in particular the kernel sequence and mixing weights) are chosen according to the recommendation of the authors as detailed in Section C.2.
- A simple (non-adaptive) Sequential Monte Carlo algorithm as described in Section C.3 with a linear tempering scheme for the potentials and a simple Metropolis-Hastings mutation kernel.

We emphasize the fact that in our implementation, all the methods, including PMH and SMC which have a linear complexity in N , are parallelized on a GPU using PyTorch [70] and KeOps [21]. All the methods are therefore comparable in terms of computation time.

5.1.2. Initialization and target distributions

We have chosen various target distributions in order to illustrate the main difficulties and to outline the strengths and weaknesses of the various methods in competition. As an illustration we start with a banana-shaped distribution in low dimension (Section 5.2) then we consider several variants of the classical Gaussian mixture example introduced by [17] in moderately high dimension (up to dimension 13), see Sections 5.3 and 5.5.

Another example of target distribution is included in Appendix E and we include results for the estimation of the normalizing constant of some of the target distributions below in Appendix D. Experiments in non-Euclidean spaces such as the Poincaré hyperbolic plane and the group of 3D rotation matrices are available online in our documentation.

We always clip distributions on the unit (hyper-)cube $[0, 1]^d$. For all the methods we have considered two types of initialization.

- A simple initialization where all the N particles are sampled uniformly and independently in the unit hyper-cube.
- A difficult initialization with N particles independently distributed according to

$$X_0^i = (0.9, \dots, 0.9)^T + 0.1U_0^i \quad (25)$$

where $U_0^i \sim \mathcal{U}([0, 1]^d)$, that is all the particles are located in a small corner of the domain.

5.1.3. Energy distance

We compare the results in term of Energy distance [72] between a true sample of the target distribution and the population of particles at each step.

The Energy Distance (ED) between two probability distributions μ, ν in \mathbb{R}^d can be defined by several equivalent formulas. Perhaps, the simplest form is

$$\mathcal{E}(\mu, \nu) = \mathbb{E}|X - Y| - \frac{1}{2}\mathbb{E}|X - X'| - \frac{1}{2}\mathbb{E}|Y - Y'|,$$

where $X, X' \sim \mu$ and $Y, Y' \sim \nu$ are independent. In the machine learning literature, it is often understood as a kernel norm comparable to other popular distances such as the Wasserstein distance [36, 87]. This interpretation comes from the formulation

$$\mathcal{E}(\mu, \nu) = \frac{1}{2}\langle \mu - \nu, K \star (\mu - \nu) \rangle,$$

where $K(x, y) = -|x - y|$ and where we recall that the convolution operator is defined by $K \star \mu(x) := \int_{\mathbb{R}^d} K(x, y)\mu(dy)$. Among all the possible kernels, this particular choice is the simplest one which ensures that the distance between two Dirac masses is equal to the distance between the points. Finally, by writing

this latter expression in the Fourier space, the ED can also be understood as a weighted L^2 distance between the characteristic functions:

$$\mathcal{E}(\mu, \nu) = C_d \int_{\mathbb{R}^d} \frac{|\varphi_\mu(z) - \varphi_\nu(z)|^2}{|z|^{d+1}} dz,$$

for a constant $C_d > 0$ and where $\varphi_\mu(z) = \mathbb{E}_{X \sim \mu}[e^{iz \cdot X}]$ is the characteristic function of the probability measure μ . Note also that in dimension 1, the Energy Distance is equal to the Cramér's squared L^2 distance between the cumulative distribution functions F_μ and F_ν of μ and ν :

$$\mathcal{E}(\mu, \nu) = \int_{\mathbb{R}} |F_\mu(z) - F_\nu(z)|^2 dz.$$

It can be shown that the ED satisfies all the properties of a metric and in particular $\mathcal{E}(\mu, \nu) = 0$ if and only if $\mu = \nu$.

The Energy Distance can serve as a baseline to test the null hypothesis that two random vectors $\mathcal{X} = (X_i)_{i \in \{1, \dots, n\}}$, $X_i \in \mathbb{R}^d$ and $\mathcal{Y} = (Y_j)_{j \in \{1, \dots, m\}}$, $Y_j \in \mathbb{R}^d$ are sampled from the same distribution. For this, we simply compute the ED between the associated empirical measures $\hat{\mu}_{\mathcal{X}} = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $\hat{\mu}_{\mathcal{Y}} := \frac{1}{m} \sum_{j=1}^m \delta_{Y_j}$, namely:

$$\begin{aligned} \mathcal{E}_{n,m}(\mathcal{X}, \mathcal{Y}) &:= \mathcal{E}(\hat{\mu}_{\mathcal{X}}, \hat{\mu}_{\mathcal{Y}}) \\ &= \frac{1}{nm} \sum_{i,j} |X_i - Y_j| - \frac{1}{2n^2} \sum_{i,j} |X_i - X_j| - \frac{1}{2m^2} \sum_{i,j} |Y_i - Y_j|. \end{aligned}$$

It is shown in [77, 72] that this quantity can serve as a test statistic for equal distributions which implies in particular that $\mathcal{E}_{n,m}$ tends to 0 as $n, m \rightarrow +\infty$ when \mathcal{X} and \mathcal{Y} are i.i.d. samples.

In order to measure the quality of the various samplers, we will compare the ED between the output samples and a true sample of the target distribution, both of size N . We consider that the sampler successfully manage to sample the target distribution when the final samples fall within the 90% prediction interval (computed with 1000 independent observations) around the average Energy Distance between two i.i.d. exact samples of size N . It means that the sample cannot be distinguished from a true sample.

5.2. Banana-shaped distribution

As a first illustrative example in dimension 2, we consider a mixture of three (equally weighted) Gaussian distributions and a banana-shaped distribution, as introduced in [43, 25]. The banana-shaped distribution is a classical test case for sampling algorithms. It is often considered difficult to sample from this distribution due to its curved, oblong shape. In dimension 2, however, it is easy to find an appropriate range of parameters such that all the methods including PMH, SAIS, SMC or a standard rejection sampler achieve very good




Fig 1: The level sets of the target distribution (red) and the final states of different algorithms after 50 iterations. The particles are displayed in blue and the rejected proposals (at the previous iteration) in green. There are $N = 10^4$ particles. (a) Vanilla CMC without an exploration proposal and a too small proposal distribution. The particles concentrate on the top-right mode and cannot escape it because the proposals cannot reach the other modes. (b) PMH with a large proposal distribution reduces to a rejection sampler (c) CMC with a large exploration proposal. Although the target distribution is correctly sampled, many proposals in green are never accepted. (d) MoKA Markov with a mixture of small and large proposal distributions. The final state for MoKA and MoKA KIDS is similar. Thanks to the automatic choice of a small interaction kernel, most of the proposals are accepted.

results within a computation time which does not exceed a few seconds. Thus we mainly use this introductory example as a qualitative illustration of the behaviour of the collective algorithms CMC, MoKA, MoKA Markov and MoKA KIDS that we simply compare to PMH. In dimension 2, it is indeed possible to directly visualize the results for these algorithms, as shown in Figure 1 and in the supplementary videos available on the GitHub page of the project. More in-depth quantitative comparisons on much more difficult test cases in higher dimensions will be presented in the following sections.

Starting from the initialization (25), the particles are initially very close to one mode and if the radius of the interaction kernel is taken too small, then the particles will stay in this mode forever, as it can be seen in Figure 1a. It is thus necessary to add an exploration mechanism. For the three MoKA variants, this can be done automatically by simply adding a large proposal distribution within the proposal mixture. Alternatively, it is possible to consider a simple upgraded version of the Vanilla CMC algorithm (18) by taking the proposal distribution

$$\Theta_{\mu}(dy|x) = (1 - \varepsilon)K \star \mu(y)dy + \varepsilon Q(dy|x),$$

where the exploration proposal Q is for instance $Q(dy|x) \propto \exp(-\frac{|y-x|^2}{2\sigma^2})$ with σ^2 large. In the remaining of this section, we will always consider this proposal (and simply call it CMC) with $\varepsilon = 0.01$, $\sigma = 0.5$ and the radius of the interaction kernel K equal to 0.1. For the MoKA algorithms, we consider a mixture of

proposal distributions of sizes 0.01, 0.03, 0.1 and 0.5. As a control system, we consider PMH with a mutation kernel of size 0.5.

The results are visualized in dimension 2 in the supplementary videos available on the GitHub page of the project. We summarize below the main observations.

- In PMH (Figure 1b), due to the large mutation kernel, blind proposals are made all over the space and the algorithm behaves essentially like a rejection algorithm. Consequently, the acceptance rate is very low, about 3%. Note that PMH with a too small mutation kernel would lead to the same situation as shown in Figure 1a.
- In CMC, the particles move collectively like a coherent swarm. They initially massively concentrate on the first mode that they find (on the top-right corner). Thanks to the large exploration proposal, a few particles are able to find the other modes. These well-located particles then slowly attract the other ones and the correct balance between the different modes is progressively attained. At the end of the simulation, there are still many proposals which are not accepted because the interaction kernel is too large. The final acceptance rate is about 16%.
- In MoKA Markov, a very large exploration proposal is initially selected (see Figure 2b) and similarly to PMH, the particles are quickly scattered all over the space. However, unlike PMH, once particles are present in all the modes, a better adapted small interaction kernel is chosen (see Figure 2b) and similarly to CMC, the balance between the modes is progressively but much more quickly attained. The final acceptance rate is about 90%.
- In MoKA, a nontrivial mixture of medium and small proposal distributions is kept throughout the simulation (see Figure 2a). However, we note that the larger ones are used more often at the beginning of the simulation, which demonstrates the ability of the algorithm to automatically explore the space. The final acceptance rate is about 80%.
- In MoKA KIDS, the behaviour is similar to MoKA except that the badly-located particles are very quickly attracted by better-located ones so that the initial phase is much faster. The final acceptance rate is about 70%.

5.3. Gaussian mixtures

5.3.1. Target distribution

As a first challenging target distribution, we revisit the classical test case introduced in [17] and consider a mixture of two Gaussian distributions in dimension $d = 12$:

$$\pi = w_1 \mathcal{N}(m + \mu_d, \sigma_d^2 I_d) + w_2 \mathcal{N}(m - \mu_d, \sigma_d^2 I_d),$$

where $m = (0.5, \dots, 0.5)^T$ is the center of the domain $[0, 1]^d$. The target distribution is parametrized $\mu_d \in \mathbb{R}^d$, $\sigma_d > 0$ and the weights $w_1, w_2 > 0$. We will




Fig 2: Weights of the different proposal distribution in MoKA (a) and MoKA Markov (b). MoKA KIDS is not shown but is similar to MoKA.

always consider the standard deviation

$$\sigma_d = \frac{1}{2} \sqrt{\frac{0.4}{d}}.$$

We will then consider two cases.

1. The example considered in the seminal article [17] is given by

$$w_1 = w_2 = \frac{1}{2}, \quad \mu_d = \frac{1}{4\sqrt{d}}(-1, 1, \dots, 1)^T. \quad (26)$$

2. A slightly more difficult example is given by

$$w_1 = 0.25, \quad w_2 = 0.75, \quad \mu_d = \frac{1}{8}(-1, 1, \dots, 1)^T. \quad (27)$$

The second example is more difficult because the distance between the two modes is larger and the two modes are not equally weighted.

5.3.2. Initialization and parameters

All the algorithms are initialized according to the difficult case (25) where all the particles initially lie in a corner of the domain. Note also that the particles are initially closer to the first mode.

The kernel size in CMC is taken equal to 0.25. For the three MoKA variants, the proposal distribution is made of a mixture of three kernels with sizes 0.25, 0.4 and 0.55. A detailed discussion on how to choose these parameters is postponed to Section 5.4. For PMH, SAIS and SMC, we have tested all the kernel sizes between 0.05 and 0.8 with step 0.05 and we show the results for the best ones. For the simpler example (26), the best kernel size for PMH is 0.25, for SAIS

it is 0.4 and for SMC it is 0.3. For the more difficult example (27), the best kernel size for PMH is 0.8, for SAIS it is 0.5 and for SMC it is 0.4. Regarding the other parameters, the ESS threshold in SMC is set to $0.75N$, there are 150 MH steps between two resampling steps and a linear tempering scheme with 25 steps is used from the initial distribution to the target distribution. For SAIS, the sequence of kernel bandwidths and mixture weights are taken according to the recommendations of the authors (see Appendix C.2). We take $N = 10^5$ particles in CMC and its variants, PMH and SMC. The number of particles added to the system at each iteration in SAIS is such that the system has N particles at the end of the burn-in phase.

5.3.3. Results

For each method in competition, we have launched 50 independent simulations. For each simulation, we have allowed a fixed GPU computation time of 60 seconds (we recall that all the methods are parallelized and thus comparable in terms of computation time). The best results are shown for the different methods in competition in dimension $d = 12$ in Figure 3 below.

For the simpler example (26), all the methods achieve excellent results, including PMH even though it is about three times slower than CMC. The adaptive methods MoKA Markov and MoKA KIDS also perform well but they suffer from a higher computational cost and they are thus slower than CMC on this example. MoKA has a computational cost comparable to the one of CMC and thus achieves equivalently good results. Note also that SMC is disadvantaged compared to the other methods in terms of computation time since due to the tempering scheme, this method effectively samples π only during the few last

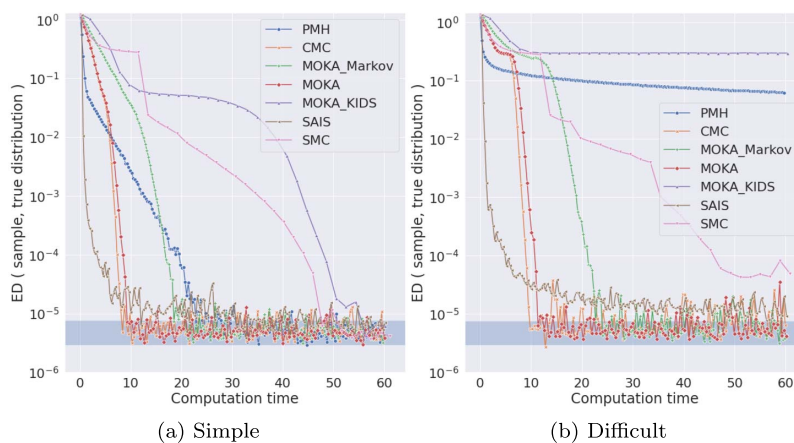


Fig 3: The best runs among 50 independent runs for the different algorithms in the simple case (a) and the difficult case (b). The blue shaded region is the 90% prediction interval of the energy distance between two true samples.

iterations. SAIS is initially the fastest method but the convergence then slows down since the computational cost increases with the iterations. Moreover, on 50 independent experiments, SAIS failed to converge in about 30% cases due to an early degeneracy of the particle weights which causes an over-concentration on only one mode. This is in contrast with all the other methods for which we have observed a very low variance in the results.

