Importance sampling of rare events in chaotic systems

Jorge C. Leitão^{1,2}, João M. Viana Parente Lopes^{3,4}, and Eduardo G. Altmann^{1,5,a}

¹ Max Planck Institute for the Physics of Complex Systems, 01187 Dresden, Germany

² DTU Compute, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark

³ Department of Physics and Center of Physics, University of Minho, 4710-057 Braga, Portugal

⁴ Physics Engineering Department, Engineering Faculty of the University of Porto, 4200-465 Porto, Portugal

⁵ School of Mathematics and Statistics, University of Sydney, 2006 NSW, Sydney, Australia

Received 23 January 2017 / Received in final form 4 May 2017

Published online 4 October 2017

© The Author(s) 2017. This article is published with open access at [Springerlink.com](https://www.springerlink.com)

Abstract. Finding and sampling rare trajectories in dynamical systems is a difficult computational task underlying numerous problems and applications. In this paper we show how to construct Metropolis-Hastings Monte-Carlo methods that can efficiently sample rare trajectories in the (extremely rough) phase space of chaotic systems. As examples of our general framework we compute the distribution of finite-time Lyapunov exponents (in different chaotic maps) and the distribution of escape times (in transient-chaos problems). Our methods sample exponentially rare states in polynomial number of samples (in both low- and high-dimensional systems). An open-source software that implements our algorithms and reproduces our results can be found in reference [J. Leitao, *A library to sample chaotic systems*, 2017, <https://github.com/jorgecarleitao/chaosppl>].

1 Introduction

Extreme events play a crucial role in our society. Landslides, floods, meteorite collisions, solar flares, earthquakes are all events that are rare but often lead to catastrophic consequences to our well being [1]. Science often studies extreme events by recreating the process that generates them sufficiently many times. While in some cases the process can be reproduced experimentally, in many others (e.g., astronomical and weather events) the best we can do is to simulate physical models.

Physical models often contain non-linearities in the equations of motion that amplify fluctuations and give origin to extreme events. A famous effect of non-linearity is chaos, the extremely sensitive dependence on initial conditions. Chaos has a fascinating history dating back to the seminal work of Poincaré in 1890 [2] and is today a well established field of research with applications in Biology, Geology, Economy, Chemistry, and Physics [3,4]. Chaotic dynamics often hinders our ability to study the evolution of the system analytically, e.g. it forbids describing trajectories in a closed formula. In this situation, again, often the best we can do is to numerically simulate the model in a computer, a paradigm popularized by Lorenz since the 1960s [3,5]. Extreme events are then studied statistically, over an ensemble of initial conditions.

In this paper we introduce a framework for performing numerical simulations in chaotic systems in such a

way that rare trajectories are generated more likely than if they would be generated by chance. This is an importance-sampling [6,7] (or rare-event-simulation [8]) strategy that builds on methods that have proven successful in the characterization of rare configurations in various problems. The distinguishing feature of our framework is that it considers general classes of observables [9] in *deterministic* chaotic [4,10] systems, being therefore able to find and sample initial conditions leading to extreme events in different problems. We apply our framework to two traditional problems in the numerical exploration of chaotic systems:

- trajectories with high or low finite-time Lyapunov exponents;
- long-living trajectories in open systems.

The results shown in this paper are general, but simulations are done in simple dynamical systems (time-discrete maps with up to 16 dimensions). Our goal is to illustrate the generic computational challenges of sampling rare events in chaotic systems, in line with the tradition that simple systems often possess the basic mechanisms responsible for extreme events.

The importance of rare trajectories in chaotic systems, including methods designed to find and sample them, have been studied through different perspectives [9–20]. For example, the method *stagger and dagger*, used to find long-living trajectories of transiently chaotic systems, has been an important tool to characterize chaotic saddles [10,11]. *Lyapunov weighted dynamics* is a population Monte Carlo

^a e-mail: eduardo.altmann@sydney.edu.au

method that was successfully applied to find trajectories with atypically low or high chaoticity [16–18,21,22]. Similarly, *genealogical particle analysis* was applied to compute rare events in simple climate models [20]. Another powerful and widely used method within trajectory sampling is *transition path sampling* [13,14], that has been used to sample rare trajectories (e.g. chemical reactions), typically influenced by thermal noise [14]. These are typically trajectories that transit from one stable configuration to another stable configuration [14,23]. These different methods achieve their goal through different, often ingenious, solutions. Here we aim to construct methods that are applicable to different classes of problems and to quantitatively understand the impact of different parameters and choices on the efficiency of the algorithm. We then show that some of the existing methods can be derived from our construction under suitable approximations.

Our methods are based on Metropolis-Hastings (MH) importance sampling [6,7], a well established numerical technique that has been used in Statistical Physics to study rare events since the 1950s. MH produces a random walk $\mathbf{x} \rightarrow \mathbf{x}'$ in the phase-space of the system that generates initial conditions leading to extreme events (rare states) more often than they would be found by chance, consequently reducing the computational cost associated with their rareness. The flexibility of MH is confirmed by its success in numerous fields in Physics, Chemistry, Finance, among many others [6,7]. While transition path sampling [14], Lyapunov weighted dynamics [16], and genealogical methods [20] already use importance sampling in deterministic chaotic systems, three fundamental questions remain largely open: (1) can MH be used efficiently to systematically sample rare trajectories of (deterministic) chaotic systems? If yes, (2) how and (3) at what (computational) cost?

We answer these questions with a framework that constructs MH algorithms to *efficiently* sample rare events in chaotic systems. The crucial step our framework solves is to derive a proposal distribution that bounds the MH acceptance rate for a broad class of chaotic systems and observables (e.g. maximal Lyapunov exponent, escape time). More specifically, we show how to incorporate properties of trajectories of the dynamical system in the MH proposal distribution and that this is a necessary condition to obtain efficient algorithms. This allows us to efficiently sample rare events in different problems and classes of chaotic systems. Proposal distributions existent in the literature [11,23–25] can be derived as limiting cases of our general framework. We expect the ideas and formalism presented here to find applications in other problems of chaotic systems and in the study of extreme events more generally. Therefore, we expect our results to be useful both to those studying extreme events in non-linear systems and to those studying numerical techniques.

The paper is organized as follows: Section 2 reviews traditional numerical problems in chaotic systems and shows how they can be formulated as a problem of sampling rare trajectories; Section 3 introduces the MH algorithm as a method to perform importance-sampling

simulations in chaotic systems, and shows that naive approaches do not lead to an efficient MH; Section 4 shows how to incorporate general features of chaotic systems, such as exponential divergence or self-similarity of some of its properties, in the proposal distribution in order to achieve an efficient MH; Section 5 presents numerical tests of the general framework that confirm the applicability of Monte Carlo algorithms to sample rare events in different classes of chaotic systems; Section 6 summarizes our results and discusses its implications.

2 Review of problems

We consider dynamical systems whose states $\mathbf{x}$ in a phase space Ω , $\mathbf{x} \in \Omega \subset \mathbb{R}^D$, evolve in time from an initial condition $\mathbf{x} = \mathbf{x}_0$ according to

$$\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t) = \mathbf{F}^{t+1}(\mathbf{x}_0), \quad (1)$$

where $\mathbf{F}(\mathbf{x}) \in \Omega$ and $\mathbf{F}^t$ is $\mathbf{F}$ composed t times, $\mathbf{F}^t(\mathbf{x}) \equiv \mathbf{F}(\mathbf{F}(\dots\mathbf{F}(\mathbf{x})\dots))$. Such discrete-time systems can be obtained from continuous-time systems through a Poincare surface of section, a stroboscopic projection, or by the time discretization of the numerical integration of differential equations. We are also interested in the dynamics in the tangent space, which quantifies the divergence between two initial conditions [4]. Specifically, the distance of a state displaced from $\mathbf{x}$ by $\mathbf{h}$, $\mathbf{x}' = \mathbf{x} + \mathbf{h}$, to the original state $\mathbf{x}$ evolves in time according to

$$\mathbf{x}'_t - \mathbf{x}_t = \mathbf{F}^t(\mathbf{x}') - \mathbf{F}^t(\mathbf{x}). \quad (2)$$

Expanding $\mathbf{F}^t(\mathbf{x}')$ around $\mathbf{F}^t(\mathbf{x})$ allows $\mathbf{x}'_t - \mathbf{x}_t$ to be written as

$$\mathbf{x}'_t - \mathbf{x}_t = J_t(\mathbf{x}) \cdot \mathbf{h} + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 \mathbf{F}^t(\mathbf{x})}{\partial x_i \partial x_j} h_i h_j + O(|\mathbf{h}|^3) \quad (3)$$

where $J_t(\mathbf{x}) \equiv d\mathbf{F}^t(\mathbf{x})/d\mathbf{x}$ is the Jacobian matrix of $\mathbf{F}^t$, and $\partial^2 \mathbf{F}^t(\mathbf{x})/(\partial x_i \partial x_j)$ is the (i, j) entry of the Hessian matrix of $\mathbf{F}$. The first term of equation (3) can be expanded using the derivative of the composition and be written as

$$D(\mathbf{x}, \mathbf{h}, t) \equiv J_t(\mathbf{x}) \cdot \mathbf{h} = \left(\prod_{i=t-1}^0 J(\mathbf{x}_i) \right) \cdot \mathbf{h} = \mathbf{h}_t \quad (4)$$

where $J \equiv J_1$ and $\mathbf{h}_t$ evolves in the tangent space according to

$$\mathbf{h}_0 = \mathbf{h}; \quad \mathbf{h}_{i+1} = J(\mathbf{x}_i) \cdot \mathbf{h}_i. \quad (5)$$

For small $|\mathbf{h}|$, the growth of $\mathbf{x}'_t - \mathbf{x}_t$ is characterized by the eigenvalues and eigenvectors of J^t [26]. The largest finite-time Lyapunov exponent (FTLE) of a point $\mathbf{x}$ can be defined¹ as

$$\lambda_t(\mathbf{x}) = \frac{\log(\mu_1)}{t}, \quad (6)$$

¹ This is sometimes denoted as the stability exponent. See Chapters 4 and 6 of reference [26].

where μ_1 is the real part of the largest eigenvalue of J^t . Thus, when at least one direction is unstable ($\mu_1 > 1$), $\mathbf{x}'_t - \mathbf{x}_t$ increases exponentially with time, and at most by

$$\mathbf{x}'_t - \mathbf{x}_t = \delta_0 e^{\lambda_t(\mathbf{x})t}. \quad (7)$$

When the system is one dimensional, the “Jacobian matrix” is a single number, the product in equation (4) is a product of numbers, and the only “eigenvalue” is the result of this product. Thus, in this case equation (6) can be written as

$$\lambda_t(\mathbf{x}) = \frac{1}{t} \sum_{i=0}^{t-1} \log \left| \frac{dF(\mathbf{x}_i)}{dx} \right|. \quad (8)$$

We now describe two computational problems in chaotic dynamical systems.

2.1 Variability of trajectories' chaoticity

For a chaotic system, $\lambda_L \equiv \lambda_{t \rightarrow \infty} > 0$ in equation (6). The variation of the (maximum) *finite-time* Lyapunov exponents λ_t (FTLE) across different initial conditions characterizes the variation of chaoticity of the system, yielding a distribution of the FTLE, $P(\lambda_{t_o})$, or, equivalently, the distribution of $E \equiv t_o \lambda_{t_o}$:

$$P(E) = \int \delta(E - t_o \lambda_{t_o}(\mathbf{x})) U(\mathbf{x}) d\mathbf{x}. \quad (9)$$

where $U(\mathbf{x})$ is a chosen probability distribution (e.g. uniform in the phase-space). The FTLE and its distribution was introduced in the 1980s [27–29] and has been used to study turbulent flows [30], Hamiltonian dynamics [31], chimera states [32], characterize dynamical trapping [27,33–36], among others [37,38]. The distribution of FTLE is related to the generalized dimensions [39] and often follows a large deviation principle, where t_o is the extensive parameter [19,39]. In strongly chaotic systems, the distribution of FTLE is Gaussian [39], whereas for intermittent chaos and other weakly chaotic systems the distribution is typically non-Gaussian [40].

Figure 1 shows the phase-space dependency and distribution of the FTLE in one chaotic system. It contains characteristic features observed in strongly chaotic systems: λ_t is a quantity that depends sensitively on the state $\mathbf{x}$, and its distribution $P(\lambda_t)$ decays exponentially to zero on both sides. In the limit $t_o \rightarrow \infty$, $P(\lambda_{t_o}) \rightarrow \delta(\lambda_{t_o} - \lambda_L)$.

The tails of the distribution $P(E)$ play a significant role in the characterization of chaotic systems, as, for example, the higher moments of the distribution are related to higher q s in the generalized dimensions D_q of the attractor [39]. Furthermore, the regions of the phase-space with small (large) finite-time Lyapunov exponent are associated with slow (fast) decay of correlations [29], and their characterization has been used to get insight on whether the system is ergodic or not [31]. Moreover, trajectories characterized by a low or high finite-time Lyapunov exponent

Fig. 1. Two main characteristics of the FTLE. Upper panel: the intricate dependency of $\lambda_{t_o=4}(\mathbf{x})$ (z -axis) with the state $\mathbf{x}$ (2D, x and y axis). Lower panel: the distribution of FTLE for different finite times t shows exponentially decaying tails with λ_{t_o} , and decay with t_o . The system used was the Standard Map (Eq. (A.4)) with $K = 6$, over the full phase-space, $\Gamma = \Omega = [0, 1]^2$. $P(\lambda_{t_o})$ was computed from 10^5 uniformly distributed initial conditions on Γ .

can play a significant role in the dynamics of interfaces in chaotic flows [30] and others [16].

A typical analysis of the FTLE is to measure how a quantity, $W(\mathbf{x})$, depends on the FTLE λ_{t_o} [29–31]. This requires estimating an integral of the form

$$W(\lambda_{t_o}) \equiv \int_{\Gamma} W(\mathbf{x}) \delta(\lambda_{t_o} - \lambda_{t_o}(\mathbf{x})) U(\mathbf{x}) d\mathbf{x} \quad (10)$$

where $W(\mathbf{x})$ is the pre-selected quantity, Γ is a pre-selected sampling region (often $\Gamma = \Omega$), and $U(\mathbf{x})$ is the weight attributed to $\mathbf{x}$ (often uniform, $U(\mathbf{x}) = 1/|\Gamma|$). Let us look for two examples. First, consider the problem of estimating the distribution $P(\lambda_{t_o})$, e.g. references [16,29,31,40]. The traditional technique is to sample M states $\mathbf{x}$ according to $U(\mathbf{x}) = 1/|\Gamma|$ and estimate $P(\lambda_{t_o})$ using the estimator M_{λ}/M , where M_{λ} is the number of samples with $\lambda_{t_o}(\mathbf{x}) \in [\lambda_{t_o}, \lambda_{t_o} + \Delta\lambda_{t_o}]$. Formally, this corresponds to $W(\lambda_{t_o})$ with $W(\mathbf{x}) = 1$. The second example is retrieved from reference [29]. There, in order to evaluate the contribution of the algebraic region

of the phase-space to the power spectra, the authors decomposed it in two terms corresponding to the power spectra of trajectories with low FTLE and trajectories with high FTLE. This required estimating the power spectra from a set of trajectories conditioned to an interval of λ s on the tails of the distribution. Associating the power spectrum ($S(f)$ in the reference) with $W(\mathbf{x}) = W_f(\mathbf{x})$ where f is the frequency, the power spectra represented in Figure 4 of reference [29] corresponds to integrals (for different frequencies f) given by

$$\mathbb{E}[W_f|\lambda_{t_o}] = \frac{1}{P(\lambda_{t_o})} \int_{\Gamma} W_f(\mathbf{x}) \delta(\lambda_{t_o} - \lambda_{t_o}(\mathbf{x})) U(\mathbf{x}) d\mathbf{x}, \quad (11)$$

which contains integrals of the form of equation (10).

Numerically estimating the integral in equation (10) is challenging for two reasons: first, $P(\lambda_{t_o})$ decays with t_o (lower panel of Fig. 1): the distribution of FTLE often follows a large deviation principle, $P(\lambda_{t_o}(\mathbf{x})) \propto \exp(t_o s(\lambda_{t_o}))$, where s is intensive in respect to t_o and is often concave [39]. Consequently, the traditional methodology of sampling states uniformly to find or sample states with increasing t_o requires an exponentially high number of initial conditions. Second, the dependency of $\lambda_{t_o}(\mathbf{x})$ on $\mathbf{x}$ has multiple local minima and maxima (upper panel of Fig. 1). Such rough (fractal [41]) landscapes are known to challenge numerical techniques (e.g., simulations get trapped in local minima or maxima) [42].