For the more difficult example, PMH fails to converge for any of the kernel sizes that we have tried. This situation can be expected because on the one hand, if the mutation kernel is small then the particle can never leave the first mode that they find and on the other hand if the mutation kernel is too large then the proposals are almost never accepted. This flaw is corrected in CMC and its MoKA variants by incorporating interactions between the particles and an adaptive choice of kernels. In particular, CMC, MoKA and MoKA Markov do not seem affected by the increased difficulty. Only MoKA KIDS fails to converge due to an early degeneracy of the particle weights which prevents them to leave the first mode. SAIS also has the same flaw as in the previous case. The performance of SMC are slightly deteriorated compared to the previous case but still achieves good results in most cases.

5.3.4. Comparison of the collective methods

In this section we compare the behaviours of the collective algorithms CMC, MoKA, MoKA Markov and MoKA KIDS. In the results below we show the convergence in terms of number of iterations rather than in terms of computation time. This is to illustrate that the collective proposals are designed to achieve better efficiency at each iteration (at the cost of a possibly increased computation time). We include PMH in the comparison as a control system. In Figure 4, we show the results for the two examples (26) and (27) during 100 iterations and for 10 independent runs.

A first observation is that, except for MoKA KIDS, the different methods typically show a very small variance in the results. Thanks to the propagation of chaos result and with a very large number of particles, the system is indeed expected to evolve according to the deterministic recurrence relation (9). It should also be noted that in all cases, the MoKA algorithms (except MoKA KIDS in the difficult case) need less iterations to converge. This shows that choosing adaptive kernels indeed increases the efficiency per iteration. Of course, it comes at the price of an increased computation cost. In this specific example and considering the results presented in Section 5.3.3, this increased efficiency does not appear necessary but in the following Section 5.5, we will investigate a situation where the MoKA algorithms perform better than CMC also in terms of computation time. Still regarding the efficiency, the collective methods including CMC are clearly better than PMH (in the simpler example, PMH needs about 2000 iterations to converge). This can be seen by comparing the acceptance rate in Figure 5. In particular, all collective methods have an acceptance rate above 0.4 while the acceptance rate for PMH drops to almost zero in the difficult case.

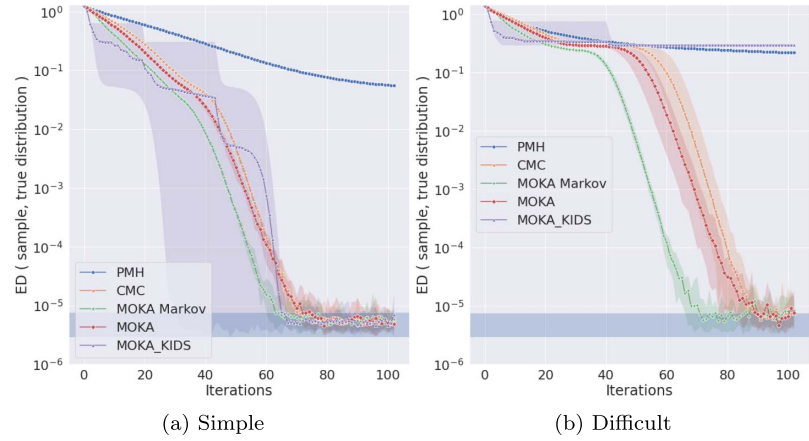


Fig 4: Mean error for 100 iterations of the Collective algorithms and PMH and min-max envelope for 10 independent runs (100% prediction interval). The horizontal blue shaded region is the 90% prediction interval of the energy distance between two true samples.

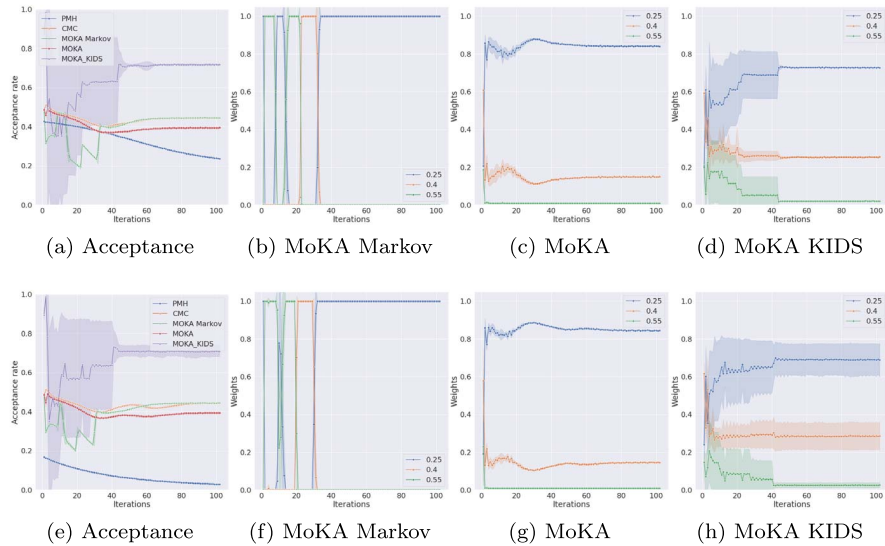


Fig 5: Mean acceptance rate and weights of the kernels over 100 iterations and standard deviations over 10 independent runs for the simple example ((a) to (d)) and for the difficult example ((e) to (h)).

In Figure 5, we can observe that the choice of the kernel weights is different in MoKA Markov compared to MoKA and MoKA KIDS. In the former case, the algorithm first chooses almost exclusively a large exploration proposal. Then, as

the particles get closer to the target distribution, the proposal kernel becomes smaller. For the other methods, a nontrivial mixture of small and medium kernels is kept during the whole simulation. This behaviour is similar to the one observed on the banana-shaped target distribution in Section 5.2.

Finally, the convergence plot (Figure 4) of the collective methods are roughly piecewise linear indicating an exponential convergence. More precisely, for CMC, MoKA and MoKA Markov we can identify two straight lines and a plateau in the difficult case. The first phase (up to iteration 30~40) corresponds to the colonization of the first mode. The following plateauing behaviour in the difficult case corresponds to a phase where the particles are almost exclusively distributed in the first mode. But as more particles enter the second mode, the Energy Distance to the target distribution decreases and a last exponentially fast converging phase begins. The duration of this plateauing exploration phase can vary, this is why a larger variance is observed for the last phase in the difficult example (27) compared to the simpler case (26). MoKA KIDS is less stable but for the simple example, it sometimes needs three times less iterations to converge compared to the other methods. Moreover, unlike the other methods, on its convergence plot (Figure 4), we can observe several plateauing phases interspersed with sudden drops. These drops corresponds to the times where a well-located particle gains most of the weight and suddenly attracts many other particles around itself. Similarly to the importance sampling based methods, this behaviour can be very efficient but also dangerous as it may impair the exploration of the space. This leads to a clear failure in the difficult example.

5.4. Tuning the hyper-parameters

As any of the other sampling methods discussed, the collective algorithms do require some appropriate tuning of the hyper-parameters. The simplest Algorithm 2 has two hyper-parameters, the number of particles and the size R of the interaction kernel K , which we take equal to the indicator of the ball centered at zero with radius R . As for the number of particles, it should simply be taken as large as possible, up to the computational capability. Note that with our implementation, one can easily handle samples of size $N \sim 10^5$ with a computation time of just a few seconds in the previous examples. The size of the interaction kernel is the most important parameter to tune. Although the convergence of the associated nonlinear Markov chain is ensured regardless of the size of K , since the algorithm uses a mean-field approximation with a finite number of particles, it is important to make sure that the particles interact in a mean-field regime. A simple criteria is simply to count the average number of neighbouring particles within the interaction radius around the proposals $(Y_t^i)_i$. This can be computed at no additional cost as it is simply a multiple of $K \star \hat{\mu}_t^N(Y_t^i)$. As a rule of thumb, particles need to interact with at least a few tens of neighbours. In our observations, the better results are obtained with $100 \sim 200$ neighbours in average and is not improved with more neighbours. In particular increasing the interaction kernel to a very large size will naturally increase the number




Fig 6: CMC for several sizes of the interaction kernel. Mean energy distance over 10 independent runs and min-max envelope

TABLE 1

Mean error, mean number of neighbors and mean acceptance rate computed over the ten last iterations of ten independent runs for each size of the interaction kernel.

Size of the interaction kernel	0.1	0.15	0.2	0.25	0.3	0.4	0.5
Mean error	7e-02	4.1e-04	1.9e-05	5.4e-06	7.3e-06	4e-04	2.5e-03
Mean number of neighbors	2	6	36	194	637	2137	3460
Mean acceptance rate	0.72	0.63	0.57	0.44	0.30	0.11	0.05

of neighbours but it may also deteriorate the performances as the acceptance rate will drop (which means that more iterations will be needed to converge). Both the number of neighbours and the acceptance rate can be easily monitored to ensure good convergence properties. This is illustrated in Figure 6 on the previous 12-dimensional simple example (26). Note that this only parameter to tune is in contrast with the other methods, in particular SAIS and SMC where the performances can also be strongly affected by the tempering scheme, the mixing weights sequence, the choice of the threshold value of the ESS, the number of mutation steps, the choice of the mutation kernel etc. All these hyper-parameters are subtle to tune in practice. Finally we also mention that although we did not explore this direction and leave it for future work, one may also consider adaptive versions of the Algorithm 2 (and its variants) where the size of the kernel is chosen in order to ensure a good balance between a good acceptance rate and a reasonable average number of neighbours. However, the convergence properties would likely be theoretically difficult to analyse as the number of neighbours is, by definition, not a mean-field quantity and the limit $N \rightarrow +\infty$ would therefore not make sense.

Finally, let us note that it is always safe to choose a large interaction kernel: in the maximal case where the kernel embraces all the domain, then the Vanilla CMC Algorithm 2 simply reduces to a standard random-walk MH algorithm with a proposal uniformly sampled in the domain. Although it may converge too slowly to be of practical use (it keeps a quadratic complexity in N), it is theoretically ensured to converge. On the other hand choosing a too small interaction kernel may lead to bad behaviours. In particular, we have observed that when a particle interact with only a few neighbours, approximately less than $10 \sim 20$, then the system may locally over-concentrate (i.e. there are localized high-density regions) and fail to converge towards the target distribution. This can be seen on Figure 6 where the smallest kernels lead to a plateauing behaviour far from the minimal error. This bad behaviour can nevertheless be easily monitored: in practice and in particular for the adaptive MoKA samplers, it is recommended to pay attention to the average number of neighbours and to either eliminate the interaction kernels which lead to a too small number of neighbours or to increase the number of particles. Since the volume of a ball of radius r behaves as r^d in dimension d , an unreasonably large number of particles or size of interaction kernel would be needed in very high-dimensional setting. This is why our collective algorithms cannot be used when d is too large. In our examples, we have observed good results up to dimension 13 (with 10^5 particles).

5.5. More complex example

5.5.1. Target distribution

In the classical test case introduced in [17, 25] and considered in Section 5.3, the target distribution is made of only two equally weighted Gaussians regardless of the dimension. Even for moderately large dimensions (around $d = 12$), good results can be obtained even with the simplest Metropolis-Hastings algorithm (with a very long computation time), although the symmetry of the target and the initialisation can be misleading. In this section we revisit this example and introduce the following more complex test case in dimension d :

$$\pi = \sum_{i=1}^d w_i \mathcal{N}(m + \mu_i, \sigma_d^2 I_d) + \sum_{j=1}^d \bar{w}_j \mathcal{N}(m - \mu_j, \sigma_d^2 I_d), \quad (28)$$

where $m = (0.5, \dots, 0.5)^T$ is the center of the domain, $\mu_i = \frac{a}{2} e_i$ with $a > 0$ and e_i is the i -th vector of the canonical basis, $\sigma_d > 0$ is a fixed standard deviation and $w_i, \bar{w}_i \in [0, 1]$ are normalized weights. This distribution is made of $2d$ Gaussian distributions, 2 per dimension, and as in [17, 25], the distance between the centers does not depend on the dimension (it is equal to a between the two distributions on the same dimension and to $a/\sqrt{2}$ between two distributions on two different dimensions). We choose the following parameters:

$$a = 0.7, \quad \sigma_d = \sqrt{\frac{0.03}{4d}}, \quad w_i = 0.25/d, \quad \bar{w}_i = 0.75/d.$$

With this choice, the $2d$ peaks are well separated and their variance decreases with the dimension. In addition half of the peaks carries the three quarters of the total weight of the distribution. This test case thus combines most of the difficulties: this is a multimodal distribution, there is a strong exploration-exploitation dilemma due to the size and the isolation of the peaks and all the modes do not have the same weight.

5.5.2. Initialization

Our samplers are initialized with all the particles in a corner (25), thus far away from all the peaks and closer to the modes with a lighter weight. For the competitors, this initialization leads to bad results and often a degeneracy of the importance weights. The behaviour of PMH does not really depend on the initialization since the large proposal distribution allows move all across the domain. Thus for PMH, SAIS and SMC the particles are initially sampled uniformly in the whole domain. This gives more chance to the importance sampling based methods since there are particles close to all the modes initially. Moreover, SMC may be quite affected by a bad initialisation due to the tempering scheme which forces the particles to stay close to their initial state during the first iterations. We have also checked that, for our samplers, the initialization (25) or a uniform initialization in the whole domain both lead to the same behaviour (with a faster convergence in the latter case).

5.5.3. Parameters

In order to tune the proposal distributions of our samplers, we introduce three different natural scales:

$$\sigma_{\text{small}} = 3.5\sigma_d, \quad \sigma_{\text{medium}} = 10\sigma_d, \quad \sigma_{\text{large}} = 20\sigma_d.$$

The small scale corresponds to roughly speaking, the size of a peak but it is very small compared to the distance between the peaks. The medium scale is larger than the typical size of a peak and allows moves between neighbouring peaks. The large scale distribution allows moves between far away peaks. The interaction kernel K in the Vanilla CMC Algorithm 2 is taken equal to the indicator of a ball of radius σ_{medium} . The different variants of the adaptive MoKA algorithm use the three sizes. For SAIS and SMC, we have tested all the kernel sizes between 0.05 and 0.8 with a step of 0.05 and we took the sizes which led to the best results. For PMH the best outcomes are obtained with a large kernel size although the different sizes essentially lead to equally bad results.

Regarding the other parameters, the ESS threshold in SMC is set to $0.75N$, there are 150 MH steps between two resampling steps and a linear tempering scheme with 80 steps is used from the initial distribution to the target distribution. For SAIS, the sequence of kernel bandwidths and mixture weights are taken according to the recommendations of the authors (see Appendix C.2).

In all samplers, the number of particles is $N = 10^5$ and the outcomes are shown for a fixed allowed GPU computation time of 180 seconds.

5.5.4. Results

The results shown in Table 2 are for dimensions ranging from $d = 4$ to $d = 13$ and computed over 50 independent runs for all samplers and each dimension. For better clarity, we have classified the results in four possible outcomes, similarly to [17], and defined as follows.

- **Disastrous.** The median Energy Distance between the final sample and a true sample is larger than $E_0/10$ where E_0 is the initial Energy Distance between a uniform sample in the domain and a true sample.
- **Excellent.** The median Energy Distance between the final sample and a true sample falls below the upper limit of the 90% prediction interval between two true samples.
- **Good.** The median Energy Distance between the final sample and a true sample falls between the upper limit of the 90% prediction interval and the average in log scale between this upper limit and $E_0/10$.
- **Mediocre.** The median Energy Distance between the final sample and a true sample falls between $E_0/10$ and the average in log scale between the upper limit of the 90% prediction interval and $E_0/10$.

The convergence plots of the best runs in dimensions 7 and 12 are shown in Figure 7.

5.5.5. Discussion

For small dimensions up to dimension 6, all samplers (including PMH) have a similar behaviour and manage to reach the smallest possible distance or at least have a good outcome. From dimension 6 for PMH, dimension 7 for CMC and MoKA KIDS and dimension 9 for SAIS, the outcome is not satisfactory anymore and even become disastrous for MoKA KIDS and SAIS. The only algorithms which have an excellent outcome up to dimension 10 are MoKA Markov and MoKA. The former remains an excellent sampler up to dimension 13. SMC also manages to have a good outcome up to dimension 13 although it sometimes, rarely, leads to mediocre or disastrous results. It is probable that with a better tempering scheme or an adaptive version, SMC would be able to reach the smallest possible distance although this may require a fine tuning of these hyper parameters. In comparison, MoKA Markov seems simpler to tune: the choice of the kernel sizes that we propose consistently leads to excellent results up to dimension 13. Without this adaptive behaviour, the fixed kernel size in CMC is initially well-adapted to exploration but it is too large compared to the size of the modes which causes a very slow convergence and leads to mediocre results (the final median error remains significantly better than PMH for the same computation time though). The adaptive procedure in MoKA is based on maximizing the acceptance and seems less efficient than the optimization procedure in MoKA Markov. It can be explained by the fact that the non-Markovian choice of the kernels tends to favour the smallest interaction kernel. Although

TABLE 2

Detailed outcomes of the different algorithms for the target distribution (28) with the dimension ranging from 4 to 13. The number in each cell corresponds to the median computed over 50 independent runs of the Energy Distance between the final sample and a true sample. The outcome is classified as Excellent (E) (green), Good (G) (light green), Mediocre (M) (yellow) or Disastrous (D) (red) as defined in Section 5.5.4. The four percentages at the bottom of each cell indicates, from left to right, the percentages of Excellent, Good, Mediocre and Disastrous results for 50 independent runs.

	PMH	CMC	MoKA Markov	MoKA	MoKA KIDS	SAIS	SMC
$d = 4$	E 4.4E-06 (98,2,0,0)	E 4.4E-06 (100,0,0,0)	E 4.6E-06 (100,0,0,0)	E 4.7E-06 (100,0,0,0)	E 4.4E-06 (100,0,0,0)	E 4.8E-06 (100,0,0,0)	E 4.5E-06 (100,0,0,0)
$d = 5$	G 3.3E-05 (0,100,0,0)	G 2.0E-05 (0,100,0,0)	E 4.6E-06 (100,0,0,0)	E 4.5E-06 (100,0,0,0)	E 4.9E-06 (96,2,2,0)	E 5.0E-06 (100,0,0,0)	E 5.3E-06 (94,6,0,0)
$d = 6$	M 5.9E-04 (0,0,100,0)	G 1.1E-04 (0,100,0,0)	E 4.7E-06 (100,0,0,0)	E 4.7E-06 (100,0,0,0)	G 4.1E-05 (16,46,32,6)	E 5.4E-06 (96,4,0,0)	G 1.0E-05 (24,70,6,0)
$d = 7$	M 1.8E-03 (0,0,100,0)	M 3.1E-04 (0,4,96,0)	E 4.9E-06 (100,0,0,0)	E 4.8E-06 (100,0,0,0)	D 3.7E-02 (0,0,0,100)	E 6.5E-06 (70,30,0,0)	G 1.5E-05 (14,80,2,4)
$d = 8$	M 2.5E-03 (0,0,100,0)	M 5.1E-04 (0,0,100,0)	E 4.8E-06 (100,0,0,0)	E 5.0E-06 (100,0,0,0)	D 9.4E-02 (0,0,0,100)	M 1.0E-03 (0,12,82,6)	G 6.5E-05 (0,96,2,2)
$d = 9$	M 3.2E-03 (0,0,100,0)	M 7.0E-04 (0,0,100,0)	E 4.8E-06 (100,0,0,0)	E 4.6E-06 (100,0,0,0)	D 1.9E-01 (0,0,0,100)	D 1.6E-02 (0,0,20,80)	G 2.8E-05 (0,92,2,6)
$d = 10$	M 3.5E-03 (0,0,100,0)	M 9.2E-04 (0,0,100,0)	E 4.9E-06 (100,0,0,0)	E 5.9E-06 (54,0,46,0)	D 2.1E-01 (0,0,0,100)	D 8.8E-02 (0,0,0,100)	G 5.4E-05 (0,100,0,0)
$d = 11$	M 3.7E-03 (0,0,100,0)	M 1.1E-03 (0,0,100,0)	E 5.1E-06 (100,0,0,0)	M 4.7E-03 (2,0,90,8)	D 2.6E-01 (0,0,0,100)	D 1.1E-01 (0,0,0,100)	G 6.8E-05 (0,96,0,4)
$d = 12$	M 3.8E-03 (0,0,100,0)	M 1.3E-03 (0,0,100,0)	E 5.1E-06 (100,0,0,0)	M 1.0E-02 (0,0,64,36)	D 2.9E-01 (0,0,0,100)	D 1.5E-01 (0,0,0,100)	G 1.2E-04 (0,100,0,0)
$d = 13$	M 3.9E-03 (0,0,100,0)	M 1.5E-03 (0,0,100,0)	E 5.3E-06 (98,0,2,0)	D 1.3E-02 (0,0,40,60)	D 5.0E-01 (0,0,0,100)	D 2.1E-01 (0,0,0,100)	G 1.1E-04 (0,98,2,0)

this kernel is also ultimately chosen by MoKA Markov, choosing this kernel “too early” does not introduce enough exploration. Note that with a uniform initialisation over the whole domain, we have observed (not shown) that this flaw disappears and MoKA and MoKA Markov have a similar behaviour. MoKA KIDS and SAIS both lead to disastrous results when the dimension increases. It is due to the early degeneracy of the weights of a small number of particles which prevents a proper mixing of the population.

6. Conclusion

Nonlinear MCMC samplers are appealing. They generalise more traditional methods and overcome many of their flaws. Getting back to the historical

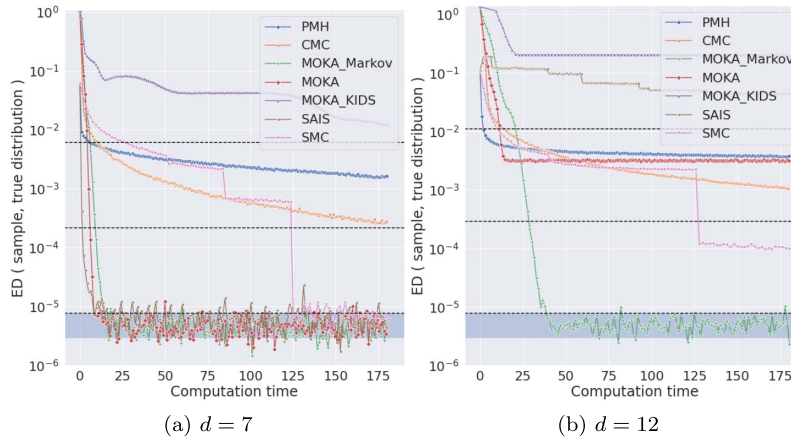


Fig 7: The best runs among 50 independent runs for the different algorithms in dimension $d = 7$ and $d = 12$. The three black dashed lines divide the y -axis into four regions for the four possible outcomes, from top to bottom: disastrous, mediocre, good, excellent. The best run is the one which leads to the smallest final energy distance or, if several runs of the same algorithm lead to an excellent outcome, the best run is the fastest one. The blue shaded region is the 90% prediction interval of the energy distance between two true samples.

development of mathematical kinetic theory, we can advantageously simulate nonlinear Markov processes using systems of interacting particles. This versatility enables the development of a wide variety of algorithms that can tackle difficult sampling problems while remaining in a traditional Markovian framework. Although the implementation may, at first sight, seem computationally demanding, we have shown that modern GPU hardware can now enable the use of interacting particles for Monte Carlo sampling at scale.

Alongside its variants, the CMC algorithm can be implemented efficiently and leads to striking reductions in global convergence times. It relies on pairwise interactions to best leverage the information that is present in any given sample swarm, and thus make the most of each evaluation of the target distribution. CMC avoids the mixing issues of classical “one particle” methods such as Metropolis-Hastings, with a notable improvement of the convergence and mixing speed. In particular when dealing with multimodal distributions, where the relative weight of each mode is difficult to estimate. In practice, we thus expect that the benefits of this improved “sample efficiency” will outweigh the (small) computational overhead of our method for most applications.

We note that the present contribution shares similarities with some well-known and recent nonlinear samplers, that are often based on non-Markovian importance sampling techniques. In the future, the joint development of Markovian and non-Markovian methods is likely to benefit both approaches: we may for instance improve the importance weights in SAIS-like methods as in the

KIDS algorithm, or construct better proposal distributions in CMC which incorporate knowledge of (part of) the past. The theoretical study of such hybrid methods would however be challenging and require the development of new analytical tools.

Finally, one may think of extending the theoretical framework introduced here to other MCMC samplers, such as *nonlinear* PDMP samplers or *nonlinear* Langevin dynamics. This could open new problems in nonlinear analysis and statistics, both on the theoretical and computational sides.

Appendix A: Proof of Theorem 2.11

Let us start with the following lemma.

Lemma A.1. *Let Θ be a proposal distribution which satisfies Assumptions 2.1, 2.2 and 2.3. Then the map*

$$\alpha : \mathcal{P}(E) \times E^2 \rightarrow \mathbb{R}, (\mu, x, y) \mapsto \alpha_\mu(x, y) := \frac{\Theta_\mu(x|y)\pi(y)}{\Theta_\mu(y|x)\pi(x)},$$

is Lipschitz in the Wasserstein-1 distance, in the sense that there exists a constant $L_\Theta > 0$ (which depends also on π) such that for all $\mu, \nu \in \mathcal{P}(E)$ and $x, x', y, y' \in E$, it holds that:

$$|\alpha_\mu(x, y) - \alpha_\nu(x', y')| \leq L_\Theta (W^1(\mu, \nu) + |x - x'| + |y - y'|).$$

Proof. By the triangle inequality, it holds that:

$$|\alpha_\mu(x, y) - \alpha_\nu(x', y')| \leq \frac{\pi(y)}{\pi(x)} \left| \frac{\Theta_\mu(x|y)}{\Theta_\mu(y|x)} - \frac{\Theta_\nu(x'|y')}{\Theta_\nu(y'|x')} \right| + \frac{\Theta_\nu(x'|y')}{\Theta_\nu(y'|x')} \left| \frac{\pi(y)}{\pi(x)} - \frac{\pi(y')}{\pi(x')} \right|.$$

We bound each of the two terms on the right-hand side:

$$\begin{aligned} \frac{\pi(y)}{\pi(x)} \left| \frac{\Theta_\mu(x|y)}{\Theta_\mu(y|x)} - \frac{\Theta_\nu(x'|y')}{\Theta_\nu(y'|x')} \right| &\leq \frac{\pi(y)}{\pi(x)\Theta_\mu(y|x)} |\Theta_\mu(x|y) - \Theta_\nu(x'|y')| \\ &\quad + \frac{\pi(y)\Theta_\nu(x'|y')}{\pi(x)\Theta_\mu(x|y)\Theta_\mu(x'|y')} |\Theta_\mu(y|x) - \Theta_\nu(y'|x')| \\ &\leq \frac{LM_0}{m_0\kappa_-} \left(1 + \frac{\kappa_+}{\kappa_-} \right) (W^1(\mu, \nu) + |x - x'| + |y - y'|). \end{aligned}$$

and

$$\begin{aligned} \frac{\Theta_\nu(x'|y')}{\Theta_\nu(y'|x')} \left| \frac{\pi(y)}{\pi(x)} - \frac{\pi(y')}{\pi(x')} \right| &\leq \frac{\Theta_\nu(x'|y')}{\pi(x)\Theta_\nu(y'|x')} |\pi(y) - \pi(y')| \\ &\quad + \frac{\Theta_\nu(x'|y')\pi(y')}{\Theta_\nu(y'|x')\pi(x)\pi(x')} |\pi(x) - \pi(x')| \\ &\leq \frac{\|\pi\|_{\text{Lip}}\kappa_+}{m_0\kappa_-} \left(|y - y'| + \frac{M_0}{m_0} |x - x'| \right), \end{aligned}$$

where $\|\pi\|_{\text{Lip}}$ denotes the Lipschitz norm of π . Gathering everything gives the result with

$$L_{\Theta} = \frac{M_0}{m_0} \left(\frac{L}{\kappa_-} \left(1 + \frac{\kappa_+}{\kappa_-} \right) + \frac{\|\pi\|_{\text{Lip}} \kappa_+}{m_0 \kappa_-} \right). \quad \square$$

Proof (of Theorem 2.11). The strategy of the proof of Theorem 2.12 will be based on coupling arguments inspired by [78] and adapted from [28]. We start by the following trajectorial representation of the nonlinear Markov chain $(\bar{X}_t)_t$ defined by the transition kernel (1).