The problem of finding and sampling states with high or low- λ has been addressed in the literature with numerical techniques that go beyond traditional uniform sampling [16,18,19]. Such techniques have been successfully applied to find [16] and sample [18,19] states with extremal λ s in different chaotic systems. References [16,19] use a population Monte Carlo canonical ensemble where λ_t plays the role of the energy E to find or sample states with high or low FTLE. The method computes stochastic trajectories that, from the numerical tests performed, are indistinguishable from (deterministic) trajectories of the system. Reference [18] proposes a flat-histogram simulation to find high or low chaotic states by developing an observable to quantify the chaoticity of the state.

2.2 Transient chaos

The best known examples of chaotic systems have a fractal attractor in the phase space [4], the Lorenz attractor being the most prominent example [5]. However, chaotic dynamics can appear also when the fractal invariant set in the phase space is not purely attracting, e.g. it may have stable and unstable directions (a saddle). Trajectories close to this set perform chaotic motion for an arbitrarily long (yet finite) time. This phenomenon appears in a variety of physical systems and is known as transient chaos [43,44].

Numerical investigations of transiently-chaotic systems are computationally difficult because most trajectories quickly escape the vicinity of the chaotic saddle (on which the chaotic dynamics is properly defined) [11,12,45].

Fig. 2. Main characteristics of the escape time of an open chaotic system. Upper panel: the dependency of $t_e(\mathbf{x})$ with the state $\mathbf{x}$, shows an intricate landscape with multiple local and global maxima. The map is the 4-dimensional coupled Hénon map (defined in Appendix A), and the two dimensions are a surface of section on the plane $x_2 = 0, y_2 = 0$. Lower panel: the exponential decay of the distribution of escape time of (1) the Coupled Hénon map with $D = 4$ and (2) the Standard Map, both defined in Appendix A. $P(t_e = t)$ was computed by uniformly drawing 10^6 states $\mathbf{x} \in \Gamma$, and measure the relative number of times that $t_e(\mathbf{x}) = t_e$.

More precisely, the escape time of a state $\mathbf{x}$ in a pre-selected region $\mathbf{x} \in \Gamma \subset \Omega$ is defined as the first passage time of its trajectory to a pre-selected exit set $\Lambda \subset \Omega$:

$$t_e(\mathbf{x}) \equiv \min\{t : \mathbf{F}^t(\mathbf{x}) \in \Lambda\}. \quad (12)$$

Almost all trajectories start at Γ and eventually leave when $\mathbf{x}_t \in \Lambda$, but they do so at different times t_e . Such variability is quantified by the distribution of escape time: the probability that a random initial condition $\mathbf{x}$ leaves at a given time t_e ,

$$P(t_e) = \int_{\Gamma} \delta_{t_e, t_e(\mathbf{x})} U(\mathbf{x}) d\mathbf{x}, \quad (13)$$

where $\delta_{t_e, t_e(\mathbf{x})}$ is the Kronecker delta and $U(\mathbf{x})$ is the probability density assigned to each state in Γ , which is often constant, $U(\mathbf{x}) = 1/\int_{\Gamma} d\mathbf{x} \equiv 1/|\Gamma|$. Figure 2 shows the

typical features of $t_e(\mathbf{x})$: it strongly depends on the initial state $\mathbf{x}$ and its distribution $P(t_e)$ decays exponentially to zero (i.e., it has a constant escape rate κ).

There are two main numerical techniques to study transiently chaotic systems. The first is to *find* one long living state $\mathbf{x}$ with $t_e(\mathbf{x}) \gg 1$ and to compute an average over states of this trajectory [10,11]. For large $t_e(\mathbf{x})$, the trajectory between times $[t_e(\mathbf{x})/2 - t_s, t_e(\mathbf{x})/2 + t_s]$ where $t_s \ll t_e(\mathbf{x})/2$ is close to a trajectory on the chaotic saddle. When the saddle is ergodic, an average over this long trajectory corresponds to an average over the natural invariant density and therefore an average over these states characterizes invariant properties of the system. The second technique, which we focus in this work, is to compute averages over an ensemble $U(\mathbf{x})$ of initial conditions $\mathbf{x}$ that leave the system at time $t_e(\mathbf{x}) = t_e$ [10,44]. For small t_e observations depend on the particular initial density $U(\mathbf{x})$ (a point of interest in itself [44]), while for large t they characterize invariant properties of the system (like in the previous technique, the states $\mathbf{F}^{t/2}(\mathbf{x})$ with $t \rightarrow \infty$ are independent samples of the natural invariant density).

A typical analysis within sampling transiently chaotic systems is to measure how a quantity, $W(\mathbf{x})$, changes with increasing escape time t_e . Numerically, this can be written as

$$W(t_e) \equiv \int_{\Gamma} W(\mathbf{x}) \delta_{t_e, t_e(\mathbf{x})} U(\mathbf{x}) d\mathbf{x}. \quad (14)$$

Let us enumerate 3 examples of computational problems that can be interpreted as numerical estimations of an integral of the form of $W(t_e)$ in equation (14).

2.2.1 Compute the escape time distribution

Numerically, $P(t_e)$ is often computed by drawing states from $U(\mathbf{x})$ (e.g. uniform density), and counting the relative number of states that exited at escape time t_e [10]. This corresponds to $W(\mathbf{x}) = 1$ in which case equation (14) reduces to equation (13).

2.2.2 Compute generalized dimensions of the chaotic saddle

The generalized dimensions are an important property of the chaotic saddle and its calculation is often performed by box counting [4,10,46]. Essentially, the phase-space is divided in equal, non-overlapping, and space-filling boxes $i = 1, \dots, B(\varepsilon)$ of linear size ε (intervals in 1 dimension, squares in 2, etc.) and the generalized dimension D_q of exponent q is proportional to $\log \sum_i^B(\varepsilon) \mu_i^q$, where μ_i is the fraction of points of the saddle that belong to the box i [46]. Numerically, μ_i is estimated by first obtaining a set of points $\mathbf{x}_j$ in the saddle and then counting how many are in the particular box i . Such an estimate can be written as the expectation of an indicator function that tells whether a state in the saddle, $\mathbf{F}^{t_e/2}(\mathbf{x})$ for $t_e(\mathbf{x}) \gg 0$, is inside the box i , $W(\mathbf{x}) = \delta_{\mathbf{F}^{t_e/2}(\mathbf{x}) \in i}$. This expectation can be written as a conditional expectation of $W(\mathbf{x})$ over states that leave at time t_e ,

$$\mathbb{E}[W|t_e] \equiv \frac{1}{P(t_e)} \int_{\Gamma} W(\mathbf{x}) \delta_{t_e, t_e(\mathbf{x})} U(\mathbf{x}) d\mathbf{x}. \quad (15)$$

Computing this essentially requires computing integrals of the form of $W(t_e)$ in equation (14).

2.2.3 Compute the distribution of FTLE on the chaotic saddle

The distribution of FTLE $P(\lambda)$ is another important property of the chaotic saddle [10,46]. The FTLE $\lambda = \lambda_t$ for a fixed t of an open system is computed for trajectories on the chaotic saddle. Like in the previous problem, each of these trajectories can be obtained by generating a state $\mathbf{x}$ according to $U(\mathbf{x})$ (e.g. uniformly distributed in Γ) that has a large escape time, $t_e(\mathbf{x}) \gg 1$, and compute $\lambda_t(\mathbf{F}^{t_e/2}(\mathbf{x}))$ using equation (6) for a fixed t . The distribution of FTLE is then computed from an ensemble of these high-escape-time states by constructing different bins of the histogram $I_\lambda = [\lambda, \lambda + \Delta\lambda]$, and numerically compute the relative number of states in each bin. Formally, this equates to compute the expected number of states with a given escape time t_e whose FTLE is in a bin, and thus corresponds to computing a conditional expectation of the form of equation (15) with $W(\mathbf{x}) = W_\lambda(\mathbf{x}) = \delta_{\lambda(\mathbf{F}^{t_e/2}(\mathbf{x})) \in I_\lambda}$.

2.3 Summary: numerical problems in the study of rare events in chaotic systems

To provide an unified treatment of the numerical problems in transient chaos and in computing FTLE of closed systems we use a common notation whenever possible. It indicates also how the methods can be generalized to different problems. Firstly, there is a quantity that we call “observable” and denote by $E(\mathbf{x})$:

- $E(\mathbf{x}) = t_o \lambda_{t_o}(\mathbf{x})$ in FTLE of closed systems;
- $E(\mathbf{x}) = t_e(\mathbf{x})$ in strongly chaotic open systems.

Secondly, we use $\delta(E, E')$ to denote both the Kronecker delta and the Dirac delta ($\delta(E - E')$), for discrete and continuous case respectively. This is needed because, even though in practice we will always consider E to be a discrete function due to binning of the histograms, formally the observable E is either discrete (t_e for maps) or continuous (λ_t or t_e for flows). Thirdly, the observable E is a function of the phase-space of the system but typically depends also on an external quantity N that parameterizes how high E can be:

- t_o in FTLE of closed systems;
- (a pre-selected) maximum escape time, $t_{\max}$, in strongly chaotic open systems.

This parameter is important to us because, the higher it is, the rarer a state with the maximum or minimum possible E is. In this notation, the two general problems presented in the previous sections can be summarized as follows:

- the analysis is conditioned to a projection of $\mathbf{x}$ into a one-dimensional observable $E(\mathbf{x})$;
- the probability distribution of E , $P(E)$, decays exponentially with increasing E ;

- the focus of the analysis is on rare states in respect to $P(E)$: states with E in one of the tails of $P(E)$;
- the rareness increases with N .

In transient chaos problems we are interested in states with increasing $N = t_{\max}$ that correspond to the tail of $P(E) = P(t_e)$, $P(t_{\max})$. In computations of the FTLE in closed systems we are interested in states with increasing $N = t_o$ and on the tails of $P(E)$. These problems share two distinct computational problems: find rare states [10–12,16] and sample rare states [18,19,24,25], which can be formalized as follows:

- Find rare states: minimize or maximize $E(\mathbf{x})$ over a constraining region Γ , $\mathbf{x} \in \Gamma \subset \Omega$, for increasing N ;
- Sample rare states: compute the integral of a function $W(\mathbf{x})$ conditioned to a particular value of E and for increasing N , over an ensemble of states $\mathbf{x}$ distributed according to $U(\mathbf{x})$ in a constraining region of the phase-space $\mathbf{x} \in \Gamma \subset \Omega$:

$$W(E) = \int_{\Gamma} W(\mathbf{x}) \delta(E, E(\mathbf{x})) U(\mathbf{x}) d\mathbf{x}. \quad (16)$$

The numerical challenges of these two numerical problems are common: the states $\mathbf{x}$ are exponentially difficult to find with increasing N (or E), the function $E(\mathbf{x})$ contains multiple local and global minima embedded in fractal-like structures, and a potential high dimensionality D of the phase-space (see Figs. 1 and 2). The goal of this paper is to develop a systematic approach to tackle the two numerical problems and these three challenges.

3 Review of methods

3.1 Importance sampling

As described in the previous section, the traditional methodology to compute an integral of the form of equation (16) is to draw m samples $\{\mathbf{x}_i\}$ distributed according to $U(\mathbf{x})$ from the relevant region Γ and approximate the integral $W(E)$ by the estimator

$$\overline{W(E)} \equiv \frac{1}{m} \sum_{i=1}^m \delta(E, E(\mathbf{x}_i)) W(\mathbf{x}_i) \quad (17)$$

where $\delta(E, E(\mathbf{x})) = 1$ when $E(\mathbf{x}_i) \in [E, E + \Delta E]$ and zero otherwise (when E is discrete, $\Delta E = 1$). The relative distance of the estimator $\overline{W(E)}$ to $\mathbb{E}[W(E)]$ is quantified by the ratio $\epsilon(E) \equiv \sigma[\overline{W(E)}] / \mathbb{E}[W(E)]$ where $\sigma[\overline{W(E)}]^2 \equiv \mathbb{E}[\overline{W(E)}^2] - \mathbb{E}[\overline{W(E)}]^2$. For the estimator in equation (17), this is given by

$$\epsilon(E) \propto \frac{1}{\sqrt{mG(E)}}, \quad (18)$$

where $G(E)$ is the density of states: the number of states per bin with an observable E .² Therefore, the number of

² $G(E) = P(E)$ when each state is equally weighted, $U(\mathbf{x}) = 1/|\Gamma|$.

samples m_* required to achieve a given precision ϵ_* for a given E_* is $m_*(E_*) \propto 1/G(E_*)$. The critical problem in sampling rare states is that, because $G(E)$ decays exponentially with E , $m_*(E)$ increases exponentially with E .

Importance sampling techniques aim to improve this scaling by drawing samples from a distribution $\pi(\mathbf{x}) \neq U(\mathbf{x})$ on the phase-space (e.g. non-uniformly in the phase-space) [7]. Specifically, consider m independent samples $\{\mathbf{x}_i\}$ drawn from $\pi(\mathbf{x})$, and the function $\pi(\mathbf{x})$ to depend only on E , $\pi(\mathbf{x}) = \pi(E(\mathbf{x})) = \pi(E)$. Because the samples are not drawn from $U(\mathbf{x})$, the estimator in equation (17) would be biased. Importance sampling uses an unbiased estimator for $W(E)$ given by [7]

$$\overline{W(E)} \equiv \frac{1}{m} \sum_{i=1}^m \delta(E, E(\mathbf{x}_i)) W(\mathbf{x}_i) \frac{U(\mathbf{x}_i)}{\pi(\mathbf{x}_i)}, \quad (19)$$

which reduces to equation (17) when $\pi(\mathbf{x}) = U(\mathbf{x})$. The advantage of importance sampling is that the relative error of the estimator in equation (19) is given by

$$\epsilon(E) \propto \frac{1}{\sqrt{mG(E)\pi(E)}}. \quad (20)$$

This is because, when sampling from $\pi(\mathbf{x}) = \pi(E(\mathbf{x}))$, the expected number of samples with a given E , $m(E)$, is equal to

$$m(E) = mG(E)\pi(E). \quad (21)$$

Equation (20) implies that the function $\pi(E)$ can be chosen to favour states $\mathbf{x}$ with observable E on the tails of $G(E)$ and therefore improve the precision of the estimator on these tails.

The standard deviations in equations (18) and (20) were obtained assuming that the m samples were independent. In traditional methodologies such as uniform sampling, this is the case. However, in the algorithms discussed below, it is not. Therefore, it is necessary to modify equation (20) for the case where the samples $\{\mathbf{x}_i\}$ are drawn from $\pi(\mathbf{x})$ and are also correlated. This modification is given by

$$\epsilon(E) \propto \sqrt{\frac{1 + 2T(E)}{mG(E)\pi(E)}}, \quad (22)$$

where $T(E)$ is the autocorrelation time [7], which increases with the correlation of the samples.

Efficient importance-sampling [6–8] techniques have to address the following three steps:

1. choose a suitable $\pi(\mathbf{x})$;
2. have a method to generate samples from $\pi(\mathbf{x})$;
3. minimize the autocorrelation time T .

A defining point in our method is our choice for the Metropolis-Hasting algorithm to address point 2, differently from references [16,20] which address similar problems through a different choice for point 2 (cloning trajectories). Below we first discuss point 2, then point 1, and finally point 3.

3.2 Metropolis-Hastings algorithm

The Metropolis-Hastings (MH) algorithm asymptotically generates states $\mathbf{x}$ according to $\pi(\mathbf{x})$ using a Markovian, ergodic and detailed balance random walk in the sampling region Γ [7]. This random walk is initialized from a random state $\mathbf{x} \in \Gamma$ and evolves to a new state $\mathbf{x}' \in \Gamma$ with a transition probability $P(\mathbf{x}'|\mathbf{x})$ chosen such that asymptotically the states $\mathbf{x}$ are drawn according to $\pi(\mathbf{x})$. In Metropolis-Hastings, $P(\mathbf{x}'|\mathbf{x})$ is written as

$$P(\mathbf{x}'|\mathbf{x}) = g(\mathbf{x}'|\mathbf{x})a(\mathbf{x}'|\mathbf{x}),$$

where $g(\mathbf{x}'|\mathbf{x})$ is the (proposal) distribution used to generate new states and $a(\mathbf{x}'|\mathbf{x})$ is the (acceptance) distribution used to select them. The random walk fulfills detailed balance because the acceptance probability is chosen as [7]

$$a(\mathbf{x}'|\mathbf{x}) = \min \left(1, \frac{g(\mathbf{x}|\mathbf{x}') \pi(\mathbf{x}')}{g(\mathbf{x}'|\mathbf{x}) \pi(\mathbf{x})} \right). \quad (23)$$

Metropolis-Hastings algorithm is not the only way to achieve this. Another popular and alternative method to sample from $\pi(\mathbf{x})$ is population Monte Carlo, which instead of a random walk, uses multiple stochastic trajectories that are cloned and destroyed. See, e.g. reference [19] for an application to the problem of the FTLE described above.

Algorithmically, the MH algorithm is implemented as follows: choose a region Γ and a random initial condition $\mathbf{x} \in \Gamma$. Evolve the random walk in time according to:

1. Propose a state $\mathbf{x}'$ drawn from $g(\mathbf{x}'|\mathbf{x})$;
2. Compute $a(\mathbf{x}'|\mathbf{x})$ replacing $\mathbf{x}$ and $\mathbf{x}'$ in equation (23);
3. Generate a random number r in $[0, 1]$. If $r < a(\mathbf{x}'|\mathbf{x})$, make $\mathbf{x}'$ to be the new $\mathbf{x}$;
4. Store $\mathbf{x}$ and go to 1.

The set of sub-steps 1-4 brings the random walk from its current state $\mathbf{x}$ to the next state and it is called a Markov step. After a transient number of steps where the algorithm converges to the asymptotic distribution, the stored states $\mathbf{x}$ are (correlated) samples drawn from $\pi(\mathbf{x})$, and can be directly used in equation (19).

3.3 Sampling distribution

3.3.1 Canonical ensemble

A sampling distribution $\pi(\mathbf{x})$ often used is the canonical distribution [6]

$$\pi(\mathbf{x}) = \pi(E(\mathbf{x})) \propto e^{-\beta E(\mathbf{x})}. \quad (24)$$

In the context of rare states, the canonical ensemble is useful because the number of sampled states $m(E)$ in equation (21) becomes

$$m(E) \propto mG(E)e^{-\beta E} \propto me^{-\beta E + S(E)}. \quad (25)$$

In particular, the maximum of m is at E^* solution of $\beta = dS/dE(E^*)$. Therefore, β tunes which value of the observable E is sampled the most. For example, the Lyapunov weighted dynamics in reference [19] uses this distribution (Eq. (15) of the reference with α replaced by β).

3.3.2 Flat-histogram ensemble

Another distribution often used in the literature of Metropolis-Hastings [47–54] is the flat-histogram (or multicanonical), given by [52,53]

$$\pi(\mathbf{x}) = \pi(E(\mathbf{x})) \propto \frac{1}{G(E(\mathbf{x}))}, \quad E \in [E_{\min}, E_{\max}], \quad (26)$$

for a given choice of $E_{\min}, E_{\max}$ that defines the region of interest on the observable E . This is known as flat-histogram because, replacing equation (26) in equation (21) leads to a constant average number of samples on each E ,

$$m(E) = \text{const}. \quad (27)$$

Consequently, the dependence of the variance in equation (22) is only due to the autocorrelation $T(E)$, which implies that the computational cost to draw a state on the tail of $G(E)$ is no longer driven by the exponential decrease of $G(E)$, but by the computational cost to draw uncorrelated samples from $\pi(\mathbf{x})$.

The main limitation of the flat-histogram is that it requires knowing $G(E)$ in advance, which is very often unknown. The most well known, that we use here, is the Wang-Landau algorithm [53], that modifies the Metropolis-Hastings algorithm to a non-Markovian chain that asymptotically converges to a flat-histogram Metropolis-Hastings. The Wang-Landau algorithm starts with an approximation of $G(E)$, $G_{WL}(E) = 1/|E_{\max} - E_{\min}|$, and, on step 4 of the MH (see algorithm in Sect. 3.2), it multiplies $G_{WL}(E(\mathbf{x}))$ ($\mathbf{x}$ is the current state of the random walk) by a constant $f > 1$. After a given number of steps, f is reduced by a factor 2 and this procedure is repeated until a final $f_{\min} \simeq 1$ is reached [53]. The value $f_{\min}$ and how f is reduced dictates how close $G_{WL}(E)$ will be from $G(E)$ [53,55,56]. An alternative approach to sample states using a broad histogram is replica exchange, see references [17,18,22] for applications in chaotic systems.

3.4 Characterisation of the efficiency

The relative error $\epsilon(E)$ depends on the value of N and on the particular value of E/N . To avoid discussing the dependency on E/N , the efficiency of the flat-histogram is often quantified in terms of the average round-trip [57–59], which is an upper bound for the number of Markov steps m (samples) required to obtain an uncorrelated sample from $\pi(\mathbf{x})$ [59,60]. The round-trip, τ , is the average number of steps required for the algorithm to go from a state $\mathbf{x}$ with $E(\mathbf{x}) = E_{\min}$ to a state $\mathbf{x}$ with $E(\mathbf{x}) = E_{\max}$ and return back.³ Reference [24] shows how the autocorrelation time of a canonical ensemble is related to the autocorrelation time of a flat-histogram, and therefore we will

³ Numerically, the round-trip time is computed by having a Boolean value tracking whether the random walk is moving in the direction $E_{\min} \rightarrow E_{\max}$ or $E_{\max} \rightarrow E_{\min}$, and using it to measure the total number of steps required to make the full round-trip.

use here the round-trip time of a flat-histogram simulation to quantify the efficiency for both distributions. The uniform sampling has a round-trip that increases exponentially with N : on average it takes $1/G(E_{\max})$ samples to get one sample with $E_{\max}$, and $1/G(E_{\max})$ increases exponentially with N .

Importance sampling Monte Carlo is widely used in statistical physics because the computational cost often scales polynomially with N [7,59], which is a dramatic improvement over uniform sampling. Under the hypothesis that (a) $\Delta E \equiv E(\mathbf{x}') - E(\mathbf{x}) \approx 1 \ll N$ and (b) the correlation between the different $E(\mathbf{x})$ of the random walk decay fast enough, it can be shown that the roundtrip τ scales as [25,59]

$$\tau(N) \sim N^2. \quad (28)$$

For example, consider the problem of sampling states of an open chaotic system with different escape times t_e from $t_e = 1$ up to $t_e = t_e^{\max}$. In this case, $E_{\min} = 1$ and $E_{\max} = N = t_e^{\max}$. Thus, under the hypothesis (a) and (b) above, the round-trip is expected to scale as

$$\tau(t_e^{\max}) \sim (t_e^{\max})^2. \quad (29)$$

There are known deviations of the scaling in equation (28) leading to a higher exponent, N^{2+z} with $z > 0$ [25,57–59] and there are two common situations where this happens: (1) the acceptance rate decreases drastically with N (hypothesis (a) above is violated); (2) autocorrelation drastically increases with N (hypothesis (b) above is violated) [25,59]. Nevertheless, these do not qualitatively change the argument: Monte Carlo with flat-histogram has the potential of generating rare states polynomially with increasing N , while uniform sampling generates rare states exponentially with increasing N .

3.5 Summary: challenges for the application of Monte Carlo methods to chaotic systems

To achieve a MH algorithm that scales polynomially with N , the autocorrelation time needs to be low. A key ingredient for an efficient algorithm is a good proposal distribution $g(\mathbf{x}'|\mathbf{x})$ [61]. The ideal proposal distribution of a Metropolis-Hastings algorithm draws $\mathbf{x}'$ independently of $\mathbf{x}$ according to $\pi(\mathbf{x}')$, $g(\mathbf{x}'|\mathbf{x}) = \pi(\mathbf{x}')$. This is because (i) the acceptance in equation (23) is always one and (ii) each step of the random walk generates an independent sample $\mathbf{x}$, which implies that the error in equation (22) is minimal. The difficulty of sampling rare events in chaotic systems is that a useful $\pi(\mathbf{x})$ to sample them is a difficult function to sample from. For concreteness, consider the problem of sampling high escape times in the open tent map defined in Appendix A.3 – equation (A.5), whose $t_e(\mathbf{x})$ is represented in Figure A.1 – and consider the canonical sampling distribution, $\pi(\mathbf{x}) \propto \exp(-\beta t_e(\mathbf{x}))$. In this example, $\pi(\mathbf{x}) \propto \exp(-\beta 1)$ between $[1/a, 1 - 1/a]$ and so forth ($\exp(-\beta t_e)$) in subsequent intervals. The number of intervals increases as 2^{t_e} . Therefore, sampling from $\pi(\mathbf{x})$ would require enumerating every interval, sample one at random according to a correct distribution that depends

on β , and then sample a uniform point within that interval. While this can be done in simple maps such as the open tent map, this is unfeasible in a general system where the $t_e(\mathbf{x})$ dependency is unknown.

One approach is to consider $g(\mathbf{x}'|\mathbf{x})$ to be uniformly distributed over Γ . One could imagine that changing the sampling distribution $\pi(\mathbf{x})$ would decrease the variance of the estimator because this gives more preference to rarer states. However, this is not the case: changing the sampling distribution alone does not decrease the scaling of the variance of the estimator. This is because changing the sampling distribution (e.g. using a canonical ensemble) leads to an exponential increase of the autocorrelation time $T(E)$ with E , see Appendix B, making it as efficient as traditional uniform sampling. While here we focus on MH, the choice of the proposed state is crucial to Monte Carlo methods more generally, and thus addressing this problem for MH can be useful more generally, e.g. for the methods used in references [17,18,22].

In summary, Metropolis-Hastings is an excellent candidate to approach the numerical challenges found in the study of rare events in chaotic systems. Firstly, because it is grounded in strong mathematical results such as importance sampling theorem and asymptotic convergence of Markov processes. Secondly, because it is formulated with very little assumptions about the system, the observable of interest or the dynamics of the system, which gives enough freedom to adapt it to the specific aim (sampling or finding), observable, and system. Thirdly, because there seems to be no theoretical reason for the sampling to be exponential; Metropolis-Hastings is used to sample rare states in polynomial time in other problems of statistical physics. Finally, because the numerical problems found in chaotic systems can be re-written as problems where Metropolis-Hastings is suitable for. On the other hand, the optimal proposal of Metropolis-Hastings is unfeasible in chaotic systems, and without any extra information about the system, Metropolis-Hastings is as efficient as the traditional uniform sampling.

4 Proposal distribution

The problem we address in this section is: how to incorporate general properties of chaotic systems into the proposal distribution in such a way that the Metropolis-Hastings algorithm becomes efficient? We first set an aim for the proposal distribution and we then show how this aim can be achieved in the different problems involving deterministic chaotic system.

4.1 Aim of the proposal distribution

The goal of the proposal distribution $g(\mathbf{x}'|\mathbf{x})$ we construct here will be to bound the acceptance rate in equation (23) away from 0 and 1. Since $a(\mathbf{x}'|\mathbf{x})$ in equation (23) depends on $\mathbf{x}'$, it is essential to look at its expectation over

the proposal,

$$\mathbb{E}[a(\mathbf{x}'|\mathbf{x})|\mathbf{x}] \equiv \int_{\Gamma} a(\mathbf{x}'|\mathbf{x})g(\mathbf{x}'|\mathbf{x})d\mathbf{x}'. \quad (30)$$

Our goal is to construct a proposal distribution such that⁴

$$\mathbb{E}[a(\mathbf{x}'|\mathbf{x})|\mathbf{x}] = a_*. \quad (31)$$

The motivation to set this as the starting point (and cornerstone) of our method is that it avoids the two typical origins of high correlation times T . When the acceptance is low, T increases because the method remains stuck in the same state $\mathbf{x}$ for long times. High acceptance typically indicates that $\mathbf{x}'$ is too close to $\mathbf{x}$ (often $E(\mathbf{x}') = E(\mathbf{x})$), which implies that the simulation moves too slowly (in E and in Ω). The goal here is not that the acceptance is exactly a_* , but rather that it remains bounded from the extremes and that it does not strongly depend on N .

In general it is non-trivial to construct a proposal distribution that guarantees a constant acceptance. Thus, the next step is to approximate equation (31) by a simpler condition. Let us first notice that $\pi(\mathbf{x})$ only depends on $\mathbf{x}$ through $E_{\mathbf{x}} \equiv E(\mathbf{x})$, $\pi(\mathbf{x}) = \pi(E_{\mathbf{x}})$. Therefore, at the very least, the proposal $g(\mathbf{x}'|\mathbf{x})$ should guarantee that $\mathbf{x}'$ is generated from $\mathbf{x}$ in such a way that $\pi(E_{\mathbf{x}'})$ is neither too close (high acceptance) nor too far (low acceptance) from $\pi(E_{\mathbf{x}})$. Quantitatively, this can be written as

$$\mathbb{E}\left[\frac{\pi(E_{\mathbf{x}'})}{\pi(E_{\mathbf{x}})}|\mathbf{x}\right] = a \quad (32)$$

where $0 < a < 1$ is a constant. When the proposal is able to achieve a small variation of E , $\pi(E_{\mathbf{x}'})$ can be expanded in Taylor series around $E_{\mathbf{x}'} = E_{\mathbf{x}}$, which allows to write

$$\frac{\pi(E_{\mathbf{x}'})}{\pi(E_{\mathbf{x}})} = 1 + \frac{d \log \pi(E_{\mathbf{x}})}{dE}(E_{\mathbf{x}'} - E_{\mathbf{x}}). \quad (33)$$

Inserting equation (33) in equation (32), an average constant acceptance is thus achieved when.⁵

$$\mathbb{E}[E_{\mathbf{x}'} - E_{\mathbf{x}}|\mathbf{x}] = \frac{a - 1}{d \log \pi(E_{\mathbf{x}})/dE} \quad (34)$$

This equation, the main result of this section, is a condition that an efficient Metropolis-Hastings imposes to the proposal distribution in terms of the average difference in the observable E . This condition is non-trivial because it depends on the particular π , E , $\mathbf{F}$, and $\mathbf{x}$. For the

sampling distributions discussed in the previous section, we obtain:

Canonical ensemble

When $\pi(\mathbf{x}) = e^{-\beta E_{\mathbf{x}}}$, the condition in equation (34) is given by

$$\mathbb{E}[E_{\mathbf{x}'} - E_{\mathbf{x}}|\mathbf{x}] = \frac{1 - a}{\beta}. \quad (35)$$

That is, the higher the β , the closer the proposed $E_{\mathbf{x}'}$ has to be from $E_{\mathbf{x}}$.

Flat-histogram

When $\pi(\mathbf{x}) \propto 1/G(E_{\mathbf{x}})$, the condition in equation (34) is given by

$$\mathbb{E}[E_{\mathbf{x}'} - E_{\mathbf{x}}|\mathbf{x}] = \frac{1 - a}{\frac{d \log G}{dE}(E_{\mathbf{x}})}. \quad (36)$$

When $E_{\mathbf{x}}$ is close to the maximum of $G(E)$, the derivative of $\log G$ approaches 0 and $E_{\mathbf{x}'}$ can be arbitrary distant from $E_{\mathbf{x}}$. As $E_{\mathbf{x}}$ deviates from the maximum of $G(E)$, smaller changes are necessary to achieve a constant acceptance. Note that equation (34) is valid for the Metropolis-Hastings algorithm in general and should be of interest also in other contexts (e.g., arbitrarily large number of spins can be flipped close to the maximum of the density of states).

4.2 Propose correlated trajectories

The proposal distribution requires correlating the trajectory starting at $\mathbf{x}'$ with a trajectory starting at $\mathbf{x}$ such that equation (34) holds. Fulfilling this requirement requires the ability to control $\mathbb{E}[E_{\mathbf{x}'} - E_{\mathbf{x}}|\mathbf{x}]$, which requires a procedure to correlate the state $\mathbf{x}'$ with $\mathbf{x}$. The aim of this section is to introduce a quantification of the correlation of two trajectories of finite-time t_o that can be related with $\mathbb{E}[E_{\mathbf{x}'} - E_{\mathbf{x}}|\mathbf{x}]$, and present two different proposal distributions that propose a state $\mathbf{x}'$ on which this correlation is controlled.