Definition A.2 (Nonlinear process). Let $\bar{X}_0 \sim \mu_0$ be an initial state where $\mu_0 \in \mathcal{P}(E)$. The state \bar{X}_t at time $t \in \mathbb{N}$, $t \geq 1$, is constructed from \bar{X}_{t-1} and the law of \bar{X}_{t-1} denoted by $\mu_{t-1} \in \mathcal{P}(E)$ as follows.

1. Take a proposal a random variable $\bar{Y}_t \sim \Theta_{\mu_{t-1}}(\cdot|\bar{X}_{t-1})$
2. Compute the ratio

$$\alpha_{\mu_{t-1}}(\bar{X}_{t-1}, \bar{Y}_t) := \frac{\Theta_{\mu_{t-1}}(\bar{X}_{t-1}|\bar{Y}_t)\pi(\bar{Y}_t)}{\Theta_{\mu_{t-1}}(\bar{Y}_t|\bar{X}_{t-1})\pi(\bar{X}_{t-1})}.$$

3. Take $\bar{U}_t \sim \mathcal{U}([0, 1])$ and if $\bar{U}_t \leq h(\alpha_{\mu_{t-1}}(\bar{X}_{t-1}, \bar{Y}_t))$, then accept the proposal, else reject it:

$$\bar{X}_t = \bar{X}_{t-1} \mathbf{1}_{\bar{U}_t \geq h(\alpha_{\mu_{t-1}}(\bar{X}_{t-1}, \bar{Y}_t))} + \bar{Y}_t \mathbf{1}_{\bar{U}_t \leq h(\alpha_{\mu_{t-1}}(\bar{X}_{t-1}, \bar{Y}_t))}.$$

From now on we consider N independent copies $(\bar{X}_t^i)_t$, $i \in \{1, \dots, N\}$, of the nonlinear process defined by Definition A.2. We then construct a coupled particle process $(X_t^i)_t$ such that for all $i \in \{1, \dots, N\}$, initially $X_0^i = \bar{X}_0^i \sim \mu_0$ and for each time $t \in \mathbb{N}$ we take:

1. the same jump decision random variables $U_t^i = \bar{U}_t^i \sim \mathcal{U}([0, 1])$,
2. optimal proposals of the form $Y_t^i = s(\bar{Y}_t^i)$ where s is an optimal transport map between $\Theta_{\mu_{t-1}}(\cdot|\bar{X}_{t-1}^i)$ and $\Theta_{\hat{\mu}_{t-1}^N}(\cdot|X_{t-1}^i)$. Since these two probability measures are absolutely continuous with respect to the Lebesgue measure, the existence of such optimal transport map (Monge problem) is proved for instance in [20] or [16]. By definition, the pathwise error between the proposals can thus be controlled by

$$\begin{aligned} \mathbb{E}[|Y_t^i - \bar{Y}_t^i| | \mathcal{F}_{t-1}] &= W^1 \left(\Theta_{\hat{\mu}_{t-1}^N}(\cdot|X_{t-1}^i), \Theta_{\mu_{t-1}}(\cdot|\bar{X}_{t-1}^i) \right) \\ &\leq W^1 \left(\Theta_{\hat{\mu}_{t-1}^N}(\cdot|X_{t-1}^i), \Theta_{\bar{\mu}_{t-1}^N}(\cdot|\bar{X}_{t-1}^i) \right) \\ &\quad + W^1 \left(\Theta_{\bar{\mu}_{t-1}^N}(\cdot|\bar{X}_{t-1}^i), \Theta_{\mu_{t-1}}(\cdot|\bar{X}_{t-1}^i) \right) \end{aligned}$$

where $\bar{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}_t^i}$ and \mathcal{F}_t is the σ -algebra generated by the processes up to time $t \in \mathbb{N}$. We conclude that:

$$\mathbb{E}[|Y_t^i - \bar{Y}_t^i| | \mathcal{F}_{t-1}] \leq W^1(\hat{\mu}_{t-1}^N, \bar{\mu}_{t-1}^N) + |X_{t-1}^i - \bar{X}_{t-1}^i| + \epsilon_t^N \quad (29)$$

where the error term ϵ_t^N only depends on (the laws of) the N independent nonlinear processes $(\bar{X}_t^i)_t$:

$$\epsilon_t^N := W^1(\bar{\mu}_{t-1}^N, \mu_{t-1}).$$

Let $t \in \mathbb{N}$, $t \geq 1$. It holds that:

$$\begin{aligned} \mathbb{E}[|X_t^i - \bar{X}_t^i| | \mathcal{F}_{t-1}] &= \mathbb{E}[|Y_t^i - \bar{Y}_t^i| \mathbf{1}_{U_t^i \leq \min(h_t^i, \bar{h}_t^i)} | \mathcal{F}_{t-1}] \\ &\quad + |X_{t-1}^i - \bar{X}_{t-1}^i| \mathbb{P}(U_t^i \geq \max(h_t^i, \bar{h}_t^i) | \mathcal{F}_{t-1}) \\ &\quad + \mathbb{E}[|X_t^i - \bar{Y}_t^i| \mathbf{1}_{h_t^i \leq U_t^i \leq \bar{h}_t^i} | \mathcal{F}_{t-1}] \\ &\quad + \mathbb{E}[|Y_t^i - \bar{X}_t^i| \mathbf{1}_{\bar{h}_t^i \leq U_t^i \leq h_t^i} | \mathcal{F}_{t-1}] \end{aligned}$$

where we write for short:

$$h_t^i \equiv h(\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i)) \quad \text{and} \quad \bar{h}_t^i \equiv h(\alpha_{\mu_t}(\bar{X}_t^i, \bar{Y}_t^i)).$$

we deduce that:

$$\begin{aligned} \mathbb{E}[|X_t^i - \bar{X}_t^i| | \mathcal{F}_{t-1}] &\leq W^1(\hat{\mu}_{t-1}^N, \bar{\mu}_{t-1}^N) + 2|X_{t-1}^i - \bar{X}_{t-1}^i| \\ &\quad + 2M_0(\mathbb{P}(h_t^i \leq U_t^i \leq \bar{h}_t^i | \mathcal{F}_{t-1}) + \mathbb{P}(\bar{h}_t^i \leq U_t^i \leq h_t^i | \mathcal{F}_{t-1})) + \epsilon_t^N \end{aligned} \quad (30)$$

The last two probabilities are bounded by $\mathbb{E}[|h_t^i - \bar{h}_t^i| | \mathcal{F}_{t-1}]$. Assuming that h is L_h -Lipschitz for a constant $L_h > 0$, it holds that:

$$|h_t^i - \bar{h}_t^i| \leq L_h \left| \alpha_{\hat{\mu}_{t-1}}(X_{t-1}^i, Y_t^i) - \alpha_{\mu_{t-1}}(\bar{X}_{t-1}^i, \bar{Y}_t^i) \right|.$$

Let $\bar{\mu}_t^N$ be the empirical measure of the N nonlinear Markov processes \bar{X}_t^i at time t . It holds that:

$$\begin{aligned} |h_t^i - \bar{h}_t^i| &\leq L_h \left| \alpha_{\hat{\mu}_{t-1}}(X_{t-1}^i, Y_t^i) - \alpha_{\bar{\mu}_{t-1}}(\bar{X}_{t-1}^i, \bar{Y}_t^i) \right| \\ &\quad + L_h \left| \alpha_{\bar{\mu}_{t-1}}(\bar{X}_{t-1}^i, \bar{Y}_t^i) - \alpha_{\mu_{t-1}}(\bar{X}_{t-1}^i, \bar{Y}_t^i) \right| \end{aligned}$$

Using Lemma A.1 we get:

$$|h_t^i - \bar{h}_t^i| \leq C_\Theta \left(W^1(\hat{\mu}_{t-1}^N, \bar{\mu}_{t-1}^N) + |X_{t-1}^i - \bar{X}_{t-1}^i| + |Y_t^i - \bar{Y}_t^i| + \epsilon_t^N \right),$$

with $C_\Theta := L_h L_\Theta$. Therefore:

$$\mathbb{E}[|X_t^i - \bar{X}_t^i| | \mathcal{F}_{t-1}] \leq (1 + C_\Theta) (|X_{t-1}^i - \bar{X}_{t-1}^i| + W^1(\hat{\mu}_{t-1}^N, \bar{\mu}_{t-1}^N)) + (1 + 2C_\Theta) \epsilon_t^N.$$

Let us define:

$$S_t := \frac{1}{N} \sum_{i=1}^N |X_t^i - \bar{X}_t^i|.$$

Summing the previous expression for i from 1 to N and dividing by N gives the following estimate for S_t :

$$\mathbb{E}[S_t | \mathcal{F}_{t-1}] \leq 2(1 + C_\Theta) S_{t-1} + (1 + 2C_\Theta) \epsilon_t^N \tag{31}$$

where we have used the fact that

$$W^1(\hat{\mu}_{t-1}^N, \bar{\mu}_{t-1}^N) \leq S_{t-1}.$$

Taking the expectation in (31), we deduce that:

$$\mathbb{E}[S_t] \leq (1 + C_\Theta) \mathbb{E}[S_{t-1}] + C_\Theta \mathbb{E}[\epsilon_t^N] \tag{32}$$

where the value of the constant C_Θ has been updated by $C_\Theta \leftarrow 1 + 2C_\Theta$.

The error term can be controlled uniformly on t using [19, Theorem 5.8] or [38]. In particular, since π is a smooth probability density function on a compact set, it has finite moments of all orders and therefore it follows from [38, Theorem 1] that:

$$\forall t \in \mathbb{N}, \mathbb{E}[\epsilon_t^N] \leq \beta(N) \tag{33}$$

where $\beta(N)$ is defined by (5).

One can easily prove by induction that:

$$\mathbb{E}[S_t] \leq C_\Theta \beta(N) \sum_{s=0}^{t-1} e^{C_\Theta s} \leq \beta(N) \frac{C_\Theta}{e^{C_\Theta} - 1} e^{tC_\Theta} \tag{34}$$

By symmetry of the processes, all the quantities $\mathbb{E}[|X_t^i - \bar{X}_t^i|]$ are equal and their common value is $\mathbb{E}[S_t]$. The result follows. \square

Remark A.3 (Moderate interaction, part 2). The result of Theorem 2.11 provides an explicit convergence rate in terms of N . This could be used to understand more precisely the moderate interaction assumption mentioned in Remark 3.1 in order to justify the study of the degenerate proposal distribution $\Theta_\mu(dy|x) = \mu(dy)$ in the nonlinear Markov chain with transition kernel (1). In the case $\Theta_\mu(dy|x) = K \star \mu(y)dy$, one can take at the particle level (*i.e* in Algorithm 1) an interaction kernel $K \equiv K^N$ which depends on N . Typically, one can consider K^N equal to the density of a centered Gaussian law with covariance $\sigma_N I_d$, $\sigma_N > 0$ or the normalized indicator of the ball centered at zero and with radius σ_N . The goal is to take the size $\sigma_N \rightarrow 0$ as $N \rightarrow +\infty$ (and thus $K^N \rightarrow \delta_0$). As a consequence the constant $C_\Theta \equiv C_\Theta^N$ in Theorem 2.11 would depend on N . Since we have a precise control on $\beta(N)$ we can choose σ_N such that the following convergence still holds for all $t \leq T$ smaller than a fixed finite value $T \in (0, +\infty)$:

$$\beta(N) e^{tC_\Theta^N} \xrightarrow{N \rightarrow +\infty} 0.$$

In particular, σ_N should not converges to zero too fast, justifying the moderate interaction terminology introduced in [68]. In the limit $N \rightarrow +\infty$ we then obtain that the empirical measure $\hat{\mu}_t^N$ converges towards the t -th iterate of

the transition operator (3) with the degenerate choice of proposal distribution $\Theta_f(dy|x) = f(y)dy$ which makes sense as soon as $f \in \mathcal{P}_0^{\text{ac}}(E)$. We refer the reader to [53] and [28] for two examples of propagation of chaos results under a moderate interaction assumption. Note that this result is mainly of theoretical interest as it does not give sharp estimates on how slow σ_N should decrease to zero. Moreover, similarly to Theorem 2.9, this result is only valid when $N \rightarrow +\infty$ on a finite bounded time interval.

Remark A.4 (About the assumptions). In order to prove a propagation of chaos property, it is usually assumed that the parameters of the problem are Lipschitz [78, 63]. This corresponds to the two Lipschitz assumptions 2.2 and 2.3. Propagation of chaos in non-Lipschitz settings is a more difficult problem (see for instance [48] for a recent result).

Assumption 1.1 and Assumption 2.1 should be understood as technical assumptions. In the proof of Theorem 2.12, we use the fact that the acceptance ratio is Lipschitz (Lemma A.1) which follows directly from Assumptions 1.1, 2.1 and 2.2. However, we could relax the compactness assumption 1.1 and keep the same Lipschitz property by replacing Assumptions 1.1, 2.1 and 2.2 by the following assumption.

Assumption A.5. *The target distribution π does not vanish on E and the map*

$$\mathcal{P}(E) \times E^2 \rightarrow \mathbb{R}, \quad (\mu, x, y) \mapsto g_\mu(y|x) := \frac{\Theta_\mu(y|x)}{\pi(y)}$$

satisfies the two following properties.