The observables E introduced in Section 2.3 are all dependent not only of $\mathbf{x}$, but also of the full trajectory of length t_o starting $\mathbf{x}$.⁶ One natural way to quantify the similarity of two trajectories is the length the two trajectories remain within a distance Δ much smaller than a characteristic length of Γ . Formally, this can be quantified by $t_*(\mathbf{x}, \mathbf{x}') = \max\{t \leq t_o : |\mathbf{F}^t(\mathbf{x}) - \mathbf{F}^t(\mathbf{x}')| \leq \Delta\}$. Under this definition, $0 \leq t_*(\mathbf{x}, \mathbf{x}') \leq t_o$. When $\mathbf{x}' = \mathbf{x}$, $t_*(\mathbf{x}, \mathbf{x}') = t_o$ because the trajectories are the same; when $\mathbf{x}'$ is far from $\mathbf{x}$, $t_*(\mathbf{x}, \mathbf{x}') = 0$. However, there is another possibility for two trajectories starting at $\mathbf{x}'$ and $\mathbf{x}$ to be similar: when the two trajectories are similar apart from a shift in time, see right panel of Figure 3. One way to include both cases in our measure of similarity is to define $t_*(\mathbf{x}, \mathbf{x}')$ as

$$t_*(\mathbf{x}, \mathbf{x}') \equiv \max\{t \leq t_o - t_{\text{shift}} : |\mathbf{F}^t(\mathbf{x}') - \mathbf{F}^{t_{\text{shift}}+t}(\mathbf{x})| \leq \Delta\}. \quad (37)$$

⁶ The escape time t_e requires that $\mathbf{F}^t(\mathbf{x}) \notin \Lambda \forall t$; the FTLE λ_{t_o} is a sum of terms computed over the trajectory.

⁴ For example, reference [61] computed $a_* = 0.234$.

⁵ In statistical physics an heuristics often used in Metropolis-Hastings is $\Delta E \equiv E_{\mathbf{x}'} - E_{\mathbf{x}} \approx 1$, e.g. it is used in the derivation of equation (28), and has been used for example in spin systems (single spin flip) [59], ensemble of complex networks (single link exchange) [62], and proteins [63].

Fig. 3. Two trajectories of length $t_o = 6$ starting at $\mathbf{x}$ and $\mathbf{x}'$ can be correlated for five steps ($t_*(\mathbf{x}, \mathbf{x}') = 5$) by two different mechanisms: (left panel) they start close from each other and are indistinguishable within Δ up to time t_* ; (right panel) they start shifted (by $t_{\text{shift}} = 1$ here) from each other and are thus indistinguishable for $t_o - t_{\text{shift}} = t_* = 5$ steps.

For $t_{\text{shift}} = 0$, this recovers the case of two trajectories starting close to each other; $t_{\text{shift}} \neq 0$, it includes situations where a trajectory starts at $\mathbf{x}'$ close to $\mathbf{F}^{t_{\text{shift}}}(\mathbf{x})$. This definition is motivated by the concept of symbolic sequences [4]. The similarity of the trajectory starting from $\mathbf{x}'$ with the one starting from $\mathbf{x}$ can be quantified by the number of symbols that both trajectories share, which corresponds to the $t_*(\mathbf{x}, \mathbf{x}')$ in equation (37). The definition in equation (37) avoids the necessity of the existence of a phase-space partition, but, for the purposes of the argument below, the two trajectories share a sequence of t_* states that are close within Δ .

The average correlation between two states whose one is drawn according to a proposal distribution is here defined by

$$t_*(\mathbf{x}) \equiv \mathbb{E}[t_*(\mathbf{x}, \mathbf{x}')|\mathbf{x}] = \int_{\Gamma} g(\mathbf{x}'|\mathbf{x}) t_*(\mathbf{x}, \mathbf{x}') d\mathbf{x}'. \quad (38)$$

Notice that this quantity does not depend on the particular sampling distribution or algorithm; it is a function of the proposal distribution and the state $\mathbf{x}$. The goal of the next sub-sections is to construct proposal distributions that guarantee a given average correlation $t_*(\mathbf{x})$. We will show, for example, that we can enforce an average correlation $t_*(\mathbf{x})$ if we use a normal distribution centered around $\mathbf{x}$ with a specific standard deviation as our proposal distribution.

4.2.1 Shift proposal

One proposal that guarantees that trajectories are correlated by t_* is the shift proposal, originally introduced in reference [13] in the context of sampling paths of chemical reactions, and has also been used in reference [23]. It consists in proposing a state $\mathbf{x}'$ that is a forward or backward iteration of $\mathbf{x}$, $\mathbf{x}' = \mathbf{F}^{t_{\text{shift}}}(\mathbf{x})$, where t_{shift} is a free parameter. The relation between $t_{\text{shift}} = t_{\text{shift}}(\mathbf{x})$ and $t_* = t_*(\mathbf{x})$

is that a shift of $\pm t_{\text{shift}}$ guarantees that $t_o - t_{\text{shift}}$ elements of the original trajectory are preserved. Therefore, this proposal guarantees that t_* elements are preserved when $|t_{\text{shift}}| = t_o - t_*$, see right panel of Figure 3. Because detailed balance has to be enforceable, the proposal must contain backward and forward shifts. A proposal that automatically fulfills detailed balance is one on which the backward and forward shifts are equally likely:

$$g(\mathbf{x}'|\mathbf{x}) = \frac{1}{2} \delta(\mathbf{x}' - \mathbf{F}^{t_{\text{shift}}}(\mathbf{x})) + \frac{1}{2} \delta(\mathbf{x}' - \mathbf{F}^{-t_{\text{shift}}}(\mathbf{x})), \quad (39)$$

with $t_{\text{shift}} = t_{\text{shift}}(\mathbf{x}) = t_o - t_*(\mathbf{x})$. Given the target average correlation $t_*(\mathbf{x})$, this proposal can be implemented as follows: generate a random number $r \in [0, 1]$; if $r < 0.5$, make $\mathbf{x}' = \mathbf{F}^{t_o - t_*(\mathbf{x})}(\mathbf{x})$, else, make $\mathbf{x}' = \mathbf{F}^{t_*(\mathbf{x}) - t_o}(\mathbf{x})$.

This proposal unfortunately has some disadvantages: (i) a priori there is no guarantee that $\mathbf{F}^{t_{\text{shift}}}(\mathbf{x}) \in \Gamma$. It is applicable when $\Gamma = \Omega$, which, e.g. is not the case in open systems; (ii) it requires the map to be invertible; (iii) the proposal can only propose states that are forward or backward iterations of $\mathbf{x}$. Consequently, for the random walk to be ergodic in the phase-space, the map itself must be ergodic, (iv) this proposal diffuses without drift on a trajectory passing through $\mathbf{x}$, by shifting the starting point forward or backward. Thus, it will always sample fewer states than a time average of a trajectory starting at $\mathbf{x}$. On the other hand, the main advantage of this proposal is that it performs non-local jumps in the phase-space. That is, it allows to jump from a region of the phase-space to another region while still maintaining $\mathbf{x}'$ correlated with $\mathbf{x}$. As shown below, in combination with other proposal, this proposal is useful to reduce correlations stemmed from local jumps.

4.2.2 Neighbourhood proposal

Another strategy to construct a proposal on which on average the states are correlated by $t_*(\mathbf{x})$ is to perturb $\mathbf{x}$ by a finite amount δ , $\mathbf{x}' = \mathbf{x} + \delta$, characterized by a direction $\hat{\delta}$ and a norm δ , $\delta \equiv \hat{\delta}\delta$. A common case is when the probability distribution is separated in two independent terms [7]:

$$P(\delta|\mathbf{x}) = P(\hat{\delta}|\mathbf{x})P(\delta|\mathbf{x}) \quad (40)$$

and that $P(\hat{\delta}|\mathbf{x})$ is uniformly distributed in the D directions and $P(\delta|\mathbf{x})$ has zero mean (i.e. an isotropic proposal). Here we restrict the analysis to this situation, and we also assume that $P(\delta|\mathbf{x})$ is characterized by a well defined scale, e.g. it is an half-normal distribution⁷ with mean $\delta_x(\mathbf{x})$:

$$P(\delta|\mathbf{x}) = \frac{\sqrt{2}}{\sqrt{\pi\sigma^2}} e^{-\frac{\pi\delta^2}{4\delta_x(\mathbf{x})^2}} \text{ for } \delta > 0. \quad (41)$$

⁷ An exponential distribution would be equally acceptable and would not change the main conclusions.

This choice makes the ratio $g(\mathbf{x}|\mathbf{x}')/g(\mathbf{x}'|\mathbf{x})$ in equation (23) to be

$$\frac{g(\mathbf{x}|\mathbf{x}')}{g(\mathbf{x}'|\mathbf{x})} = \frac{\delta_x(\mathbf{x})}{\delta_x(\mathbf{x}')} \exp \left[-\frac{\pi|\mathbf{x}' - \mathbf{x}|^2}{4\delta_x(\mathbf{x})^2} \left(1 - \frac{\delta_x(\mathbf{x})^2}{\delta_x(\mathbf{x}')^2} \right) \right]. \quad (42)$$

The main motivation for this choice is that the proposal distribution is described by a single function, $\delta_x(\mathbf{x})$, that quantifies the distance $\mathbf{x}' - \mathbf{x}$, $\mathbb{E}[|\mathbf{x}' - \mathbf{x}||\mathbf{x}] = \delta_x(\mathbf{x})$.

The goal is now to relate $\delta_x(\mathbf{x})$ with $t_*(\mathbf{x})$. Let us start to describe two important limits: in the limit $\delta_x(\mathbf{x}) \rightarrow 0$, the states are the same and therefore $\lim_{\delta_x(\mathbf{x}) \rightarrow 0} t_*(\mathbf{x}) = t_o$. In the limit $\delta_x(\mathbf{x}) \rightarrow |\Gamma|$, the proposal is approximately equal to draw $\mathbf{x}'$ uniformly from Γ , and $\mathbf{x}'$ is independent of $\mathbf{x}$ and $t_*(\mathbf{x}) = 0$. To preserve a correlation of $t_*(\mathbf{x})$, it is necessary that $\delta_x(\mathbf{x})$ is such that the two trajectories starting at $\mathbf{x}$ and $\mathbf{x}'$ are close together up to a time $t_*(\mathbf{x})$, see left panel of Figure 3. Because the system is chaotic, for small $\delta_x(\mathbf{x})$, two trajectories diverge exponentially in time according to equation (4), and, in particular, their maximal distance is given by equation (7). Therefore, to guarantee that two trajectories are distanced at most by Δ after a time $t_*(\mathbf{x})$, $\delta_x(\mathbf{x})$ must be given by

$$\delta_x(\mathbf{x}) = \Delta e^{-\lambda_{t_*}(\mathbf{x})t_*(\mathbf{x})}. \quad (43)$$

This equation relates the parameter of the proposal distribution, $\delta_x(\mathbf{x})$, with the average correlation $t_*(\mathbf{x})$ of the two states $\mathbf{x}$ and $\mathbf{x}'$.

The neighborhood proposal derived above is closely related to a proposal described in reference [23] as “precision shooting”. Precision shooting constructs a trajectory $\{\mathbf{x}'_i\}$ with $\mathbf{x}'_0 \equiv \mathbf{x}' = \mathbf{x} + \delta\hat{\mathbf{e}}$ (where δ is a free parameter) that, within the numerical precision of a computer, is indistinguishable (in the system considered) from a trajectory starting at $\mathbf{x}'$ with δ small. The trajectory $\{\mathbf{x}'_i\}$ shadows a true trajectory starting at $\mathbf{x}'$, in the same spirit as the algorithm used in reference [64] to construct a pseudo-trajectory. Thus, precision shooting can be interpreted as the neighborhood proposal, equation (43), with t_* a free parameter (related to δ via Eq. (43)), that assumes shadowing theorem to simplify the construction $\mathbf{x}'$. Reference [23] discusses how the acceptance rate depends on δ , suggesting that the acceptance rate increases with decreasing δ (Fig. 9 of the reference). In light of the discussion in Section 4.2, this result is interpreted as follows: as δ decreases, $\mathbf{x}'$ becomes more correlated with $\mathbf{x}$ (since t_* is related with δ by Eq. (43)), and therefore the acceptance is expected to increase, as indicated in Figure 9 of reference [23]. This discussion is unfortunately insufficient to us because it does not allow to derive δ (or t_*) that fulfills the condition in equation (34). The crucial advantage of equation (43) is that it allows to relate $\mathbb{E}[E(\mathbf{x}') - E(\mathbf{x})|\mathbf{x}]$ (in Eq. (34)) with $\mathbf{x}' - \mathbf{x}$ (in Eq. (43)) through t_* (in Eq. (37)). This is the goal of the next section.

4.3 Guarantee local proposals

Now that we derived proposal distributions that enforce an average correlation $t_*(\mathbf{x})$, the next (and final)

step is to obtain a relationship between $t_*(\mathbf{x})$ and $\mathbb{E}[E(\mathbf{x}') - E(\mathbf{x})|\mathbf{x}]$. Because the computation of $t_*(\mathbf{x})$ depends on the particular observable E , a different derivation is presented for two observables, t_e and λ_{t_o} . Given the limitations of the shift proposal described above, the argumentation below is limited to neighborhood proposals (an equivalent argumentation can be made to the shift proposal).

4.3.1 FTLE in closed systems

As introduced in Section 2.1, the observable in this case is given by $E(\mathbf{x}) = t_o\lambda_{t_o}(\mathbf{x})$, where t_o is the finite time. The aim in this case is thus to write $\mathbb{E}[t_o\lambda_{t_o}(\mathbf{x}') - t_o\lambda_{t_o}(\mathbf{x})|\mathbf{x}]$ as a function of $t_*(\mathbf{x})$. The finite-time Lyapunov exponent considered in equation (8) is a sum of t_o terms⁸ and thus $t_o\lambda_{t_o}(\mathbf{x})$ can be written as the sum of the FTLE up to time t_* and the FTLE from t_* up to t_o ,

$$t_o\lambda_{t_o}(\mathbf{x}) = t_*\lambda_{t_*}(\mathbf{x}) + (t_o - t_*)\lambda_{t_o-t_*}(\mathbf{x}_{t_*}), \quad (44)$$

where $\mathbf{x}_{t_*} \equiv \mathbf{F}^{t_*}(\mathbf{x})$. Likewise for the trajectory starting from $\mathbf{x}'$,

$$t_o\lambda_{t_o}(\mathbf{x}') = t_*\lambda_{t_*}(\mathbf{x}') + (t_o - t_*)\lambda_{t_o-t_*}(\mathbf{x}'_{t_*}). \quad (45)$$

Because $\mathbf{x}'$ is proposed according to equation (43), by construction, the first t_* states of the trajectory starting at $\mathbf{x}'$ are close (within Δ) to the states of the trajectory starting at $\mathbf{x}$ up to $t_*(\mathbf{x})$. Therefore, we can approximate that the respective Lyapunovs up to time t_* are equal,

$$\mathbb{E}[\lambda_{t_*}(\mathbf{x}')|\mathbf{x}] \approx \lambda_{t_*}(\mathbf{x}). \quad (46)$$

Subtracting equation (44) from equation (45) and using equation (46) gives

$$\begin{aligned} & \mathbb{E}[t_o\lambda_{t_o}(\mathbf{x}') - t_o\lambda_{t_o}(\mathbf{x})|\mathbf{x}] \\ &= (t_o - t_*) (\mathbb{E}[\lambda_{t_o-t_*}(\mathbf{x}'_{t_*})|\mathbf{x}] - \lambda_{t_o-t_*}(\mathbf{x}_{t_*})). \end{aligned} \quad (47)$$

The left side of this equation is the same as in equation (34) and thus the aim now is to write the right side as a function of properties of the system. Let us focus on the calculation of $\mathbb{E}[\lambda_{t_o-t_*}(\mathbf{x}'_{t_*})|\mathbf{x}]$ first. By construction, $\mathbf{x}'$ is generated such that $|\mathbf{x}'_{t_*} - \mathbf{x}_{t_*}| \approx \Delta$. Because the system is chaotic, one can approximate that $\mathbf{x}'_{t_*}$ is sufficiently separated from $\mathbf{x}_{t_*}$ such that $\lambda_{t_o-t_*}(\mathbf{x}'_{t_*})$ is independent of $\mathbf{x}$. Under this approximation, $\mathbf{x}'_{t_*}$ is essentially a random state from the phase-space, and thus $\lambda_{t_o-t_*}(\mathbf{x}'_{t_*})$ will be a drawn from the distribution of FTLE at time $t_o - t_*$. Denoting the mean of this distribution by λ_{L,t_o-t_*} , we get

$$\mathbb{E}[\lambda_{t_o-t_*}(\mathbf{x}'_{t_*})|\mathbf{x}] = \lambda_{L,t_o-t_*}. \quad (48)$$

Replacing equation (48) in equation (47) gives

$$\begin{aligned} & \mathbb{E}[t_o\lambda_{t_o}(\mathbf{x}') - t_o\lambda_{t_o}(\mathbf{x})|\mathbf{x}] \\ &= (t_o - t_*) (\lambda_{L,t_o-t_*} - \lambda_{t_o-t_*}(\mathbf{x}_{t_*})). \end{aligned} \quad (49)$$

⁸ This equation is strictly valid for 1D, but the idea is that we split the FTLE between a component up to t_* and the remaining component.

This equation relates the expected change in the observable with properties of the system ($\lambda_{L,t_o-t_\star}$), of the trajectory $\mathbf{x}$, $\lambda_{t_o-t_\star}(\mathbf{x}_{t_\star})$, and $t_\star(\mathbf{x})$ and it can thus be used in the energy condition we obtained earlier, equation (34). Replacing the left side of equation (34) by the expectation in equation (49) and solving to $t_\star(\mathbf{x})$ gives

$$t_\star(\mathbf{x}) = t_o - \frac{a-1}{d \log \pi(\lambda_t(\mathbf{x}))/d\lambda_t} \frac{1}{\lambda_{L,t_o-t_\star} - \lambda_{t_o-t_\star}(\mathbf{x}_{t_\star})}. \quad (50)$$

We can further simplify this relation with two approximations: (a) for large $t_o - t_\star$, the mean FTLE at time $t_o - t_\star$, $\lambda_{L,t_o-t_\star}$, is approximately the Lyapunov exponent of the system, λ_L ,

$$\lambda_{L,t_o-t_\star} \approx \lambda_L, \quad (51)$$

(b) because trajectories of chaotic systems are short-correlated in time, the FTLE of the trajectory $\mathbf{x}$ up to $t_\star$ will be approximately equal to the FTLE of the trajectory starting at $\mathbf{x}_{t_\star}$. Thus,

$$\lambda_{t_o-t_\star}(\mathbf{x}_{t_\star}) \approx \lambda_{t_o}(\mathbf{x}). \quad (52)$$

Using equations (51) and (52), $t_\star(\mathbf{x})$ in equation (50) can be written as

$$t_\star(\mathbf{x}) = t_o - \frac{a-1}{d \log \pi(\lambda_t(\mathbf{x}))/d\lambda_t} \frac{1}{\lambda_L - \lambda_{t_o}(\mathbf{x})}. \quad (53)$$

To enforce that $t_\star \in [0, t_o]$, we use

$$t_\star(\mathbf{x}) = \max \left\{ 0, t_o - \left\lfloor \frac{a-1}{d \log \pi(\lambda_t(\mathbf{x}))/d\lambda_t} \frac{1}{\lambda_L - \lambda_{t_o}(\mathbf{x})} \right\rfloor \right\}, \quad (54)$$

which is the main result of this section. This equation provides an expression to $t_\star(\mathbf{x})$ that can be inserted in the parameter of the proposal distribution, equation (43), that under the approximations used above achieves a constant acceptance rate.⁹

The derivation of equation (54) can be generalized to other observables which, as λ_{t_o} , can be written as an average over the trajectory: consider

$$e_{t_o}(\mathbf{x}) \equiv \frac{1}{t_o} \sum_{i=1}^{t_o} f(\mathbf{x}_i) = E_{t_o}(\mathbf{x})/t_o \quad (55)$$

where $f(\mathbf{x}_i)$ is an arbitrary function of the phase-space (the logarithm of the derivative of the map corresponds to $e_{t_o}(\mathbf{x}) = \lambda_{t_o}(\mathbf{x})$, $E_{t_o}(\mathbf{x}) = \lambda_{t_o}(\mathbf{x})t_o$). Replacing this quantity in the derivation of equation (54) mutatis mutandis and without using the approximation in equation (52), one obtains

$$t_\star(\mathbf{x}) = \max \left\{ 0, t_o - \left\lfloor \frac{a-1}{d \log \pi(E_{t_o}(\mathbf{x}))/dE} \frac{1}{E_L - E_{t_o-t_\star}(\mathbf{x})} \right\rfloor \right\}, \quad (56)$$

⁹ In reference [25] we used $t_\star = t_o - 1$, different from equation (54). The derivation of equation (54) uses stronger approximations than the derivation of $t_\star = t_o - 1$. See argumentation after equation (14) of reference [25].

Fig. 4. Illustration of the core idea of our proposal for the case of the escape time: a trajectory $\mathbf{x}'$ starting close to $\mathbf{x}$ and that remains similar to $\mathbf{x}$ up to a time $t_\star = 2$ will likely leave at a time $t_e(\mathbf{x}') > t_e(\mathbf{x}) = 3$. The black (red) lines represent the iteration of the state $\mathbf{x}(\mathbf{x}')$ until it leaves at time $t_e(\mathbf{x}) = 3$ ($t_e(\mathbf{x}') > 3$) by entering the shaded exit region Λ . The squares represent zooms of $\exp(\lambda_t(\mathbf{x})t)$ to the phase-space and thus correlated trajectories up to $t_\star$ correspond to select a specific zoom. Our proposal is constructed in such a way that (1) the distance between $\mathbf{x}$ and $\mathbf{x}'$ at time $t_\star$ is on average Δ and (2) $t_\star$ is chosen to be neither too large nor too small.

where E_L is approximately the mean of the distribution $P(E_{t_o})$ (using the approximation in Eq. (48)). This generalizes equation (54) for an arbitrary average over trajectories of size t_o , $e_{t_o}(\mathbf{x})$, and it should be useful to sample rare states in respect to observables correspondent to expected values over trajectories.

4.3.2 Escape time in strongly chaotic open systems

As introduced before, in strongly chaotic open systems $E(\mathbf{x}) = t_e(\mathbf{x})$ and $P(t_e) \sim \exp(-\kappa t_e)$. The aim here is to compute $t_\star(\mathbf{x})$ that fulfills equation (34) by taking into account that $\mathbf{x}$ is given and $\mathbf{x}' = \mathbf{x} + \delta\delta_x(\mathbf{x})$ with $\delta_x(\mathbf{x})$ given by equation (43). In this case the trajectory's length is not a constant but instead it is given by the escape time $t_e(\mathbf{x})$.

A trajectory starting at $\mathbf{x}'$ proposed according to equation (43) fulfills $|\mathbf{x}'_{t_\star} - \mathbf{x}_{t_\star}| \leq \Delta$. Therefore, up to $t_\star$, the two trajectories are indistinguishable. Under this assumption, and because from the definition of $t_\star$ in equation (37), $t_e(\mathbf{x}') = t_\star(\mathbf{x}) + t_e(\mathbf{x}'_{t_\star})$ and therefore

$$\mathbb{E}[t_e(\mathbf{x}')|\mathbf{x}] = t_\star(\mathbf{x}) + \mathbb{E}[t_e(\mathbf{x}'_{t_\star})|\mathbf{x}], \quad (57)$$

where $t_e(\mathbf{x}'_{t_\star})$ is the escape time of the state $\mathbf{x}'$ iterated $t_\star$ times and $\mathbb{E}[t_e(\mathbf{x}'_{t_\star})|\mathbf{x}] = \int_{\Omega} g(\mathbf{x}'|\mathbf{x})t_e(\mathbf{x}'_{t_\star})d\mathbf{x}'$ is the expected $t_e(\mathbf{x}'_{t_\star})$ over a neighborhood of $\mathbf{x}$ of size $\delta_x(\mathbf{x})$

(given by Eq. (43)). The idea behind this equation is represented in Figure 4. The proposal density $g(\mathbf{x}'|\mathbf{x})$ at $\mathbf{x}$ is such that, at t_* , the two trajectories distance themselves on average by Δ . After t_* , the trajectory starting at $\mathbf{x}'_{t_*}$ will approximately be independent of $\mathbf{x}$ and thus $t_e(\mathbf{x}') = t_* + t_e(\mathbf{x}'_{t_*})$. Moreover, proposing according to equation (43) is equivalent to “zoom” the landscape of $t_e(\mathbf{x})$ around $\mathbf{x}$ with a scale correspondent to t_* ’th iteration of the construction of the landscape. Given the self-similarity of the landscape (see e.g. Fig. A.2), under this zoom, the landscape of $t_e(\mathbf{x}_{t_*})$ is equal to the landscape of $t_e(\mathbf{x})$ and therefore $\mathbb{E}[t_e(\mathbf{x}'_{t_*})|\mathbf{x}]$ should be a constant independent of t_e . Moreover, because $\mathbf{x}'_{t_*}$ is approximately independent of $\mathbf{x}$, $\mathbb{E}[t_e(\mathbf{x}'_{t_*})|\mathbf{x}]$ is the just average escape time of an independent state $\mathbf{x}'_{t_*}$, which is the average of $P(t_e)$ and is given approximately by $1/\kappa$. Thus, $\mathbb{E}[t_e(\mathbf{x}'_{t_*})|\mathbf{x}]$ is the average escape time of an independent state $\mathbf{x}'_{t_*}$, which is the average of $P(t_e)$ and is given by $1/\kappa$:

$$\mathbb{E}[t_e(\mathbf{x}'_{t_*})|\mathbf{x}] = 1/\kappa. \quad (58)$$

It is this result that incorporates the self-similarity of the escape time function: the value of t_* chooses the particular zoom of the landscape (Fig. 4), and this equation assumes that, as long as the zoom is proportional to $\exp(\lambda t_*)$, the average escape time of the phase-space of the zoomed region is still $1/\kappa$. Replacing equation (58) in equation (57) gives

$$\mathbb{E}[t_e(\mathbf{x}')|\mathbf{x}] = t_*(\mathbf{x}) + 1/\kappa, \quad (59)$$

and subtracting $t_e(\mathbf{x})$ on both sides of equation (59) gives

$$\mathbb{E}[t_e(\mathbf{x}') - t_e(\mathbf{x})|\mathbf{x}] = t_*(\mathbf{x}) + \frac{1}{\kappa} - t_e(\mathbf{x}). \quad (60)$$

The left side of this equation is the left side of the condition of constant acceptance rate, equation (34). Equating both left sides and solving for $t_*(\mathbf{x})$ gives

$$t_*(\mathbf{x}) = t_e(\mathbf{x}) - \frac{1}{\kappa} - \frac{a-1}{d \log \pi(t_e)/dt_e}, \quad (61)$$

which is analogous to equation (54) and is the central result of this section. Together with equation (43), it is the proposal distribution we had the goal of constructing.

4.4 Summary: how to propose

The argument in this section can be summarized as follows:

- Metropolis-Hastings requires a proposal that guarantees a specific variation of E , which we approximate by equation (34);
- the proposal with a scale given by equation (43) guarantees that on average the trajectory starting at $\mathbf{x}'$ stays close to the trajectory starting at $\mathbf{x}$ up to $t_*(\mathbf{x})$;
- the variation of E is related to $t_*(\mathbf{x})$ via equations (49) and (60).

These three steps led to analytical formulas for $t_*(\mathbf{x})$, equations (61) and (54), that make the proposal satisfy equation (34). Together with equation (43), they define proposal distributions required for an average constant acceptance.

There are 3 points that deserve to be noted: the first point is that the formulas we obtained for $t_*(\mathbf{x})$ are also valid for the problem of *finding rare states* discussed, e.g. in references [11,12,16,18,65]. Specifically, these formulas were derived to guarantee equation (34), which dictates how different $\mathbb{E}[E(\mathbf{x}')|\mathbf{x}]$ has to be from $E(\mathbf{x})$ to guarantee a constant acceptance. This condition is stronger than the condition required for an algorithm to find minima or maxima of E , which requires only proposing states $\mathbf{x}'$ such that $\mathbb{E}[E(\mathbf{x}')|\mathbf{x}] > E(\mathbf{x})$ (or vice-versa for minimizing E). This is independent of the particular minimization algorithm (e.g. stimulated annealing, step descent, stagger and dagger in open systems) because it only discusses which new state $\mathbf{x}'$ should be *tried*, given the current state $\mathbf{x}$. The second point is that equations (61) and (54) reduce the proposal to a uniform distribution when the sampling distribution is the uniform distribution: $d \log \pi(E)/dE(E) \rightarrow 0$ implies $\delta_x(\mathbf{x}) \rightarrow \infty$. The third point is that different proposals in the literature, precision shooting [23,25], exponential proposal distribution [11,18], and the one in reference [24], can be obtained from the proposals derived in this section through different approximations, as shown in Appendix C.

Overall, this section described a framework to add information of chaotic systems, in this case the self-similar properties of the landscape, the exponential divergence of trajectories, and the exponential decay of correlations, to construct an efficient Metropolis-Hastings algorithm to sample them. A simplified description of the algorithm is given in Appendix D and an open-source implementation in reference [66]. The next section is devoted to test the assumptions used here on each of the problems, escape time and FTLE, and confirm the practical usefulness of the framework.

5 Numerical tests

The previous section concluded with a set of formulas – the proposal in equation (43) combined with the formulas for $t_*(\mathbf{x})$ equation (54) or equation (61) – for proposing states $\mathbf{x}'$ that are expected to guarantee the desired acceptance rate (bounded from 0 and 1). In this section we test some of the approximations made in the derivation of $t_*(\mathbf{x})$ (in a simple system) and we analyze the efficiency of the algorithm (confirming the polynomial scaling) in the computation of the FTLE in closed system – introduced in Section 2.1 – and of the escape time in open systems – introduced in Section 2.2. The tests are performed in the skewed tent map and open skewed tent map, and the efficiency is tested in numerous maps (chain of couple Hénon maps, Standard map, Logistic map). All these systems are introduced in detail in Appendix A.

Fig. 5. The finite-time Lyapunov exponent of the first half of the trajectory is independent from the one of the second half of the trajectory. The graph was obtained by sampling 10^5 random initial conditions $\mathbf{x}_i$ and compute $(\lambda_t(\mathbf{x}_i), \lambda_t(\mathbf{F}^t(\mathbf{x}_i))) = (\text{first half, second half})$ of a $2t = 16$ steps trajectory. The y axis represents the mean (full black) ± 2 standard deviations (full blue) of $\lambda_t(\mathbf{F}^t(\mathbf{x}_i))$ conditioned to a given λ_t . The dashed lines represent the same mean and standard deviation, but over all points (without conditioning). When $\lambda_t(\mathbf{F}^t(\mathbf{x}_i))$ is independent of $\lambda_t(\mathbf{x}_i)$, the dashed and full lines are the same within fluctuations, as observed. A Kolmogorov-Smirnov test comparing the un-conditioned and conditioned distributions gives a p -value higher than 0.001 (hypothesis that they are independent is not rejected).

5.1 Finite-time Lyapunov exponent

The first approximation made in the derivation of equation (54) is that when $\delta_x \equiv |\mathbf{x}' - \mathbf{x}|$ is drawn from a half-normal distribution with scale parameter $\delta_x(\mathbf{x})$ given by equation (43), $\mathbf{F}^{t_*}(\mathbf{x}')$ is sufficiently close from $\mathbf{F}^{t_*}(\mathbf{x})$ such that $\mathbb{E}[\lambda_{t_*}(\mathbf{x}')|\mathbf{x}] = \lambda_{t_*}(\mathbf{x})$, equation (46), holds. We test this numerically by randomly drawing 10^5 states $\mathbf{x}_i$ in the tent map with $a = 3$ (see Appendix A), and, for each, propose a state $\mathbf{x}'_i = \mathbf{x}_i + \delta(\mathbf{x})$ according to equation (41) with $\delta_x(\mathbf{x})$ given by our equation (43). From the 10^5 pairs of states $(\mathbf{x}_i, \mathbf{x}'_i)$, we estimate the difference of the observables: $E' \equiv \mathbb{E}[t\lambda_t(\mathbf{x}'_i)|\mathbf{x}]$ and $E \equiv t\lambda_t(\mathbf{x}_i)$. Our expectation is that for a fixed $t_* = t$, $E' - E$ should be much smaller than E (the relevant scale in Eq. (47)). We numerically obtain that within the 99% quantile, $E' - E \approx 1$ independently of E and t . This value is much smaller than $E \in [0.6t, 1.1t]$, specially since we are interested in large t . This strongly supports the approximation we make in equation (46).