- (**Boundedness**). *There exists two constants $\kappa_- > 0$ and $\kappa_+ > 0$ such that*

$$\forall (x, y) \in E^2, \quad \kappa_- \leq g(y|x) \leq \kappa_+.$$

- (**Lipschitz**). *There exists a constant $L > 0$ such that*

$$\begin{aligned} \forall (\mu, x, y), (\nu, x', y') \in \mathcal{P}(E) \times E^2, \\ |g_\mu(y|x) - g_\nu(y'|x')| \leq \left(W^1(\mu, \nu) + |x - x'| + |y - y'| \right). \end{aligned}$$

In practice, this would necessitate a precise control of the tails of π and of the proposal distribution. It seems easier for us to check the compactness and boundedness assumptions 1.1, 2.1 and 2.2 (possibly up to truncating the support of π and replacing it by a compact set).

Appendix B: Continuous time version

As mentioned in Section 2.4, entropy methods are widely used in a continuous-time context, in particular for the long-time analysis of nonlinear PDEs [54]. In this section, we define a continuous-time version of Algorithm 1, we show that the mean-field limit towards the solution of a nonlinear PDE and we prove the

exponential long-time convergence of its solution towards π . This last step also motivates the discrete-time analysis in Section 2.4. Finally, the application of these ideas to the classical linear Metropolis-Hastings case leads (formally) to a convergence result obtained earlier in [27] with other techniques (Section B.4).

B.1. Continuous-time particle system

The N -particle system constructed by Algorithm 1 is a standard Markov chain which can be turned into a continuous-time Markov process by subordinating it to a Poisson process, as explained in [15, Chapter 8, Definition 2.2]. Namely, let $(N_t)_{t \geq 0}$ be a Poisson process independent of all the other random variables and with constant parameter equal to 1. The continuous-time particle system is defined at time $t \geq 0$ by $\tilde{X}_t^i := X_{N_t}^i$ for $i \in \{1, \dots, N\}$, where $X_{N_t}^i$ is the i -th particle constructed by Algorithm 1 at iteration N_t . For notational simplicity, in the remaining of this section, we drop the tilde notation and simply write $X_t^i \equiv \tilde{X}_t^i$ with $t \in [0, +\infty)$.

The first step is the mean-field limit analog of Theorem 2.11.

Theorem B.1. *Let the particles $X_0^i \sim f_0$ be initially i.i.d. with common law $f_0 \in \mathcal{P}_0^{\text{ac}}(E)$. Under the assumptions of Theorem 2.11, the (random) continuous-time empirical measure $\hat{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}$ at any time $t \geq 0$ satisfies*

$$\mathbb{E}W^1(\hat{\mu}_t^N, f_t) \leq C_1 \beta(N) e^{tC_2},$$

where $C_1, C_2 > 0$ are two absolute constants, $\beta(N)$ is defined by (5) and f_t is the solution of the nonlinear PDE

$$\partial_t f_t = \mathcal{T}[f_t] - f_t, \quad (35)$$

with initial condition f_0 .

Proof. It is enough to prove the analog of (6) in the time continuous framework. The end of the proof will then follow similarly as in Corollary 2.12. The nonlinear process $(\bar{X}_t)_t$ with law f_t is defined as follows.

1. Let $(T_n)_{n \in \mathbb{N}}$ the increasing sequence of jump times of the Poisson process $(N_t)_t$ with $T_0 = 0$ and initially $\bar{X}_0 \sim f_0$.
2. Between two jump times, at any $t \in [T_n, T_{n+1})$, $\bar{X}_t = \bar{X}_{T_n}$.
3. At each jump time T_n , we sample a proposal random variable $\bar{Y}_{T_n} \sim \Theta_{f_{T_n}}(\cdot | \bar{X}_{T_n})$ and accept or reject it as in Definition A.2.

Then we construct two coupled systems of N independent nonlinear processes $(\bar{X}_t^i)_t$ and N particles as in the proof of Theorem 2.11. We define similarly $S_t := \frac{1}{N} \sum_{i=1}^N \mathbb{E}|X_t^i - \bar{X}_t^i|$. As in the discrete time case (31), at any jump time T_n , it holds that

$$\mathbb{E}[S_{T_n} | \mathcal{F}_{T_n}^-] \leq 2(1 + C_\Theta) S_{T_{n-1}} + C_\Theta \beta(N).$$

Taking the conditional expectation with respect to $\mathcal{G}_n = \sigma(T_1, T_2 - T_1, \dots, T_n - T_{n-1})$, we obtain:

$$\mathbb{E}[S_{T_n} | \mathcal{G}_n] \leq 2(1 + C_\Theta) \mathbb{E}[S_{T_{n-1}} | \mathcal{G}_{n-1}] + C_\Theta \beta(N).$$

And thus, for all $n \in \mathbb{N}$,

$$\mathbb{E}[S_{T_n} | \mathcal{G}_n] \leq \beta(N) \frac{C_\Theta}{e^{C_\Theta} - 1} e^{C_\Theta n}.$$

Since N_t follows a Poisson law with parameter t , it holds that:

$$\begin{aligned} \mathbb{E}[S_{T_n} \mathbf{1}_{N_t=n}] &= \mathbb{E}[\mathbb{E}[S_{T_n} \mathbf{1}_{N_t=n} | \mathcal{G}_{n+1}]] \\ &= \mathbb{E}[\mathbf{1}_{N_t=n} \mathbb{E}[S_{T_n} | \mathcal{G}_n]] \\ &\leq \beta(N) \frac{C_\Theta}{e^{C_\Theta} - 1} e^{C_\Theta n} \mathbb{P}(N_t = n), \end{aligned}$$

where the second inequality comes from the fact that the event $\{N_t = n\}$ is \mathcal{G}_{n+1} measurable and the fact that S_{T_n} is independent from $T_{n+1} - T_n$. As a consequence, since $S_t = S_{T_{N_t}}$ and $\mathbb{P}(N_t = n) = e^{-t} t^n / n!$, we conclude that:

$$\mathbb{E}[S_t] = \mathbb{E}[S_{T_{N_t}}] = \sum_{n=0}^{+\infty} \mathbb{E}[S_{T_n} \mathbf{1}_{N_t=n}] \leq \beta(N) \frac{C_\Theta}{e^{C_\Theta} - 1} \exp(t(e^{C_\Theta} - 1)). \quad \square$$

B.2. Convergence of the nonlinear process

The analog of Theorem 2.15 in the continuous-time setting is the following theorem.

Theorem B.2 (Convergence of the Nonlinear Process). *Let Θ a proposal distribution which satisfies Assumption 2.6. Let $f_0 \in \mathcal{P}_0^{\text{ac}}(E)$ and let f_t be the solution at time $t \in [0, +\infty)$ of the nonlinear PDE (35). Then for all $t \geq 0$, it holds that:*

$$\|f_t - \pi\|_{\text{TV}} \leq C_0 e^{-\lambda t},$$

where $C_0 > 0$ depends only on f_0 and π and where

$$\lambda := c^- \left(\inf_x \frac{f_0(x)}{\pi(x)} \right) h(1) > 0. \tag{36}$$

Note that in the continuous-time case, Assumption 2.4 is not required.

We recall that given a convex function $\phi : [0, +\infty) \rightarrow [0, +\infty)$ such that $\phi(1) = 0$, the relative entropy $\mathcal{H}[f|\pi]$ and dissipation $\mathcal{D}[f|\pi]$ of a probability density $f \in \mathcal{P}^{\text{ac}}(E)$ with respect to π are defined respectively by

$$\mathcal{H}[f|\pi] := \int_E \pi(x) \phi\left(\frac{f(x)}{\pi(x)}\right) dx, \quad \mathcal{D}[f|\pi] := - \int_E \phi'\left(\frac{f(x)}{\pi(x)}\right) (\mathcal{T}[f](x) - f(x)) dx.$$

In the continuous time framework when f_t solves (35), the entropy and dissipation are simply linked by the relation

$$\frac{d}{dt} \mathcal{H}[f_t|\pi] = -\mathcal{D}[f_t|\pi] \leq 0, \quad (37)$$

and therefore we conclude immediately that the entropy is non increasing (which is Lemma 2.18 in the discrete setting). A consequence of this fact is an alternative proof of Lemma 2.14.

Lemma B.3. *Let f_t be the solution of the integro-differential equation (35) with initial condition $f_0 \in \mathcal{P}_0^{\text{ac}}(E)$. Then*

$$\inf_{x \in E} \frac{f_t(x)}{\pi(x)} \geq \inf_{x \in E} \frac{f_0(x)}{\pi(x)}, \quad \sup_{x \in E} \frac{f_t(x)}{\pi(x)} \leq \sup_{x \in E} \frac{f_0(x)}{\pi(x)}.$$

Proof. Let us denote

$$m := \inf_{x \in E} \frac{f_0(x)}{\pi(x)} \quad \text{and} \quad M := \sup_{x \in E} \frac{f_0(x)}{\pi(x)}.$$

Let us take $\phi : [0, +\infty) \rightarrow [0, +\infty)$ a convex function such that $\phi \equiv 0$ on the segment $[m, M]$ and $\phi > 0$ elsewhere. Note that since f_0 and π are both probability densities, it holds that $m < 1$ and $M > 1$ and thus $\phi(1) = 0$. The entropy-dissipation relation (37) gives:

$$\frac{d}{dt} \mathcal{H}[f_t|\pi] \leq 0$$

and therefore for all $t \geq 0$,

$$\mathcal{H}[f_t|\pi] \leq \mathcal{H}[f_0|\pi] = 0$$

by definition of ϕ . As a consequence and since $\phi \geq 0$ and $\pi > 0$ on E , it holds that for all $t \geq 0$ and all $x \in E$,

$$\phi\left(\frac{f_t(x)}{\pi(x)}\right) = 0,$$

which implies that

$$\forall x \in E, \quad m \leq \frac{f_t(x)}{\pi(x)} \leq M. \quad \square$$

In order to prove Theorem B.2, we follow the classical steps which are detailed for instance in [54, Section 1.3] and can be applied to various linear and nonlinear jump and diffusion processes.

1. Compute the dissipation $\mathcal{D}[f_t|\pi] = -\frac{d}{dt} \mathcal{H}[f_t|\pi]$.
2. Prove that the dissipation can be bounded from below by a multiple of the entropy: for a constant $\lambda > 0$,

$$\mathcal{D}[f_t|\pi] \geq \lambda \mathcal{H}[f_t|\pi].$$

3. Apply Gronwall lemma to the relation $\frac{d}{dt}\mathcal{H}[f_t|\pi] \leq \lambda\mathcal{H}[f_t|\pi]$ to obtain the exponential decay of the entropy:

$$\mathcal{H}[f_t|\pi] \leq \mathcal{H}_0 e^{-\lambda t}.$$

4. Show that the entropy controls the TV distance and conclude that for some constants $c, C_0 > 0$:

$$\|f_t - \pi\|_{\text{TV}} \leq c\mathcal{H}[f_t|\pi] \leq C_0 e^{-\lambda t}.$$

Proof (of Theorem B.2). Let $\phi(s) = \frac{1}{2}(s-1)^2$. Most of the computations are the same as in the discrete-time case. From the entropy-dissipation relation (37), Lemma B.3 and Lemma 2.19, it follows that

$$\begin{aligned} \frac{d}{dt}\mathcal{H}[f_t|\pi] &= -\mathcal{D}[f_t|\pi] \\ &\leq -\frac{c^-(m)h(1)}{2} \iint_{E \times E} \pi(x)\pi(y) \left| \frac{f_t(x)}{\pi(x)} - \frac{f_t(y)}{\pi(y)} \right|^2 dx dy = -2c^-(m)h(1)\mathcal{H}[f_t|\pi], \end{aligned}$$

where $m := \inf_{x \in E} f_0(x)/\pi(x)$. Using Gronwall's inequality we then deduce that:

$$\mathcal{H}[f_t|\pi] \leq \mathcal{H}[f_0|\pi] e^{-2c^-(m)h(1)t}.$$

The conclusion follows from the Cauchy-Schwarz inequality by writing

$$\|f_t - \pi\|_{\text{TV}} = \int_E |f_t(x) - \pi(x)| dx = \int_E \sqrt{\pi} \sqrt{\pi} \left| \frac{f_t(x)}{\pi(x)} - 1 \right| dx \leq \sqrt{2\mathcal{H}[f_t|\pi]},$$

where we have used the fact that the TV norm is equal to the L^1 norm of the probability density functions. \square

Remark B.4. Another natural choice for ϕ would be $\phi(s) = s \log s - s + 1$. The relative entropy is in this case equal to the Kullback-Leibler divergence. However, the dissipation term becomes in this case:

$$\begin{aligned} \mathcal{D}[f_t|\pi] &= \frac{1}{2} \iint_{E \times E} W_{f_t}(x \rightarrow y) \pi(x) \left(\frac{f_t(x)}{\pi(x)} - \frac{f_t(y)}{\pi(y)} \right) \\ &\quad \times \left(\log \left(\frac{f_t(x)}{\pi(x)} \right) - \log \left(\frac{f_t(y)}{\pi(y)} \right) \right) dx dy, \end{aligned}$$

and it is not clear that it can be bounded from below by the relative entropy. Note that this dissipation functional is very similar to the one obtained in the study of the Boltzmann equation (in this context, the Kullback-Leibler divergence is also called the Boltzmann entropy). The long-time asymptotics of this equation is a long-standing problem and the specific question of whether the dissipation controls the entropy is the object of a famous conjecture by Cercignani [26, 84]. In our case, we know that the Kullback-Leibler divergence is decreasing with time but all this suggests that its exponential decay could be harder to obtain or could hold only in specific cases.

Putting together Theorem B.1 and Theorem B.2 leads to the following corollary which is the continuous-time analog of Theorem 2.9.

Corollary B.5. *Under the assumptions of Theorem B.1 and Theorem B.2, for any $t > 0$ and $N > 1$ it holds that*

$$\mathbb{E}[W^1(\hat{\mu}_t^N, \pi)] \leq C_1\beta(N)e^{C_2t} + C_3e^{-\lambda t},$$

where $\beta(N)$ is given by (5), $\lambda > 0$ is given by (36) and $C_1, C_2, C_3 > 0$ are absolute constants.

Proof. By the triangle inequality, it holds that

$$\mathbb{E}[W^1(\hat{\mu}_t^N, \pi)] \leq \mathbb{E}[W^1(\hat{\mu}_t^N, f_t)] + \mathbb{E}[W^1(f_t, \pi)].$$

The first term on the right-hand side is bounded by $C_1\beta(N)e^{C_2t}$ by Theorem B.1. For the second term on the right-hand side, we first note that on the compact set E , the total variation norm controls the Wasserstein-1 distance [85, Theorem 6.15]. The conclusion therefore follows from Theorem B.2. \square

B.3. Links between the discrete- and continuous-time versions

The discrete-time counterpart of the entropy-dissipation relation (37) is the relation (16). The difference $\mathcal{H}[\mu_{t+1}|\pi] - \mathcal{H}[\mu_t|\pi]$ is the discrete analog of a time derivation. The main difference with the continuous-time entropy-dissipation relation is the additional non negative term on the right-hand side of (16). The role of the technical Assumption 2.4 is to ensure that this non negative term remains smaller than the dissipation in order to close the argument as in the continuous-time case. Since this term does not appear in the continuous-time setting, Assumption 2.4 is not required to prove Theorem B.2.

Assumption 2.4 is actually better understood when the discrete-time relation (9) is seen as an explicit Euler discretization scheme (11) of the nonlinear PDE (35). More precisely, the numerical scheme (11) can be re-written

$$\mu_{t+1} = \mu_t + \Delta t Q[\mu_t], \tag{38}$$

where

$$Q[\mu](dx) := \int_E \pi(x) W_\mu(x \rightarrow y) \left(\frac{\mu(dy)}{\pi(y)} dx - \frac{\mu(dx)}{\pi(x)} dy \right).$$

Note that $\mathcal{T}[\mu] = \mu + Q[\mu]$. In order to check that this numerical scheme preserves the continuous-time entropy-dissipation relation (37), let us write for a general function the second-order Taylor expansion:

$$\mathcal{H}[\mu_{t+1}|\pi] \simeq \mathcal{H}[\mu_t|\pi] - \Delta t \mathcal{D}[\mu_t|\pi] + \frac{\Delta t^2}{2} \int_E \pi(x) \phi'' \left(\frac{\mu_t(x)}{\pi(x)} \right) \left| \frac{Q[\mu_t](x)}{\pi(x)} \right|^2 dx.$$

Since ϕ is a convex function, the second-order term of the Taylor expansion is non negative but it is dominated by the non positive first-order term (equal to

minus the dissipation) for Δt small enough. For this reason, the explicit Euler discretization scheme does not unconditionally preserve the entropy structure of the continuous-time PDE. Since $W_\mu(x \rightarrow y) = \Theta_\mu(y|x)h(\alpha_\mu(x, y))$, the operator $Q[\mu]$ is proportional to h and it is equivalent to assume $\Delta t < 1$ or $h < 1$. With the notations of Section 2.4.3, the time-step is $\Delta t \equiv \eta$ and Assumption 2.4 can therefore be interpreted as a “numerical” condition to preserve the entropy structure of the discrete scheme. We did not manage to prove the exponential convergence of the discrete scheme when $\Delta t \equiv \eta = 1$. However, it is still possible to prove the convergence of the discrete scheme without rate and for a weaker topology using Theorem B.2 and a compactness argument.

Corollary B.6 (Convergence of the discrete scheme). *Let $(\mu_t)_{t \in \mathbb{N}}$ be the sequence of probability laws defined by the recurrence relation (9) (i.e. by the discrete scheme (38) with $\Delta t = 1$). Let $\mu_0 \in \mathcal{P}_0^{\text{ac}}(E)$ be the initial condition and let Θ satisfy Assumption 2.6. Then it holds that $\mu_t \rightarrow \pi$ as $t \rightarrow +\infty$ for the weak convergence of probability measures.*

Proof. The sequence $(\mu_t)_{t \in \mathbb{N}}$ is tight because E is compact so it admits a converging subsequence. The limit of any converging subsequence is a fixed point of the operator \mathcal{T} . Since the convergence result stated in Theorem B.2 does not depend on the initial condition in $\mathcal{P}_0^{\text{ac}}(E)$, it implies that π is the unique fixed point of the operator \mathcal{T} in $\mathcal{P}_0^{\text{ac}}(E)$ and therefore, all the converging subsequences of $(\mu_t)_t$ converges towards π , which implies the convergence of the whole sequence. \square

B.4. The Metropolis-Hastings case

The following theorem revisits the main result of [27] regarding the convergence rate of the Metropolis-Hastings algorithm, here formally proved with the entropy techniques introduced in Section B.2 and in the continuous-time setting for simplicity.

Theorem B.7 (Formal). *Let us consider the linear case outlined in Section 3.1 with $q(y|x) = K_\sigma(x - y)$, where K_σ is a fixed symmetric random walk kernel of size $\sigma > 0$ (typically a Gaussian kernel with standard deviation σ). Assume that, as $\sigma \rightarrow 0$, the random-walk kernel K_σ satisfies for any smooth function $\varphi \in C_0^\infty(\mathbb{R}^d)$:*

$$\int_E \varphi(x) K_\sigma(x) dx = \varphi(0) + \frac{1}{2} \sigma^2 \Delta \varphi(0) + o(\sigma^2). \quad (39)$$

Assume that π and E are such that the following Poincaré inequality holds:

$$\int_E u(x)^2 \pi(x) dx - \left(\int_E u(x) \pi(x) dx \right)^2 \leq \frac{1}{\lambda_P} \int_E |\nabla u|^2(x) \pi(x) dx \quad (40)$$

for all functions u in the weighted Sobolev space $H_\pi^1(E)$ and for a constant $\lambda_P > 0$. Then, as $\sigma \rightarrow 0$, it holds that:

$$\|f_t - \pi\|_{\text{TV}} \leq C_0 e^{-\frac{1}{2} \sigma^2 (\lambda_P + o(1)) t},$$

for a constant $C_0 > 0$ which depends only on f_0 and π .

Proof. The formal Taylor expansion (39) as $\sigma \rightarrow 0$ applied to (37) leads to

$$\mathcal{D}_\phi[f_t|\pi] = \frac{\sigma^2}{2} \int_E \pi(x) \phi'' \left(\frac{f_t(x)}{\pi(x)} \right) \left| \nabla_x \left(\frac{f_t}{\pi} \right) \right|^2 dx + o(\sigma^2). \quad (41)$$

Taking $\phi(s) = \frac{1}{2}(s-1)^2$ gives:

$$\mathcal{D}_\phi[f_t|\pi] = \frac{\sigma^2}{2} \int_E \pi \left| \nabla \left(\frac{f_t}{\pi} \right) \right|^2 dx + o(\sigma^2).$$

Using the Poincaré inequality with the function $u = f_t/\pi$, we obtain:

$$\mathcal{D}_\phi[f_t|\pi] \geq \frac{\sigma^2 \lambda_P}{2} \left(\int_E \left(\frac{f_t}{\pi} \right)^2 \pi dx - 1 \right) + o(\sigma^2) = \sigma^2 \lambda_P \mathcal{H}_\phi[f_t|\pi] + o(\sigma^2).$$

From the entropy-dissipation relation (37) and Gronwall lemma, we deduce that

$$\mathcal{H}_\phi[f_t|\pi] \leq C_0 e^{-\sigma^2(\lambda_P + o(1))t}.$$

The conclusion follows from the Cauchy-Schwarz inequality as in the proof of Theorem B.2. \square

A similar result has been obtained rigorously in [27] using linear spectral theory. In particular, the fact that a Poincaré inequality holds depends on the regularity of the boundary of E . In our case, the argument could be made rigorous by studying in details the wellposedness of (35) in Sobolev spaces. The argument could also lead to a more detailed analysis of the convergence rate of the Metropolis-Hastings algorithm in other metrics. In particular, the same formal argument holds when we take $\phi(s) = s \log(s) - s + 1$ in (41). In this case, the relative entropy is the Kullback-Leibler divergence and its time derivative is controlled by the following dissipation:

$$\mathcal{D}_\phi[f_t|\pi] = \frac{\sigma^2}{2} \int_E \pi \left| \nabla \sqrt{\frac{f_t}{\pi}} \right|^2 dx + o(\sigma^2).$$

In order to apply Gronwall lemma and obtain the exponential decay of the Kullback-Leibler divergence, we need the following convex Sobolev inequality:

$$\int_E u(x)^2 \log \left(\frac{u(x)^2}{\|u\|_{L^2_\pi}^2} \right) \pi(x) dx \leq \frac{1}{\lambda_S} \int_E |\nabla u(x)|^2 \pi(x) dx,$$

for all $u \in H_\pi^1$ and for a constant $\lambda_S > 0$. The conclusion follows by applying this inequality to $u = \sqrt{f_t/\pi}$ and as before from Gronwall lemma and from the classical Csiszár-Kullback-Pinsker inequality [54, Theorem A.2] which shows that the Total Variation norm is controlled by the Kullback-Leibler divergence. When

the above convex Sobolev inequality holds, then the Poincaré inequality (40) also holds with $\lambda_P \geq 2\lambda_S$. More details on convex Sobolev inequalities can be found in [8] or [54, Section 2.2]. Their application to the rigorous computation of (optimal) convergence rates for the Metropolis-Hastings algorithm is left for future work.

Appendix C: Related works

C.1. Another nonlinear MCMC sampler

A nonlinear kernel which does not fit into the “collective proposal” category has been introduced in [4] and is defined by:

$$K_\mu(x, dy) = (1 - \varepsilon)K^{MH}(x, dy) + \varepsilon Q_\mu(x, dy),$$

where K^{MH} is the Metropolis-Hastings kernel and

$$Q_\mu(x, dy) = \left(1 - \int_E \alpha(x, u)\mu(du)\right) \delta_x(dy) + \alpha_\eta(x, y)\mu(dy).$$

The function α is defined by: $\alpha_\eta(x, y) = \eta(x)\pi(y)/(\eta(y)\pi(x))$, that is, $\alpha_\eta(x, y)$ is the Metropolis ratio associated to **another** distribution $\eta \in \mathcal{P}_0^{\text{ac}}(E)$. In [4], the authors investigated the case $\eta = \pi^{\tilde{\alpha}}$ for $\tilde{\alpha} \in (0, 1)$. This kernel satisfies:

$$\iint_{E \times E} \phi(y)K_\eta(x, dy)\pi(dx) = \int_E \phi(y)\pi(dy).$$

The sampling procedure is therefore quite different as it requires an auxiliary chain to build samples from η first in order to construct a sample from the desired nonlinear kernel. More precisely, the authors propose the following iterative procedure to construct a couple of Markov chains (X_t, Y_t) :

$$(X_{t+1}, Y_{t+1}) \sim \left((1 - \varepsilon)K^{MH}(X_t, dx_{t+1}) + \varepsilon Q_{\hat{\mu}_t^Y}(X_t, dx_{t+1})\right)P(Y_t, dy_{t+1}),$$

where P is a (linear) Markov transition kernel with invariant distribution η and $(Y_t)_t$ is a Markov chain with transition kernel P . The empirical measure of this chain is denoted by:

$$\hat{\mu}_t^Y = \frac{1}{t+1} \sum_{s=0}^t \delta_{Y_s}.$$

The final MCMC approximation of an observable φ is given in this case by:

$$\int_E \varphi(x)\pi(dx) \simeq \frac{1}{t+1} \sum_{s=0}^t \varphi(X_s). \quad (42)$$

In this empirical sum, the successive iterations of the single chain $(X_t)_t$ are used. In the collective proposal framework introduced in Section 2, the algorithm

produces N (asymptotically) independent copies of a nonlinear chain $(X_t^i)_t$, $i \in \{1, \dots, N\}$ and we have at our disposal a *sequence* of MCMC approximations of the form:

$$\int_E \varphi(x) \pi(dx) \simeq \frac{1}{N} \sum_{i=1}^N \varphi(X_t^i), \quad (43)$$

as $t \rightarrow +\infty$. We can therefore interpret the sum (42) as a *time average* and the sum (43) as an *ensemble average*.

C.2. Links with importance sampling based methods

Even though CMC does not use importance weights, it shares some similarities with importance sampling methods, in particular SMC [24] and PMC methods [18]. We will discuss the links with SMC in Section C.3 and focus here on PMC.

According to [18], in PMC methods, without the importance correction, a regular acceptance step — as in Metropolis-Hastings — for each mutation would lead to a simple parallel implementation of N Metropolis-Hastings algorithm. Under the same parallel, we can compare the mutation step in PMC with the proposal step of CMC.

In the first implementation of PMC, at each mutation step, each particle X_t^i is updated independently from the others, according to a kernel $q_{it}(\cdot)$ (that can depend on t and i), the new particle X_{t+1}^i is then associated with a weight proportional to the ratio $\pi(X_{t+1}^i)/q_{it}(X_{t+1}^i)$. In PMC, a mutation therefore occurs according to q_{it} , that is only depends on the position of the ancestor particle. In CMC, the mutation occurs according to $\Theta_{\hat{\mu}_t^N}(\cdot | X_t^i)$, the update thus depends on the position of all the particles as Θ depends on the empirical measure of the system $\hat{\mu}_t^N$. This additional dependency is particularly emphasized in the case of Algorithm 2. This corresponds to the Rao-Blackwellised version of PMC, in which we integrate over the position of all the particles.

Recently, [25] also proposed to Rao-Blackwellise the mutation kernel in PMC while keeping an importance sampling framework. The resulting algorithm is non-Markovian and does not conserve the number of particles. At each iteration a batch of particles is added to the system according to the previous estimation of the target density. The number of particles N grows with the number of iterations. Once a particle is added, its position does not change, only its weight is updated along the iterations as explained in the following Algorithm 6.