The second approximation made in the derivation of equation (54) is that there is no dependence between $\lambda_{t_*}(\mathbf{x})$ and $\lambda_{t_o-t_*}(\mathbf{F}^{t_*}(\mathbf{x}))$, equation (48). In other words, that the FTLE of the trajectory starting at $\mathbf{F}^{t_*}(\mathbf{x})$ and ending at t_o is indistinguishable from the one drawn from the distribution of FTLE with finite-time $t_o - t_*$, $P(\lambda_{t_o-t_*})$. This approximation was numerically tested by drawing points $\mathbf{x}_i$, computing the pairs $\lambda_t(\mathbf{x}_i), \lambda_t(\mathbf{F}^t(\mathbf{x}_i))$ (i.e. $2t_* = t_o = 2t$), and testing whether the conditional probability of $\lambda_t(\mathbf{F}^t(\mathbf{x}_i))$ equals the (unconditional) probability of $\lambda_t(\mathbf{x}_i)$. The results in Figure 5 confirm the equality of these probabilities.

Fig. 6. The $t_*(\mathbf{x})$ given by equation (54) guarantees a bounded acceptance ratio independently of t_o . The graph represents the average acceptance as a function of $\lambda_{t_o}(\mathbf{x})$ obtained from uniformly sample of 10^6 pairs of states $(\mathbf{x}, \mathbf{x}'(\mathbf{x}))$ (see text for details) in the tent map with $a = 3$, for different finite-times t_o and sampling distributions. Top panel: $\pi(\mathbf{x})$ is the canonical ensemble, equation (24), with $\beta = 1$. Bottom panel: $\pi(\mathbf{x})$ is the flat-histogram ensemble, equation (26).

We finally test whether a proposal using $t_*(\mathbf{x})$ given by equation (54) guarantees a constant acceptance rate, the original motivation for our calculation. The test consisted in sampling 10^6 states according to the following procedure: (1) uniformly draw a state $\mathbf{x}_i$ and compute $\lambda_i \equiv \lambda_{t_o}(\mathbf{x}_i)$; (2) generate a state $\mathbf{x}'_i$ according to the proposal distribution equation (41) with $\delta_x(\mathbf{x})$ given by equation (43) and t_* given by equation (54) ($\delta_0 = 1$), and compute $\lambda'_i \equiv \lambda_{t_o}(\mathbf{x}'_i)$; (3) store $G_i \equiv g(\mathbf{x}|\mathbf{x}')/g(\mathbf{x}'|\mathbf{x})$ computed from equation (42): δ_x is given by equation (43), and $|\mathbf{x}' - \mathbf{x}|$ is given by storing $\delta_i = |\mathbf{x}'_i - \mathbf{x}_i|$. The ratio of the target distribution is given by $r_i \equiv \pi(E')/\pi(E) = \exp(-\beta t_o(\lambda'_i - \lambda_i))$ for the canonical ensemble and $r_i = G(E)/G(E')$, where $G(E) = G(t_o\lambda_i)$, is given by equation (A.3) for the flat-histogram. The numerically estimated acceptance ratio, $A(\lambda_t) \equiv \langle \min(1, r_i \times G_i) \rangle$ is shown in Figure 6. It is not independent of λ_{t_o} – the expected outcome based on the assumption of constant acceptance used in our derivations – but it is bounded from 0, 1 for increasing $N = t_o$ – the original requirement for an efficient simulation set in Section 4.1. In the canonical ensemble, there is linear dependency of Π with λ_{t_o} ,

Fig. 7. Average round-trip of a flat-histogram in the tent map with a mixed proposal is polynomial, as opposed to uniform sampling, shift proposal, and neighborhood proposal. The simulation was made using a flat-histogram simulation on the tent map with $\alpha = 3$, where $G(\lambda_t)$ is given from equation (A.3), and where the round-trip time (efficiency) was defined as going from $\lambda_{\min}$ to $\lambda_{\max}$ and return. The dashed black line represents $1/P(\lambda_{\max})$, the expected number of samples required in uniform sampling; the bottom full line is proportional to t_o^2 , the upper line is proportional to t_o^3 . The shift proposal used $t_{\text{shift}} = 1$ and for the backward iteration in time it used one (randomly chosen) of the two pre-images of the state. In the mixed proposal the shift and neighborhood proposals were chosen with probability $1/2$.

and in the flat-histogram ensemble, the ratio is 0.8 in the maximum of $P(\lambda_{t_o})$, and decays to about 0.1 on the tails.

5.1.1 Efficiency of the flat-histogram ensemble

The success in achieving a bounded acceptance independent of N , as the one confirmed in the previous section for the tent map, does not guarantee the efficiency of the method (see Sect. 3.4). Here we test the efficiency of flat-histogram computations of $P(\lambda_t)$ in the tent map that use the neighborhood proposal, the shift proposal, and both (mixed proposal). The results shown in Figure 7 suggest that only when both proposals are used the efficiency scales polynomially with $N = t_o$.

This result can be understood looking at the landscape of λ_t , as illustrated in Figure 8. Imagine a flat-histogram simulation on this system, for $t_o = 4$ (black curve in the figure), and analyze what happens to it in terms of a round-trip. Lets suppose that the simulation was recently at the minimum λ_t (0.41) and that the next round-trip is made by going to the maximum λ_t (1.09) and return back. Lets further suppose that the simulation eventually got to a state with $\lambda_t \approx 0.92$. Because $\pi(\mathbf{x}) = \pi(\lambda_t(\mathbf{x}))$, every state at that λ_t is equiprobable. Therefore, the state can be at any plateau (of the 4, see Fig. 8), proportionally to their plateau-size. However, not every plateau contains, on its neighborhood, a neighbor plateau with higher λ_t , for example, the plateau around 0.3. Therefore, a local proposal would never be able to reach a higher plateau from a state on such a plateau. First, it would need to go backward, reach the maximum of $P(\lambda_t)$ (around 0.69), where

Fig. 8. The finite-time Lyapunov exponent $\lambda_{t_o}(x)$ in the tent map, equation (A.2), already contains an increasing number of local minima and maxima with increasing t_o , a crucial feature present in Figure 1. In both cases, the maximum $\lambda_{t_o}(x)$ is $\lambda_{\max} \approx 1.09$, and the minimum is $\lambda_{\min} \approx 0.4$.

the proposal proposes any other state, and then try to find another path towards a higher λ_t . This would already be the problem if the simulation would be a canonical ensemble with a β favoring higher λ_t 's, since it would require decreasing λ_t by an amount $\Delta\lambda$, and this happens with a probability that decays exponentially with $\Delta\lambda$. This is solved by using a flat-histogram ensemble. However, the crucial challenge here is that as t increases, the number of local maxima also increases, but the number of maxima connected with the global maximum is constant: the red curve, with $t = 6$, contains now 11 plateaus for $\lambda_t \approx 0.87$, but only 1 is locally connected to the maximum λ_t , the one around 0. This means that, as t increases, it becomes more difficult to perform a round-trip: not only because the expected time to diffuse increases (see Eq. (28)), but also because there are more times where the simulation diffuses forward and backward until it reaches the global maxima. The shift proposal alleviates this problem by allowing non-local proposals in the phase-space. A shift proposes a state $\mathbf{x}'$ on a non-neighborhood of $\mathbf{x}$, which improves the probability of reaching higher or lower λ 's, which explains why a combined proposal has such a low round-trip time in this system.

Finally, we confirm that a polynomial efficiency is obtained in the computation of $P(\lambda_{t_o})$ more generally. We performed flat-histogram simulations, using the Wang-Landau algorithm to estimate $P(\lambda_{t_o})$ on different chaotic systems: the tent map, the logistic map, and the standard map (see Appendix A for details). The results are shown in Figure 9 and confirm the dramatic improvement and generality of using Metropolis-Hastings to sample rare states. The simulations in Figure 9 use $t_*(\mathbf{x}) = t_o - 1$ instead of the one given by equation (54). This is computationally always more expensive because the correlations due to the neighborhood proposal are maximal (see the discussion after Eq. (54)), but on the other hand this proposal is simpler because it does not require estimating $d \log P / dE$ and λ_L .

Fig. 9. The number of samples required to sample a rare state, proportional to the round-trip time, scales polynomially with $N = t_o$ in Metropolis-Hastings with the proposal distribution derived in Section 3.4, as opposed to the exponential increase in uniform sampling. Tent map: equation (A.1) with $a = 3$; Logistic map: $F(x) = 4x(1 - x)$; Standard map [4]: $K = 8$ and $U(x) = \text{const.}$ was used in every case. The Wang-Landau algorithm was used to estimate the distribution prior to perform the flat-histogram and the distribution agrees with the analytical one when available [39,40]. The proposal distribution used was a mixed proposal composed by 50% chance of being the neighborhood proposal with $t_\star = t_o - 1$ and $\Delta = 1$, and 50% chance of being the shift proposal with $t_{\text{shift}} = 1$. Adapted from reference [25].

5.2 Transient chaos

The derivation of equation (61) uses the assumption that when $\mathbf{x}'$ is proposed with a scale given by equation (43) with $t_\star = t_e$, $\mathbb{E}[t_e(\mathbf{x}')|\mathbf{x}] = t_e(\mathbf{x})$, as per equation (57). We tested assumption in a similar way we tested the approximation of equation (46) for the FTLE, and consisted in uniformly drawing states $\mathbf{x}_i$, compute their escape time $t_e \equiv t_e(\mathbf{x}_i)$ and, for each, generate a state $\mathbf{x}'_i = \mathbf{x}_i + \delta(\mathbf{x})$ with a scale $\delta_x(\mathbf{x})$ given by equation (43) with $t_\star(\mathbf{x}) = t_e(\mathbf{x})$ and compute its escape time $t'_e \equiv t_e(\mathbf{x}'_i)$. The assumption is valid when, on average, $t'_e - t_e \ll t_e$ for large t_e . We did this experiment in the following systems: open tent map with $a = 3$ and $b = 5$, standard map with $K = 6$, and coupled Hénon maps with $D = 2, 4, 6, 8$ (see Appendix A) for $\Delta = 1$. We observed that in all 6 cases, the average $t'_e - t_e$ is smaller than 1 for all $t_e > 2/\kappa$. These observations show that the assumption of equation (57) is valid for a broad class of chaotic systems.

A second assumption tested here is the self-similarity argument used in deriving equation (58). The self-similarity assumption we use in equation (58) is that the landscape is self-similar such that, irrespectively of the particular scale we choose (by decreasing $t_\star$), the properties of the zoomed phase-space remain the same. In particular, we are interested in checking that $t'_e - t_e$ does not depend on t_e when we choose a larger scale, i.e. when $t_\star$ is decreased by a constant. We thus repeat the experiment above but we decrease $t_\star$ to $t_\star(\mathbf{x}) = t_e(\mathbf{x}) - T$ with $T = 0, 1/\kappa, 5/\kappa$ to check that increasing T decreases the average $t'_e - t_e$ without changing its independency with t_e .

Fig. 10. The acceptance rate of a Monte Carlo flat-histogram simulation is constant as a function of the escape time t when the Lyapunov exponent of the system λ_L is used. The simulation was made on the open tent map with $a = 3$ and $b = 5$ with the exact $P(t_e)$ given by equation (A.6). Different curves represent using the proposal with $\delta_x(\mathbf{x})$ given by equation (C.5) with three different exponents. When the exponent is larger than λ_L , $\delta_x(\mathbf{x})$ effectively aims for a smaller $t_\star$ and therefore a larger distance $t'_e(\mathbf{x}) - t_e(\mathbf{x})$, consequently decreasing the acceptance. When the exponent is smaller than λ_L , $\delta_x(\mathbf{x})$ aims for a larger $t_\star$ than the one given by equation (61), and therefore $t'_e(\mathbf{x}) - t_e(\mathbf{x})$ decreases to 0 as $t_e(\mathbf{x}) \rightarrow \infty$, and the acceptance converges to 1. Adapted from [24].

Our numerical tests in the same systems as before show that $t'_e - t_e$ decreases with increasing T and it remains independent of t_e , confirming our hypothesis.

The previous tests indicate that the proposal distribution should induce a constant acceptance rate when the Lyapunov exponent of the system, equation (C.2), is used. To confirm that this is the case, a flat-histogram simulation with an isotropic proposal distribution with width $\delta_x(\mathbf{x}) = \delta_x(t_e(\mathbf{x}))$ given by equation (C.2) with $\lambda_L(t_e) = \lambda_L$ was made. The results, Figure 10, reproduced from reference [24], confirm that proposing with λ_L guarantees a constant acceptance, and that any other exponent in equation (C.2) fails to achieve so.

The above tests confirm that the derivation made in Section 4.3.2 holds for a paradigmatic strongly chaotic open system. These tests also present a major advantage of using the approach in this paper: it allows to test the assumptions made on each step, something that other approaches, such as the ones in references [11,18,19], do not explicitly allow.

The efficiency of the simulation is tested in Figure 11 as a function of the maximal escape time considered, t_{max} , $\tau(t_{\text{max}})$ in the generic coupled Hénon maps defined by equation (A.8). It confirms the dramatic improvement of Metropolis-Hastings with the proposal derived in Section 2.2 over uniform sampling: the scaling is polynomial using importance sampling, and exponential in uniform sampling.

The derivation in Section 4.3.2, the tests presented above, and results in Figure 11, show the why and how importance sampling Metropolis-Hastings can efficiently sample long-living trajectories in strongly chaotic open systems. The proposal distribution should be applicable to strongly chaotic open systems more generally, as the approximations made are expected to be valid in other strongly chaotic systems.

Fig. 11. Polynomial scaling of the number of samples required to perform a round-trip ($1 \rightarrow t_{\max} \rightarrow 1$) as a function of $t_{\max}$ of the Metropolis-Hastings algorithm with the proposal derived in Section 4.3.2, as opposed to the exponential scaling in uniform sampling. This plot represents the average round-trip time of a flat-histogram simulation with the proposal given in Section 4.4 and number of samples required to sample $t_{\max}$ in using uniform sampling (line) in the coupled Hénon map, equation (A.8), for different dimensions. The two full lines represent $t_{\max}^2$ (lower) and $t_{\max}^3$ (upper), and the dashed line represents $1/P(t_e)$ for $D = 4$ (e.g. from Fig. 2; $D > 4$ have a even higher exponent). The flat-histogram was obtained by first running a Wang-Landau algorithm for 10 refinement steps, each with 100 round-trips. Each point represents the average round-trip time over 100 round-trips after the 10 refinement steps. The proposal distribution used was the isotropic, equation (41), with $\delta_x(\mathbf{x})$ given by equation (C.5) with $\delta_0 = 10$.

6 Conclusions

6.1 Summary of results

We have introduced a framework to sample rare trajectories in different classes of chaotic systems. It is based on the Metropolis-Hasting algorithm, a flexible and well-established Monte Carlo method suitable for the investigation of many numerical problems in chaotic dynamical systems (as shown in Sect. 2). Our main contribution is a procedure (see Sect. 4) to construct the proposal step of the Metropolis-Hasting algorithm (which proposes a new state $\mathbf{x}'$ given the current state $\mathbf{x}$) that ensures the efficiency of the sampling. The main arguments in the construction of this procedure are:

- (i) Set (in Sect. 4.1) as an heuristic goal to have a bounded (or constant) acceptance rate (23). This generalizes the traditional heuristic [59] used in Metropolis-Hastings, $E(\mathbf{x}') - E(\mathbf{x}) \sim 1$ and can therefore be used more generally in Metropolis-Hastings simulations.¹⁰

¹⁰ For example, using this heuristics, in a Metropolis-Hastings flat-histogram in the Ising model (in this case $\mathbf{x}$ corresponds to a list of all spins, $E(\mathbf{x})$ corresponds to the energy of the configuration), when the configuration of the system is close to the maximum of the density of states (maximal energy, no magnetization), there is no need to flip just one spin at the

Table 1. The four t_* derived for the two problems (escape time and FTLE as the observable) and two target distributions (canonical and flat-histogram). The four values of $t_*(\mathbf{x})$ reported in the table should be used in $\delta_x(\mathbf{x})$ of equation (43) and specify the time the proposed trajectory $\mathbf{x}'$ should stay close to the trajectory $\mathbf{x}$ for the acceptance to be bounded.