The main hyper parameters to tune are the sequence of kernels $(K_t)_t$ and mixing weights $(\lambda_t)_t$. Typically the sequence of kernels is taken equal to $K_t(x) \equiv K_{h_t}(x) := K_0(x/h_t)/h_t^d$ where K_0 is a fixed kernel and $(h_t)_t$ is a sequence of bandwidths. In [25], the authors recommend to use

$$h_t = h_0 \left(1 + \frac{2Nt}{T_0} \right)^{-1/(4+d)},$$

Input: A target distributions π ,
 an initial distribution q_0 ,
 a number of particles N ,
 a sequence of transition kernels K_t
 a sequence of mixing weights $(\lambda_t)_t$
 a number of iterations T

Output: A weighted sample $(X_t^i, w_t^i)_{1 \leq i \leq N; 0 \leq t \leq T}$

Sample independently N particles $X_0^i \sim q_0, i \in \{1, \dots, N\}$;
 Compute the importance weights $w_0^i = \pi(X_0^i)/q_0(X_0^i)$;
for $t = 1$ **to** T **do**
 Sample independently N particles from $X_t^i \sim q_t := (1 - \lambda_t)f_t + \lambda_t q_0$,
 $i \in \{1, \dots, N\}$ where

$$f_t(x) = \sum_{s=0}^{t-1} \sum_{i=1}^N \tilde{w}_s^i K_t(x - X_s^i),$$

 and the weights $\tilde{w}_s^i = w_s^i / \sum_{s=0}^{t-1} \sum_{i=1}^N w_s^i$ are normalized.
 Compute the new importance weights $w_t^i = \pi(X_t^i)/q_t(X_t^i)$;
end

Algorithm 6: Safe and Adaptive Importance Sampling (SAIS)

$$\lambda_t = \begin{cases} 0.25 \left(1 + \frac{2Nt}{T_0}\right)^{-1/(4+d)} & \text{if } t \geq T_0, \\ 0.5 & \text{if } T_0/2 < t < T_0, \\ 1 & \text{if } t \leq T_0/2. \end{cases}$$

where $T_0 \sim T/10$ is the duration of the burn-in phase and h_0 is the initial bandwidth size which becomes the only parameter to choose. During the burn-in phase the authors also recommend to raise the importance weights to the power $3/4$ in order to prevent an early degeneracy of a small number of weights. We will always use these recommendations.

This method shares some similarities with ours, an important difference being that an old particle cannot be improved through time and a “bad” particle will indefinitely remain in the system at the same place, while its weight decreases through time, eventually increasing the computation cost. A second difference is that it is necessary to tune the sequence of bandwidths and mixture weights.

Importance sampling based methods output unbiased estimators. For CMC, as stated before, except for the Metropolis-Hastings proposal, each one of the methods previously described is biased. Indeed, for a fixed number N of particles, the algorithm does not converge to the target distribution. For a large number of particles, the algorithm provides however a good approximation of the target density, according to Theorem 2.9. In addition, as a byproduct, we can re-use this approximation to provide an *unbiased* estimator by simply using the (sequence of) collective proposal distributions as importance distributions in any importance sampling based sampler. A more thorough study is left for future work, but some elements can be found in Appendix D.

C.3. Adaptive sequential Monte Carlo methods

A classical version of the Sequential Monte Carlo method well suited for sampling problems is given by Algorithm 7.

<p>Input: A sequence of target distributions π_0, \dots, π_T, a number of particles N, a sequence of transition kernels $K_{t,\xi}$ parametrized by $\xi \in E^N$, a threshold ESS_{\min}, a given allowed number S of mutation steps</p> <p>Output: A weighted sample $(X_t^i, w_t^i)_{1 \leq i \leq N; 1 \leq t \leq T}$</p> <p><i>(Initialization);</i> for $i = 1$ to N do Sample $X_0^i \sim \pi_0$. ; Set the weights $w_0^i = 1/N$; end for $t = 1$ to T do <i>(Correction) ;</i> Update the weights $w_t^i = w_{t-1}^i \frac{\pi_t(X_{t-1}^i)}{\pi_{t-1}(X_{t-1}^i)}$ for $i \in \{1, \dots, N\}$; Normalize the weights ; <i>(Resampling) ;</i> Compute $ESS = (\sum_i (w_{t-1}^i)^2)^{-1}$; if $ESS < ESS_{\min}$ then Resample the particles $X_{t-1}^i \sim \sum_j w_{t-1}^j \delta_{X_{t-1}^j}$ for all $i \in \{1, \dots, N\}$; Set $w_{t-1}^i = 1/N$ for all $i \in \{1, \dots, N\}$ end <i>(Mutation) ;</i> Set $\mathcal{X}_{t-1}^N = (X_{t-1}^1, \dots, X_{t-1}^N) \in E^N$; for $i = 1$ to N do Set $X_{0,t}^i = X_{t-1}^i$; for $s = 1$ to S do Sample $X_{s,t}^i \sim K_{t,\mathcal{X}_{t-1}^N}(X_{s-1,t}^i, dx)$; end Set $X_t^i = X_{S,t}^i$; end end</p>

Algorithm 7: Sequential Monte Carlo (SMC)

In this version, the sequence of target distributions is typically given by a tempering scheme between π_0 , which is easy to sample from, and the final distribution $\pi_T = \pi$, the target distribution that we wish to sample from. The SMC algorithm alternates between two steps. During the mutation step, the particles evolve independently using a π_t -invariant random-walk kernel. During this step, the weights of the particles are updated using a classical importance sampling procedure. Then, during the resampling step, the particles with a small weight

are eliminated and replaced by particles with a higher weight. The SMC algorithm is said to be adaptive when the transition kernel $K_{t, \mathcal{X}_{t-1}^n}$ at iteration t depends on the population of particles at the previous iteration.

This framework is general enough to encompass our main Algorithm 1, by simply taking un-weighted particles (or $ESS_{\min} = N$), a constant sequence of target distributions and $K_{t, \mathcal{X}_{t-1}^N}$ equal to (1) with μ equal to the empirical measure of the particle system. However our setting has several differences with the classical SMC methods advocated in the literature.

In all the adaptive methods proposed in the literature and that we are aware of, the choice of the random-walk kernel does not depend on the whole distribution of the particles but rather on summary statistics, that is on real-valued quantities of the form $\langle \Psi_t, \hat{\mu}_t^N \rangle$ for a well-chosen test function Ψ_t . A typical choice is to consider the covariance. Considering only summary statistics has several advantages, in particular the algorithm keeps a linear footprint in N and can be easily parallelized. Then, the mutation step is typically a standard MH algorithm whose flaws are expectedly corrected by the resampling procedure. However, the full knowledge of the particle swarm is not exploited and choosing the right summary statistics and the right mutation kernel are difficult tasks which may cause instability in practice. One goal in this article is precisely to avoid resampling and tempering by considering better mutation kernels thanks to wish we can safely rely on ergodicity only. Our main algorithm can be seen as the most general case where the mutation kernel depends on the whole distribution of the particle system and thus fully exploit the knowledge of the particle swarm. Even if the complexity of the algorithms that we propose becomes quadratic in N , we show that at least equally good results can be achieved.

On the theoretical side, this choice of mutation kernel is much more difficult to study and requires specific tools. On the contrary, due to their linear nature, the convergence properties of adaptive SMC methods with a mutation kernel which depends on summary statistics are typically studied with completely different tools. One important difference to note is that, unlike most of the particle-based algorithms in the literature, the target distribution $\pi^{\otimes N}$ is *not* an invariant distribution of the particle system. This issue, which may seem critical at first sight, is actually solved asymptotically when $N \rightarrow +\infty$ thanks to a propagation of chaos result. This enlarge the class of mutation kernels.

Finally, we mention that although we did not investigate this direction, CMC within SMC would be a very natural idea.

Appendix D: Importance sampling plug-in of CMC

Importance sampling methods are unbiased, by definition, while CMC presents a (usually small) bias depending on the number of particles. To remove this bias, while keeping the original form of the algorithm, we present in Algorithm 8 a small plug in, which output an unbiased sample without additional cost.

```

Input: An initial population of particles  $(X_0^1, \dots, X_0^N) \in E^N$ , a
maximum time  $T \in \mathbb{N}$ , a proposal distribution  $\Theta$  and an
acceptance function  $h$ 
Output:  $T$  estimators  $\hat{\varphi}_t$  of  $E_\pi[\varphi(X)]$ 
for  $t = 0$  to  $T - 1$  do
  for  $i = 1$  to  $N$  do
    Draw  $Y_t^i \sim \Theta_{\hat{\mu}_t^N}(\cdot | X_t^i)$  a proposal for the new state of particle  $i$ ;
    Compute  $\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i) = \frac{\Theta_{\hat{\mu}_t^N}(X_t^i | Y_t^i) \pi(Y_t^i)}{\Theta_{\hat{\mu}_t^N}(Y_t^i | X_t^i) \pi(X_t^i)}$ ;
    Store  $W_t^i = \pi(Y_t^i) / \Theta_{\hat{\mu}_t^N}(Y_t^i | X_t^i)$  and  $Y_t^i$ ;
    Draw  $U_t^i \sim \mathcal{U}([0, 1])$ ;
    if  $U_t^i \leq h(\alpha_{\hat{\mu}_t^N}(X_t^i, Y_t^i))$  then
      | Set  $X_{t+1}^i = Y_t^i$ ;
    else
      | Set  $X_{t+1}^i = X_t^i$ ;
    end
  end
end
Set  $\hat{\varphi}_t = \sum_i W_t^i \varphi(Y_t^i)$ .

```

Algorithm 8: Collective Monte Carlo with IS output

At each step, we can store Y_t^i the proposed state for each particle, and the numerator of the acceptance ratio $W_t^i = \pi(Y_t^i) / \Theta_{\hat{\mu}_t^N}(Y_t^i | X_t^i)$. An estimator of $\mathbb{E}_\pi[\varphi(X)]$ is then $\sum_i W_t^i \varphi(Y_t^i)$. This term corresponds exactly to the importance ratio with π as target and $\Theta_{\hat{\mu}_t^N}$ as proposal distribution. This is similar in a sense to PMC, but the main difference is that *we do not propagate* Y_t^i unless it is accepted, and thus the weight cannot degenerate.

This addition does not interfere with the choice of the interaction Θ , and ensures that the resulting estimator is unbiased. We do not have to use all the points (Y_t^i) as we can only use some of the t . Notice that each of the $(Y_t^i)_i$ are independent by construction, but that $(Y_t^i)_{t,i}$ are not independent.

To measure the quality of a numerical method, it is common to use the Effective Sample Size (ESS) [73] that represents the equivalent number of i.i.d. samples needed to build an estimator of $\mathbb{E}_\pi[X]$ with the same variance — notice that this quantity is mostly informative. In Importance sampling, the ESS is commonly estimated by $ESS(t) = (\sum_i (W_t^i)^2)^{-1}$. Interestingly enough, Theorem 2.9 shows that precisely, as the number of particles increases,

$$\mathbb{E}[(\sum_i (W_t^i)^2)^{-1}] / N \rightarrow_{N \rightarrow \infty} 1,$$

as it tends to an i.i.d. sample distributed according to π . The precise convergence speed would depend on the constants appearing in the theorems.

As usually in importance sampling, this allows to estimate at no cost the normalizing constant — or in a Bayesian framework, the marginal likelihood of a




Fig 8: Log-MSE of the normalizing constant computed for the banana-shaped distribution (left) and the 8-dimensional Gaussian mixture (right).




Fig 9: True sample and levels of the mixture of Cauchy target distribution

mode: $\int_x \pi(x)dx$, where we recall that π is non-normalised. Classical Metropolis-Hastings methods cannot estimate this quantity, while importance sampling methods (SMC, PMC) can provide an estimator whose variance may be too large to be usable. Figure 8 shows the result for the experiments of Section 5. These are only results on two “simple” targets, further numerical simulations would be needed to confirm the efficiency of our method to estimate this quantity.

Appendix E: Numerical experiments on a Cauchy mixture

Our last target is a simple mixture of two two-dimensional Cauchy distributions with means $(0.2, 0.8)$ and $(0.8, 0.2)$ and the same scale parameter 0.01. We represent in Figure 9 a sample and the level sets of the target distribution. This distribution has “heavy tails” which should reduce the efficiency of our uniform proposals on balls. For this example, we choose 0.1 for the radius of the unique




Fig 10: Mean energy distance to a true sample on 10 repetitions of the algorithm, Cauchy mixture. The dotted line represents the mean distance between two iid exact samples, computed over 100 independent realisations, and the coloured area is the corresponding 90% prediction interval.

proposals, and 0.01, 0.05, 0.1, 0.3 for the MoKAs. The results presented in Figure 10 show that none of the method reaches the minimal value possible for the Energy distance. Furthermore, MoKA-KIDS dramatically fails, this is probably because of the underlying assumptions of the deconvolution algorithm we use. This last example confirms the reliability of MoKA-Markov as an efficient and secure method. The variance of the energy distance is given in Tabular 3.

TABLE 3
Variance of the Energy distance at the last iteration for the Cauchy target.

	CMC	MoKA	MoKA-Markov	MoKA-KIDS	PMH	SAIS
Cauchy mixture	8.01e-05	6.59e-05	6.86e-05	7.40e-02	3.31e-4	8.66e-05

Acknowledgments

The authors wish to thank Pierre Degond, Robin Ryder and Christian Robert for their support and useful advice. A.D. acknowledges the hospitality of the CEREMADE, Université Paris-Dauphine where part of this research was carried out. G.C. acknowledges the hospitality of the Mathematics Department at Imperial College London where part of this research was carried out.