		Proposal correlation time $t_*(\mathbf{x})$	
Target distribution		Escape time $E(\mathbf{x}) = t_e(\mathbf{x})$	FTLE $E(\mathbf{x}) = t_{\lambda_t}(\mathbf{x})$
Canonical	$\pi(\mathbf{x}) \propto e^{-\beta E(\mathbf{x})}$	$t_e(\mathbf{x}) + \frac{a-1}{\beta} - \frac{1}{\kappa}$	$t - \frac{a-1}{\beta} \frac{1}{\lambda_L - \lambda_t(\mathbf{x})}$
Flat-histogram	$\pi(\mathbf{x}) \propto \frac{1}{G(E(\mathbf{x}))}$	$t_e(\mathbf{x}) - \frac{a}{\kappa}$	$t - \frac{a-1}{d \log G(\lambda_t)/d\lambda_t} \frac{1}{\lambda_L - \lambda_t(\mathbf{x})}$

- (ii) Introduce (in Sect. 4.2) an auxiliary quantity, the correlation time t_* , that quantifies the similarity between any two states. It is motivated by the notion that the observables considered in Section 2.3 are computed over trajectories, and $t_*(\mathbf{x}, \mathbf{x}')$ quantifies the time in which trajectories remain close to each other. We then showed how two proposal distributions, shift and neighborhood, can be used to control the average of $t_*(\mathbf{x}, \mathbf{x}')$ over $\mathbf{x}'$, $t_*(\mathbf{x})$.
- (iii) Derive (in Sect. 4.2.1) an expression for the values of $t_*(\mathbf{x})$ – equations (54) and (61) – which should be used in order to guarantee a constant acceptance. This is done for two observables of interest in chaotic systems (see Sect. 2) – the escape time $E(\mathbf{x}) = t_e(\mathbf{x})$ and the finite-time Lyapunov exponent (FTLE) $E(\mathbf{x}) = \lambda_{t_o}(\mathbf{x})$ – and two target distributions of the Metropolis-Hasting method (see Sect. 3.1) – canonical and flat-histogram. These results are summarized in Table 1.

A successful application of our framework leads to an algorithm in which the number of samples to obtain an independent rare sample scales polynomially with the difficulty of the problem, as opposed to the exponential increase observed in traditional uniform sampling (see Figs. 9 and 11).

6.2 Comparison to previous results

The importance of the proposal distribution has long been emphasized for Monte Carlo methods [7,61] and for sampling chaotic systems [11,13,14,16]. The questions that remain open from these works, and that we tackle in our paper, are how the efficiency of the sampling method is related to the proposal distribution and how an efficient proposal can be constructed from assumptions about the system. The proposals derived in this paper can be mapped, under appropriate simplifications, to known results from the literature, specifically to the proposals derived in references [24,25] and the stagger part of the algorithm presented in reference [11]. We make this connection explicitly in Appendix C. Another example of an algorithm that can be directly analyzed by the framework developed here

time: one can flip all spins at once because, there, correlating such states brings no advantage to increase the acceptance rate (it is going to be accepted anyway), but it increases r .

is the *precision shooting* proposed in references [23,67] and used in numerous applications of transition path sampling to chemical reactions [67–70]. The precision shooting method proposes a state $\mathbf{x}'$ isotropically distanced from $\mathbf{x}$ by $\delta_x(\mathbf{x})$ given by equation (43), with $t_\star(\mathbf{x}) = t_o$ where t_o is the length of the trajectory. From our results we conclude that this proposal is sub-optimal because it over-correlates the proposed state. Another application of the results of Section 4 is in the algorithm *Lyapunov Weighted Dynamics* of references [16,19], which uses a population Monte Carlo algorithm. As mentioned in these references (see also Ref. [20]), there is a parameter ε that controls how far new clones $\mathbf{x}'$ should be distanced from the existing clone $\mathbf{x}$, and that it should be neither too small nor too large. This plays a role similar to $\delta_x(\mathbf{x})$ in Section 4 and an optimal ε should therefore be related to our results.

In comparison to previous sampling methods in dynamical systems, including those mentioned above and others (e.g., Ref. [11]), the distinguishing feature of our results is that they provide an explicit connection between the proposal distribution and the acceptance rate. This connection, which is typically absent in Monte Carlo methods more generally, is extremely powerful because failures of the algorithm can be related to violations of the hypothesis (about the method and dynamical system) that we used in our derivations. Such violations should then be understood, and this understanding can then be inserted back in this methodology to generate new methods adapted for that situation. We hope this process will increase the range of applicability of our framework to other classes of dynamical systems (e.g., non-hyperbolic systems [33,44], see Ref. [71]) and observables E [72]. Possible improvements of our results can be obtained considering anisotropic search domains, an idea that has shown to be essential in the case of finding chaotic saddles in systems with more than one positive Lyapunov exponents [45].

6.3 Discussion

The proposal distribution is a way of moving in the phase-space stochastically and how to select a new state $\mathbf{x}'$ from a given state $\mathbf{x}$ is a general problem in different numerical techniques. Some of the most successful numerical algorithms in the literature, such as the golden section search or gradient descent, are essentially generic and efficient ways of selecting a new state. The results in Figures 7, and 9 show that the proposal distribution strongly influences the computational cost of the different procedures, often irrespectively of the particular sampling procedure (canonical or flat-histogram) and problem (sampling or finding). This reinforces the notion that the proposal of the new tentative state from the current state is a crucial factor when developing numerical techniques for optimization and numerical integration. Therefore we can expect that the main insights of this papers, e.g. the direct connection between fundamental properties of chaotic systems and the optimal proposal distribution, to be useful also for other problems (e.g., to the optimization problem

in which one is interested in maximizing or minimizing the observable).

The development of a numerical algorithm requires compromising between how fast it solves a particular problem, and how it is able to solve different problems. One interesting aspect of the algorithms (i.e. proposal distributions) introduced in this paper is that even though they can be made very specific (e.g. propose with the FTLE of the trajectory, Eq. (43)), they can also be made more general (e.g. propose using the Lyapunov exponent of the system, or the power-law proposal distribution, that does not use any specific information about the system). That is, more specificity requires more information (the FTLE of the trajectory) and makes the algorithm more efficient, and less information (only the Lyapunov of the system) makes the algorithm less specific, but typically also less efficient (proposals using less information, such as a power-law proposal, can be more robust in complicated cases [45]). This demonstrated capability of this methodology shows how it is not only useful to study a particular system on which some information about it is known, but also useful to situations on which less is known about the system. This does not mean that the methods apply to all problems, as there are important classes of chaotic systems on which some of the assumptions used in Section 4 are violated. For example, in non-hyperbolic systems [33,44] trajectories may remain correlated for a long time, which implies that one cannot assume that after $t_\star$ the trajectories are independent. Nevertheless, because the framework was outlined in the form of adding known information about the system, it is possible that improved insights about a class of chaotic systems can be translated to a faster algorithm.

One advantage of the methodology presented here is that it is not restricted to specific observables $E(\mathbf{x})$. In principle, it can be used to construct proposal distributions to sample rare states in different observables E , $E(\mathbf{x}) = t\lambda_t$, and $E(\mathbf{x}) = t_e(\mathbf{x})$. While it remains unclear what is the precise class of observables for which our methodology allows to construct an efficient proposal distribution, the different derivations of the proposal distribution do provide insights on the observables for which a proposal distribution could be constructed from properties of the system. As argued at the end of Section 4.3.1, observables computed as an average along the trajectories are similar to the FTLE and therefore the proposal distribution, derived in equation (56) should lead to efficient algorithms in these cases.

Altogether, our results reveal a fascinating interplay between the chaotic nature of some non-linear systems and the numerical techniques available to study rare events. It reinforces the idea that an efficient proposal *requires information about the system* (e.g., our derivation of the proposal distribution used the fact that “trajectories diverge exponentially” and that “the escape time function is a fractal-like function”). The analysis of this interplay allows to both better understand these systems and better understand these numerical techniques. This understanding opens perspectives to develop better

techniques to numerical study rare trajectories and extreme events in non-linear systems more generally [72].

J.C.L. was funded by FCT (Portugal), Grant No. SFRH/BD/90050/2012. E.G.A. thanks J. Wouters and T. Tél for helpful discussions. Open access funding provided by Max Planck Society.

Author contribution statement

All authors designed the research, developed the methodology, and wrote the manuscript. J.C.L. implemented the algorithms and performed the simulations.

Open Access This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Appendix A: Dynamical systems

In this appendix we describe the different dynamical systems that we use throughout the paper.

A.1 Skewed tent map

A paradigmatic example of a strongly chaotic system is the tent map [4], defined on $\Omega = [0, 1]$ by

$$F(x) = \begin{cases} ax & \text{for } x \leq 1/a \\ b(1-x) & \text{for } x > 1/a \end{cases} \quad (\text{A.1})$$

where $a > 1$ is a constant and $b \equiv a/(a-1)$. This map contains the main features of a chaotic system: it has a positive Lyapunov exponent ($\lambda_L = \frac{a \log(b) + b \log(a)}{a+b}$) and a positive measure. The finite-time Lyapunov exponent is given by

$$t_o \lambda_{t_o}(x) = i(x) \log a + (t_o - i(x)) \log b, \quad (\text{A.2})$$

where $i(x)$ is the number of times $x_t \in [0, 1/a]$. Its distribution of the finite-time Lyapunov exponent for $U(\mathbf{x}) = 1$, $P(E) = G(E)$, can be computed analytically and is a binomial,

$$\begin{aligned} G(E) &\equiv \int_0^1 \delta(E - t_o \lambda_{t_o}(\mathbf{x})) d\mathbf{x} \\ &= \sum_{i=0}^{t_o-1} \delta_{E, \lambda_{t_o}(i)} \binom{t_o}{i} \frac{1}{a^i} \left(1 - \frac{1}{a}\right)^{t_o-i}. \end{aligned} \quad (\text{A.3})$$

Fig. A.1. The open tent map, and its corresponding surviving set, equal to the construction of the cantor set. Left: the open tent map, where the escape correspond to states inside the interval in the middle, that maps to outside the unit interval. Right: an iteration of the open tent map with $a = b = 3$ corresponds to remove the middle third of each of the plateaus of the surviving set at time t and the set that survives this removal is the surviving set at $t + 1$. The third middle Cantor set is the surviving set at $t \rightarrow \infty$.

A.2 Standard map

The standard map considered here is defined by $\mathbf{x} \equiv (p, \theta) \in \Omega = [0, 1] \times [0, 1]$ that evolves in time according to

$$F(p, \theta) = \begin{cases} p + K/(2\pi) \sin(2\pi\theta) & \text{mod } 1 \\ \theta + p + K/(2\pi) \sin(2\pi\theta) & \text{mod } 1. \end{cases} \quad (\text{A.4})$$

We focus on the parameter $K = 6$ which leads to a phase-space with no visible KAM islands. We also consider the leaked (open) version this map by introducing a hole into the system at $\Lambda = [0.1, 0.1]$.

A.3 Skewed open tent map

The paradigmatic example of a strongly chaotic open system is the open tent map [10], defined on $\Omega = [0, 1]$ by

$$F(x) = \begin{cases} ax & \text{for } x \leq b/(a+b) \\ b(1-x) & \text{for } x > b/(a+b) \end{cases} \quad (\text{A.5})$$

where $a > 1$ and $b > a/(a-1)$. The state exits the system when it leaves the unit interval, i.e. $\Lambda = \mathbb{R} - \Omega$. This map contains the main features of an open chaotic system: it has a positive Lyapunov exponent ($\lambda_L = \frac{a \log(b) + b \log(a)}{a+b}$), an exponential decay of the escape time distribution,

$$P(t_e) = \kappa e^{-\kappa t_e}, \quad (\text{A.6})$$

with $\kappa = -\log(1/a + 1/b)$, and a conditionally invariant measure that is fractal with a (non-integer) fractal dimension D_0 given implicitly by

$$a^{-D_0} + b^{-D_0} = 1. \quad (\text{A.7})$$

A.4 Coupled open Hénon map

As a generic example of a high-dimensional strongly chaotic open system, we consider a set of d coupled Hénon maps on a ring, defined by a state $\mathbf{x} =$

Fig. A.2. The escape time function, $t_e(x)$, of the tent map for a generic a and b . There are 2^t intervals, and the size of each plateau can be analytically computed from the one at a previous time, and therefore be written analytically. Specifically, a given interval has size $\varepsilon(x) = (1 - 1/a - 1/b)a^{-i(x)}b^{-t_e(x)+i(x)}$, where $i(x)$ is the number of times $0 < F^t(x) < 1/a$, for $t = 1, \dots, t_e(x)$. The crucial observation is that $\lambda_{t_e(x)}(x)$ is proportional to $\log(\varepsilon(x))/t_e(x)$. Adapted from reference [24].

$(x_1, y_1, \dots, x_{d/2}, y_{d/2}) \in \Omega = \mathbb{R}^d$ where each individual map (x_i, y_i) evolves according to

$$\begin{pmatrix} x_i \\ y_i \end{pmatrix} = \begin{pmatrix} A_i - x_i^2 + B y_i + k(x_i - x_{i+1}) \\ x_i \end{pmatrix}, \quad (\text{A.8})$$

for $i = 1, \dots, d/2$, $d/2+1 \equiv 1$, and with parameters $k = 0.4$, $B = 0.3$, $A_1 = 3$ (if $d > 1$), $A_{d/2} = 5$, and $A_i = A_1 + (A_{d/2} - A_1)(i - 1)/(d/2 - 1)$. This choice of parameters ensures that a chaotic map is obtained in the $d = 2$ case and corresponds to the map studied in reference [11] for $d = 4$. The constraining region is $\Gamma = [-4, 4]^d$ because it covers the chaotic saddle of the system, and $\Lambda = \Omega - \Gamma$, i.e. the trajectory leaves the system if the absolute value of any of the coordinates is higher than 4. The escape function $E(\mathbf{x}) = t_e(\mathbf{x})$ for $d = 4$ is represented in Figure 2.

Appendix B: Efficiency of the uniform proposal

Here we show that the choice of sampling distribution does not necessarily decrease the scaling of the variance of the estimator. Our goal is to compute the average acceptance rate for a given escape time t_e in the canonical ensemble with a uniform proposal distribution, $g(\mathbf{x}'|\mathbf{x}) = 1/|\Gamma|$. The acceptance is given by $a(\mathbf{x}'|\mathbf{x}) = \min\{1, \exp(-\beta(t_e(\mathbf{x}') - t_e(\mathbf{x})))\}$ where $\beta < 0$ is used to reach higher $E(\mathbf{x}) = t_e(\mathbf{x})$. The acceptance of a state $\mathbf{x}$ is given by

$$a(\mathbf{x}) = \int_{\Gamma} d\mathbf{x}' a(\mathbf{x}'|\mathbf{x}) g(\mathbf{x}'|\mathbf{x}). \quad (\text{B.1})$$

Because $g(\mathbf{x}'|\mathbf{x}) = 1/|\Gamma|$ and π only depends on $t_e(\mathbf{x})$, $a(\mathbf{x})$ does not depend on $\mathbf{x}$, only on t_e : $a(\mathbf{x}) = a(t_e(\mathbf{x}))$. The average acceptance rate at a given t_e , $A(t_e) \equiv \mathbb{E}[a(\mathbf{x})|t_e]$, is given by

$$A(t_e) = \frac{1}{m(t_e)} \int_{\Gamma} d\mathbf{x} \delta(t_e - t_e(\mathbf{x})) e^{-\beta t_e(\mathbf{x})} a(\mathbf{x}). \quad (\text{B.2})$$

Because $a(\mathbf{x})$ only depends on t_e , it can be pulled out of the integral, and thus $A(t_e) = a(t_e)$. Taking into account

that $P(t'_e) = \kappa \exp(-t'_e \kappa)$, the acceptance rate can be computed analytically by integrating equation (B.1) and leads to

$$A(t_e) = \frac{e^{-\kappa t_e} \beta + e^{t_e \beta} \kappa}{\beta + \kappa}. \quad (\text{B.3})$$

This shows that the acceptance rate decays exponentially with increasing t_e (recall that $\beta < 0$). Because low acceptance implies that the random walk stays on the same state for a long time, this leads to an exponential increase of the autocorrelation time $T(E)$ and therefore an increase of the variance in equation (22). This same argument applies to a flat-histogram simulation, where $\pi(\mathbf{x}) \propto \exp(\kappa t_e(\mathbf{x}))$.

Appendix C: Simplified proposals

This section presents approximations that can be used to simplify both the implementation time and the computational cost of the proposals derived in Section 4. These approximations often reduce the general proposal derived in Section 4 to particular proposals already found in the literature, and therefore explains such proposals in this wider context.