References

- [1] AGRAWAL, S., VATS, D., ŁATUSZYŃSKI, K. and ROBERTS, G. O. (2021). Optimal scaling of MCMC beyond Metropolis. *arXiv preprint arXiv: 2104.02020*.
- [2] ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2010). Particle Markov Chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **72** 269–342.
- [3] ANDRIEU, C., JASRA, A., DOUCET, A. and DEL MORAL, P. (2007). Non-linear Markov Chain Monte Carlo. *ESAIM: Proc.* **19** 79–84.
- [4] ANDRIEU, C., JASRA, A., DOUCET, A. and DEL MORAL, P. (2011). On nonlinear Markov Chain Monte Carlo. *Bernoulli* **17** 987–1014.
- [5] ATCHADÉ, Y., FORT, G., MOULINES, E. and PRIOURET, P. (2011). Adaptive Markov Chain Monte Carlo: Theory and Methods. In *Bayesian Time Series Models* (D. Barber, A. Taylan Cemgil and S. Chiappa, eds.) 32–51. Cambridge Univ. Press.
- [6] ATCHADÉ, Y. F. and ROSENTHAL, J. S. (2005). On adaptive Markov Chain Monte Carlo algorithms. *Bernoulli* **11** 815–828.
- [7] BARNES, J. and HUT, P. (1986). A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* **324** 446–449.
- [8] BARTIER, J.-P. and DOLBEAULT, J. (2006). Convex Sobolev inequalities and spectral gap. *C. R. Math. Acad. Sci. Paris* **342** 307–312.
- [9] BELLOMO, N., DEGOND, P. and TADMOR, E., eds. (2017). *Active Particles, Volume 1: Advances in Theory, Models, and Applications*. Birkhäuser.
- [10] BELLOMO, N., DEGOND, P. and TADMOR, E., eds. (2019). *Active Particles, Volume 2: Advances in Theory, Models, and Applications*. Birkhäuser.
- [11] BESAG, J. (1994). Comments on “Representations of knowledge in complex systems” by U. Grenander and MI Miller. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **56** 591–592. [MR1293234](#)
- [12] BHATNAGAR, P. L., GROSS, E. P. and KROOK, M. (1954). A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems. *Phys. Rev.* **94** 511.
- [13] BOLLEY, F. and VILLANI, C. (2005). Weighted Csiszár-Kullback-Pinsker inequalities and applications to transportation inequalities. *Ann. Fac. Sci. Toulouse Math.* **14** 331–352.
- [14] BORNN, L., JACOB, P. E., DEL MORAL, P. and DOUCET, A. (2013). An adaptive interacting Wang–Landau algorithm for automatic density exploration. *J. Comput. Graph. Statist.* **22** 749–773. [MR3173740](#)
- [15] BRÉMAUD, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulation and Queues. Texts in Applied Mathematics* **31**. Springer-Verlag New York. [MR1689633](#)
- [16] CAFFARELLI, L., FELDMAN, M. and MCCANN, R. (2002). Constructing optimal maps for Monge’s transport problem as a limit of strictly convex costs. *J. Amer. Math. Soc.* **15** 1–26.
- [17] CAPPÉ, O., DOUC, R., GUILLIN, A., MARIN, J.-M. and ROBERT, C. P. (2008). Adaptive importance sampling in general mixture classes. *Stat.*

- Comput.* **18** 447–459.
- [18] CAPPÉ, O., GUILLIN, A., MARIN, J.-M. and ROBERT, C. P. (2004). Population Monte Carlo. *J. Comput. Graph. Statist.* **13** 907–929.
 - [19] CARMONA, R. and DELARUE, F. (2018). *Probabilistic Theory of Mean Field Games with Applications I*. Springer.
 - [20] CHAMPION, T., DE PASCALE, L. et al. (2011). The Monge problem in \mathbb{R}^d . *Duke Math. J.* **157** 551–572. [MR2785830](#)
 - [21] CHARLIER, B., FEYDY, J. and GLAUNES, J. (2018). Kernel operations on the GPU, with autodiff, without memory overflows. <http://www.kernel-operations.io>.
 - [22] CRAIU, R. V., ROSENTHAL, J. and YANG, C. (2009). Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *J. Amer. Statist. Assoc.* **104** 1454–1466. [MR2750572](#)
 - [23] DEL MORAL, P. (2013). *Mean Field Simulation for Monte Carlo Integration*. Chapman and Hall/CRC.
 - [24] DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 411–436. [MR2278333](#)
 - [25] DELYON, B. and PORTIER, F. (2020). Safe and adaptive importance sampling: a mixture approach. *arXiv preprint arXiv: 1903.08507*.
 - [26] DESVILLETES, L., MOUHOT, C. and VILLANI, C. (2001). Celebrating Cercignani’s conjecture for the Boltzmann equation. *Kinet. Relat. Models* **4** 277–294.
 - [27] DIACONIS, P., LEBEAU, G. and MICHEL, L. (2011). Geometric analysis for the Metropolis algorithm on Lipschitz domains. *Invent. Math.* **185** 239–281.
 - [28] DIEZ, A. (2020). Propagation of chaos and moderate interaction for a piecewise deterministic system of geometrically enriched particles. *Electron. J. Probab.* **25** 1–38.
 - [29] DOBRUSHIN, R. L. (1979). Vlasov equations. *Funct. Anal. Appl.* **13** 115–123.
 - [30] DOUC, R., GUILLIN, A., MARIN, J. M. and ROBERT, C. P. (2007). Convergence of adaptive mixtures of importance sampling schemes. *Ann. Statist.* **35** 420–448.
 - [31] DOUCET, A., DE FREITAS, N. and GORDON, N. (2013). *Sequential Monte Carlo Methods in Practice*. Springer Science & Business Media.
 - [32] DUANE, S., KENNEDY, A. D., PENDLETON, B. J. and ROWETH, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B* **195** 216–222.
 - [33] ETHERIDGE, A. (2011). *Some Mathematical Models from Population Genetics*. *École D’Été de Probabilités de Saint-Flour 2012*. Springer Science & Business Media.
 - [34] FEARNHEAD, P., BIERKENS, J., POLLOCK, M., ROBERTS, G. O. et al. (2018). Piecewise deterministic Markov processes for continuous-time Monte Carlo. *Statist. Sci.* **33** 386–412.
 - [35] FEYDY, J. (2020). Geometric data analysis, beyond convolutions, PhD thesis, Université Paris-Saclay.

- [36] FEYDY, J. (2020). Geometric data analysis, beyond convolutions, PhD thesis, Université Paris-Saclay Gif-sur-Yvette, France.
- [37] FEYDY, J., GLAUNÈS, J., CHARLIER, B. and BRONSTEIN, M. (2020). Fast geometric learning with symbolic matrices. *Proc. NeurIPS* **2** 6.
- [38] FOURNIER, N. and GUILLIN, A. (2015). On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Related Fields* **162** 707–738.
- [39] GORDON, N. J., SALMOND, D. J. and SMITH, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE proceedings F (radar and signal processing)* **140** 107–113. IET.
- [40] GRAHAM, C. (1992). McKean–Vlasov Itô–Skorohod equations, and nonlinear diffusions with discrete jump sets. *Stochastic Process. Appl.* **40** 69–82.
- [41] GRAHAM, C. (1992). Nonlinear diffusion with jumps. *Ann. Inst. Henri Poincaré Probab. Stat.* **28** 393–402.
- [42] GREENGARD, L. and ROKHLIN, V. (1987). A fast algorithm for particle simulations. *J. Comput. Phys.* **73** 325–348.
- [43] HAARIO, H., SAKSMAN, E. and TAMMINEN, J. (1999). Adaptive proposal distribution for random walk Metropolis algorithm. *Comput. Statist.* **14** 375–396.
- [44] HAARIO, H., SAKSMAN, E. and TAMMINEN, J. (2001). An adaptive Metropolis algorithm. *Bernoulli* **7** 223–242.
- [45] HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109.
- [46] HAURAY, M. and MISCHLER, S. (2014). On Kac’s chaos and related problems. *J. Funct. Anal.* **266** 6055–6157.
- [47] HOCKNEY, R. W. and EASTWOOD, J. W. (1988). *Computer Simulation Using Particles*. CRC Press.
- [48] JABIN, P.-E. and WANG, Z. (2016). Mean field limit and propagation of chaos for Vlasov systems with bounded forces. *J. Funct. Anal.* **271** 3588–3627.
- [49] JASRA, A., STEPHENS, D. A. and HOLMES, C. C. (2007). On population-based simulation for static inference. *Stat. Comput.* **17** 263–279.
- [50] JIN, S., LI, L. and LIU, J.-G. (2020). Random Batch Methods (RBM) for interacting particle systems. *J. Comput. Phys.* **400** 108877. [MR4019084](#)
- [51] JOURDAIN, B., LELIÈVRE, T. and MIASOJEDOW, B. (2014). Optimal scaling for the transient phase of Metropolis Hastings algorithms: The longtime behavior. *Bernoulli* **20**.
- [52] JOURDAIN, B., LELIÈVRE, T. and MIASOJEDOW, B. (2015). Optimal scaling for the transient phase of the random walk Metropolis algorithm: The mean-field limit. *Ann. Appl. Probab.* **25**.
- [53] JOURDAIN, B. and MÉLÉARD, S. (1998). Propagation of chaos and fluctuations for a moderate model with smooth initial data. *Ann. Inst. Henri Poincaré Probab. Stat.* **34** 727–766.
- [54] JÜNGEL, A. (2016). *Entropy Methods for Diffusive Partial Differential Equations*. Springer. [MR3497125](#)
- [55] KAC, M. (1956). Foundations of kinetic theory. In *Proceedings of the Third*

- Berkeley Symposium on Mathematical Statistics and Probability* **3** 171–197. University of California Press Berkeley and Los Angeles, California.
- [56] KLAAS, M., BRIERS, M., DE FREITAS, N., DOUCET, A., MASKELL, S. and LANG, D. (2006). Fast particle smoothing: If I had a million particles. In *Proceedings of the 23rd international conference on Machine learning* 481–488.
- [57] ŁATUSZYŃSKI, K. and ROBERTS, G. O. (2013). CLTs and asymptotic variance of time-sampled Markov chains. *Methodol. Comput. Appl. Probab.* **15** 237–247.
- [58] LEE, A., YAU, C., GILES, M. B., DOUCET, A. and HOLMES, C. C. (2010). On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *J. Comput. Graph. Statist.* **19** 769–789.
- [59] LUCY, L. B. (1974). An iterative technique for the rectification of observed distributions. *Astron. J.* **79** 745.
- [60] The Mathworks, Inc. (2017). MATLAB Natick, Massachusetts.
- [61] MCKEAN, H. P. (1966). A class of Markov processes associated with non-linear parabolic equations. *Proc. Nat. Acad. Sci.* **56** 1907.
- [62] MCKEAN, H. P. (1967). Propagation of chaos for a class of non-linear parabolic equations. *Stochastic Differential Equations (Lecture Series in Differential Equations, Session 7, Catholic Univ., 1967)* 41–57.
- [63] MÉLÉARD, S. (1996). Asymptotic behaviour of some interacting particle systems; McKean–Vlasov and Boltzmann models. In *Probabilistic models for nonlinear partial differential equations* 42–95. Springer.
- [64] MENGERSEN, K. L. and TWEEDIE, R. L. (1996). Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.* **24** 101–121.
- [65] METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087–1092.
- [66] METROPOLIS, N. and ULAM, S. (1949). The Monte Carlo method. *J. Amer. Statist. Assoc.* **44** 335–341.
- [67] NATTERER, F. and WÜBBELING, F. (2001). *Mathematical Methods in Image Reconstruction. SIAM Monographs on Mathematical Modeling and Computation* **5**. Society for Industrial and Applied Mathematics. [MR1828933](#)
- [68] OELSCHLÄGER, K. (1985). A law of large numbers for moderately interacting diffusion processes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **69** 279–322.
- [69] PASZKE, A., GROSS, S., CHINTALA, S., CHANAN, G., YANG, E., DEVITO, Z., LIN, Z., DESMAISON, A., ANTIGA, L. and LERER, A. (2017). Automatic differentiation in PyTorch. In *Proceedings of Neural Information Processing Systems*.
- [70] PASZKE, A., GROSS, S., CHINTALA, S., CHANAN, G., YANG, E., DEVITO, Z., LIN, Z., DESMAISON, A., ANTIGA, L. and LERER, A. (2017). Automatic differentiation in PyTorch.
- [71] RICHARDSON, W. H. (1972). Bayesian-Based Iterative Method of Image Restoration. *J. Opt. Soc. Am.* **62** 55–59.

- [72] RIZZO, M. L. and SZÉKELY, G. J. (2016). Energy distance. *Wiley interdisciplinary reviews: Computational Statistics* **8** 27–38. [MR3457239](#)
- [73] ROBERT, C. and CASELLA, G. (2013). *Monte Carlo statistical methods*. Springer Science & Business Media.
- [74] ROKHLIN, V. (1985). Rapid solution of integral equations of classical potential theory. *J. Computat. Phys.* **60** 187–207.
- [75] SCHMEISER, C. (2018). Entropy methods <https://homepage.univie.ac.at/christian.schmeiser/Entropy-course.pdf>.
- [76] SIGURGEIRSSON, H., STUART, A. and WAN, W.-L. (2001). Algorithms for Particle-Field Simulations with Collisions. *J. Comput. Phys.* **172** 766–807.
- [77] SZÉKELY, G. J., RIZZO, M. L. et al. (2004). Testing for equal distributions in high dimension. *InterStat* **5** 1249–1272.
- [78] SZNITMAN, A.-S. (1991). Topics in propagation of chaos. In *Éc. Été Probab. St.-Flour XIX—1989* 165–251. Springer.
- [79] R CORE TEAM (2020). R: A language and environment for statistical computing, Vienna, Austria.
- [80] TOLMAN, R. C. (1979). *The principles of statistical mechanics*. Courier Corporation.
- [81] VAN DER WALT, S., COLBERT, S. C. and VAROQUAUX, G. (2011). The NumPy array: a structure for efficient numerical computation. *Computing in science & engineering* **13** 22–30.
- [82] VANETTI, P., BOUCHARD-CÔTÉ, A., DELIGIANNIDIS, G. and DOUCET, A. (2017). Piecewise-Deterministic Markov Chain Monte Carlo. *arXiv preprint arXiv: 1707.05296*.
- [83] VERLET, L. (1967). Computer” experiments” on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Phys. Rev.* **159** 98.
- [84] VILLANI, C. (2003). Cercignani’s conjecture is sometimes true and always almost true. *Comm. Math. Phys.* **234** 455–490.
- [85] VILLANI, C. (2009). *Optimal Transport, Old and New. Grundlehren der mathematischen Wissenschaften* **338**. Springer-Verlag Berlin Heidelberg.
- [86] YANG, T., LI, Y.-F., MAHDAVI, M., JIN, R. and ZHOU, Z.-H. (2012). Nystrom method vs random Fourier features: A theoretical and empirical comparison. In *Advances in neural information processing systems* 476–484.
- [87] ZIEL, F. (2021). The energy distance for ensemble and scenario reduction. *Philos. Trans. Roy. Soc. A* **379** 20190431.