C.1 Propose with the Lyapunov exponent in open systems

Equation (43) requires computing $\lambda_{t_e(\mathbf{x})}(\mathbf{x})$, even though the main interest is in $t_e(\mathbf{x})$. This calculation requires using a numerical algorithm or multiply a product of matrices [26], both of which have an associated computational cost. A simplification to this proposal is to approximate $\lambda_{t_e(\mathbf{x})}(\mathbf{x})$ by the maximum of the distribution of FTLE with finite-time t_e , $\lambda_L(t_e)$,

$$\lambda_{t_e(\mathbf{x})}(\mathbf{x}) \approx \lambda_L(t_e(\mathbf{x})). \quad (\text{C.1})$$

This approximation is valid as long as $\lambda_{t_e(\mathbf{x})}$ is not on the tails of the distribution of FTLE with finite-time t_e , $P(\lambda_{t_e})$. A sampling distribution that only depends on t_e , $\pi(\mathbf{x}) = \pi(t_e(\mathbf{x}))$, guarantees that states with the same t_e are equally sampled, $P(\mathbf{x}|t_e) = U(\mathbf{x})$, and therefore $P(\lambda_{t_e(\mathbf{x})}|t_e) = P(\lambda_{t_e})$. Furthermore, the approximation of using the maximum of the distribution holds because the tails of $P(\lambda_{t_e})$ decay exponentially with increasing t_e (see Sect. 2.1). Under these approximations, equation (43) can be simplified to

$$\delta_{\mathbf{x}}(\mathbf{x}) = \Delta e^{-\lambda_L(t_e(\mathbf{x})) t_{\star}(\mathbf{x})}. \quad (\text{C.2})$$

Furthermore, $\lambda_L(t_e)$ converges to the Lyapunov exponent of the system, λ_L , with increasing t_e . Therefore, a further simplification is to use the Lyapunov exponent of the system instead of $\lambda_L(t_e)$ in equation (C.2),

$$\delta_{\mathbf{x}}(\mathbf{x}) = \Delta e^{-\lambda_L t_{\star}(\mathbf{x})}. \quad (\text{C.3})$$

Furthermore, the $t_*(\mathbf{x})$ we derived in equation (61), in the flat-histogram ensemble and with $\log G(t_e) \propto -\kappa t_e$, can be written as

$$t_*(\mathbf{x}) = t_e(\mathbf{x}) - \frac{a}{\kappa}. \quad (\text{C.4})$$

Defining the constant $\delta_0 \equiv \Delta e^{-a\lambda_L/\kappa}$, we can write

$$\delta_x(\mathbf{x}) = \delta_0 e^{-\lambda_L t_e(\mathbf{x})}. \quad (\text{C.5})$$

This equation is exactly the proposal derived in reference [24], and shows that the proposal derived in Section 4.3.2 generalizes the proposal in reference [24].

C.2 Adaptively estimate the Lyapunov exponent

Using the proposal distribution with λ_L requires a priori knowledge of it, which typically is not available. This difficulty resembles the same problem that flat-histogram simulations have: $G(E)$ is required, but it is typically unknown a priori. This analogy motivates a Monte Carlo procedure that on the fly computes $\delta_x(t)$ that scales with λ_L .

Consider an hypothetical simulation with an isotropic proposal distribution (Eq. (41)) with

$$\delta_x(\mathbf{x}) = \sigma(t_e(\mathbf{x})), \quad (\text{C.6})$$

where $\sigma(t)$ is initially set to be $\sigma(t) = 1$ for every t . Consider also that the simulation reached a state $\mathbf{x}$ with a high escape time (e.g. $t_e = t_e(\mathbf{x}) = 10/\kappa$). A proposed state, $\mathbf{x}' = \mathbf{x} + \hat{\mathbf{h}}\sigma(t_e)$, will most likely have a much lower escape time (e.g. $t_e(\mathbf{x}') = 1/\kappa$). From equations (43) and (61), this indicates that $\sigma(t_e)$ is much higher than the “correct” proposal, $\delta_x(\mathbf{x})$, and therefore it should be reduced in the next proposal. The opposite is also true: when $\sigma(t_e)$ is much smaller than $\delta_x(\mathbf{x})$, $t_e(\mathbf{x}') = t_e(\mathbf{x})$ and it should be increased. This hypothetical simulation suggests that, in the same spirit as the Wang-Landau algorithm to approximate the density $P(t_e)$, there is the possibility to approximate $\delta_x(\mathbf{x})$ using an update scheme that can be inserted in the Metropolis-Hastings algorithm, and that is given by the same algorithm as the Wang-Landau (see Sect. 3.3.2), but instead of updating $P_{WL}(t)$, it updates also $\sigma(t)$ [24]:

$$\sigma(t_e) = \begin{cases} \sigma(t_e)f & \text{for } t_e(\mathbf{x}') = t_e, \\ \sigma(t_e)/f & \text{for } t_e(\mathbf{x}') < t_e. \end{cases} \quad (\text{C.7})$$

This update scheme generalizes the Wang-Landau procedure to the proposal distribution. It is expected to converge to a function $\sigma(t)$ that decays exponentially with the Lyapunov exponent of the system, and a proposal distribution with a constant acceptance rate in a flat-histogram simulation. It was extensively tested in different systems (tent map, full chaotic standard map with a leak, Coupled Hénon map with different D s), see references [24,45,71].

C.3 Power-law proposal distribution

The proposal distributions derived in the previous sections requires some knowledge about the state and the system: $\lambda_{t_o}(\mathbf{x})$, $t_e(\mathbf{x})$ (in open systems), λ_L of the system, and, in some situations, $G(E(\mathbf{x}))$. Another alternative to avoid computing $\delta_x(\mathbf{x})$ is to consider a proposal on which the time t_* that the two trajectories remain together is not imposed by a given $t_*(\mathbf{x})$, but that is a uniformly random variable between $[0, t_o]$ (FTLE) or $[0, t_e(\mathbf{x})]$ (open systems) that is generated on each proposal. Some values of t_* will be far from the optimal $t_*(\mathbf{x})$ and the corresponding $\mathbf{x}'$ will be rejected or it will be too close from $\mathbf{x}$, but others t_* will still be close from the optimal $t_*(\mathbf{x})$ and therefore useful.

Having a uniformly distributed correlation t_* still requires computing $\delta_x(\mathbf{x})$ in equation (43), which requires $\lambda_{t_o}(\mathbf{x})$. In the case $\lambda_t(\mathbf{x})$ is unknown (e.g. in open systems one could be only interested in the escape time and therefore not compute $\lambda_t(\mathbf{x})$), one may further approximate it by an uniform distribution between two extremes. Because the product of two uniformly distributed random variables is also uniformly distributed, this leads to a proposal distribution where $\delta_x(\mathbf{x})$ is given by $\exp(-U(a, b))$ where a and b are free parameters. By standard transformation of variables, this leads to a scale $\delta_x(\mathbf{x})$ that is power-law distributed and given by

$$P(\delta_x|\mathbf{x}) = P(\delta_x) = \frac{1}{\delta_x} \frac{1}{s_{\max} - s_{\min}} \in [\delta_{\min}, \delta_{\max}], \quad (\text{C.8})$$

where $\delta_{\min}$ and $\delta_{\max} \approx |\Gamma|$ are two free parameters. The δ in $\mathbf{x}' = \mathbf{x} + \mathbf{h}\delta$ is an half-normal distribution with a scale δ_x , but since this scale is now power-law distributed, it is no longer necessary to use the half-normal distribution altogether; instead, it is possible to just use the power-law proposal distribution where δ is drawn from $P(\delta_x)$ in equation (C.8), i.e. $|\mathbf{x}' - \mathbf{x}|$ is power-law distributed according to equation (C.8). Without the half-normal distribution the proposal distribution no longer depends on $\mathbf{x}$ and therefore $g(\mathbf{x}'|\mathbf{x})/g(\mathbf{x}|\mathbf{x}') = 1$.

The stagger part of the algorithm of reference [11] proposes exactly with a scale given by equation (C.8) and the argumentation above explains why the proposal distribution used in reference [11] to find states with high-escape time t_e is reported to work well: it is a proposal distribution that proposes $\mathbf{x}'$ correlated with $\mathbf{x}$ with a correlation t_* that is uniformly distributed, which eventually proposes $\mathbf{x}'$ with the optimal correlation $t_*(\mathbf{x})$. To confirm this explanation, let us consider a flat-histogram simulation with a power-law proposal distribution on the open tent map and consider the measurement of $\mathbb{E}[\log \delta_x|A^*, t_e]$, where A^* is the condition $\varepsilon < a(\mathbf{x}'|\mathbf{x}) < 1 - \varepsilon$ (of bounded acceptance). Under the above argumentation, the scale $\log \delta_x$ that contributes to a bounded acceptance is given by $-\lambda_{t_e} t_e$, per equation (C.5). This is confirmed by numerical simulation, shown in Figure C.1, and was obtained also for the problem of finding rare states, as reported in reference [45]. This result, combined with the derivation of $t_*(\mathbf{x})$, explains the success of the proposal (the stag-

Fig. C.1. The relevant proposals of the power-law proposal are those on which the scale is given by $\delta_x(\mathbf{x})$ in equation (C.5). The x axis is the escape time; the y axis is the best estimator of $\mathbb{E}[\log \delta_x | A^*, t_e]$ (black dots, 2σ) over 4×10^5 samples obtained from a flat-histogram simulation with the power-law proposal given by equation (C.8), for each escape time t_e and conditional to on acceptance $A^* = \{\varepsilon < a(\mathbf{x}'|\mathbf{x}) < 1 - \varepsilon\}$ with $\varepsilon = 0.1$. The power-law proposal distribution samples all scales, but the scales suitable for Metropolis-Hastings depend on t_e as $\exp(-\lambda t_e)$, as expected from the results of Section C.1. This simulation was made on the open tent map, equation (A.5), with $a = 3$ and $b = 5$. The best estimate of $-\lambda_{t_e}(\mathbf{x})t_e + \log(\Delta)$ with $\Delta = 50$ corresponds to the red line (2σ). The correspondence of the two curves indicates that the scale of the power-law proposal whose acceptance rate is bounded corresponds to the scale given by $\delta_x(\mathbf{x})$. The parameters used in the power-law proposal were $\delta_{\max} = 1$, $\delta_{\min} = 2^{-40}$.

ger) used in reference [11] from basic notions of chaotic systems and numerical methods.

Appendix D: Algorithmic description

It is useful to summarize the different proposals in an algorithmic form so they can be easily implemented. In all cases, the proposal requires the current state of the random walk, denoted by $\mathbf{x}$. In our implementation (see Ref. [66]) we use arbitrary precision libraries that allow us to use proposals with arbitrary small widths and therefore to test our results beyond the precision of particular computers. For finite-precision implementations it is important to verify that the proposal width $\delta_x(\mathbf{x})$ is larger than round-off errors and one could also consider alternative proposals that avoid this limitation (e.g., to propose trajectories using the time-evolved $\mathbf{x}$).

D.1 FTLE in closed systems

1. Generate a unitary vector $\hat{\delta}$ in D dimensions;
2. Compute $t_* = t_o - \frac{\log(a)}{d \log \pi(E)/dE(E)} \frac{1}{|\lambda_L - \lambda_{t_o}(\mathbf{x})|}$, equation (54);
3. Compute $\delta_x = \Delta \exp(-t_* \lambda_{t_o}(\mathbf{x}))$, equation (43);
4. Generate a random number δ from a normal distribution with mean 0 and variance δ_x^2 ;

5. Make $\mathbf{x}' = \mathbf{x} + \hat{\delta}|\delta|$.

Where $\Delta \in \mathbb{R}$ is free parameter (e.g. $\Delta = 0.1$) and a is the chosen average acceptance (e.g. 0.5). For example, in a canonic ensemble with parameter β , $d \log \pi(E)/dE(E) = \beta$ and therefore $t_* = t_e(\mathbf{x}) - \frac{\log(a)}{\beta} \frac{1}{|\lambda_L - \lambda_{t_o}(\mathbf{x})|}$. The value of λ_L can be estimated using e.g. the first samples of the random walk. In the flat-histogram ensemble, λ_L is the maximum of $G(\lambda_{t_o})$ and $\log \pi(E)/dE(E) = \log G(E)/dE(E)$, or by an approximation of it, e.g. using G_{WL} of the Wang-Landau algorithm.

D.2 Open systems

1. Generate a unitary vector $\hat{\delta}$ in D dimensions;
2. Compute $t_* = t_e(\mathbf{x}) - 1/\kappa - \frac{a-1}{d \log \pi/dt_e}$, equation (61);
3. Compute $\delta_x = \Delta \exp(-t_* \lambda_{t_o}(\mathbf{x}))$, equation (43);
4. Generate a random number δ from a normal distribution with mean 0 and variance δ_x^2 ;
5. Make $\mathbf{x}' = \mathbf{x} + \hat{\delta}|\delta|$.

Where δ_0 is a free parameter (e.g. $\delta_0 = 0.1$). Both $\lambda_{t_e}(\mathbf{x})$ and $t_e(\mathbf{x})$ are required by the proposal and both can be computed during the same evolution of the system: $t_e(\mathbf{x})$ is the time until the trajectory enters the exit region Λ , $\lambda_{t_e}(\mathbf{x})$ is the FTLE of this trajectory.

References

1. S. Albeverio, V. Jentsch, H. Kantz, *Extreme Events in Nature and Society* (Springer, 2006), ISBN 978-3-540-28610-3
2. P. Holmes, Phys. Rep. **193**, 137 (1990)
3. A.E. Motter, D.K. Campbell, Phys. Today **66**, 27 (2013)
4. E. Ott, *Chaos in Dynamical Systems*, 2nd edn. (Cambridge University Press, Cambridge, 1993), ISBN 0 521 43215 4
5. E.N. Lorenz, J. Atmos. Sci. **20**, 130 (1963)
6. M.E.J. Newman, G.T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, USA, 2002), ISBN 0198517971
7. C.P. Robert, G. Casella, *Monte Carlo statistical methods*, Springer texts in statistics, 2nd edn. (Springer, Berlin, 2005), ISBN 0387212396
8. J.A. Bucklew, *Introduction to Rare Event Simulation*, Springer Series in Statistics (Springer New York, New York, NY, 2004), ISBN 978-1-4419-1893-2, <http://link.springer.com/10.1007/978-1-4757-4078-3>
9. V. Lucarini, D. Faranda, A.C.M. Freitas, J.M. Freitas, T. Kuna, M. Holland, M. Nicol, M. Todd, S. Vaienti, [arXiv:1605.07006](https://arxiv.org/abs/1605.07006), 2016
10. Y.C. Lai, T. Tél, *Transient chaos: complex dynamics in finite time scales*, 1st edn. (Springer, New York, 2010), Vol. 173
11. D. Sweet, H.E. Nusse, J.A. Yorke, Phys. Rev. Lett. **86**, 2261 (2001)
12. E.M. Bollt, Int. J. Bifurc. Chaos **15**, 1615 (2005)
13. C. Dellago, P.G. Bolhuis, P.L. Geissler, Adv. Chem. Phys. **123**, 1 (2002)

14. P.G. Bolhuis, D. Chandler, C. Dellago, P.L. Geissler, *Annu. Rev. Chem.* **53**, 291 (2002)
15. C. Giardinà, J. Kurchan, L. Peliti, *Phys. Rev. Lett.* **96**, 120603 (2006)
16. J. Tailleur, J. Kurchan, *Nat. Phys.* **3**, 203 (2007)
17. T. Yanagita, Y. Iba, *J. Stat. Mech.: Theory Exp.* **2009**, P02043 (2009)
18. A. Kitajima, Y. Iba, *Comput. Phys. Commun.* **182**, 251 (2011)
19. T. Laffargue, K.D.N.T. Lam, J. Kurchan, J. Tailleur, *J. Phys. A: Math. Theor.* **46**, 254002 (2013)
20. J. Wouters, F. Bouchet, [arXiv:1511.02703](https://arxiv.org/abs/1511.02703), 2015, pp. 1–29
21. P. Geiger, C. Dellago, *Chem. Phys.* **375**, 309 (2010)
22. S.I. Sasa, K. Hayashi, *Europhys. Lett.* **74**, 156 (2006)
23. M. Grünwald, C. Dellago, P.L. Geissler, *J. Chem. Phys.* **129**, 194101 (2008)
24. J.C. Leitão, J.M.V.P. Lopes, E.G. Altmann, *Phys. Rev. Lett.* **110**, 220601 (2013)
25. J.C. Leitão, J.M.V.P. Lopes, E.G. Altmann, *Phys. Rev. E* **90**, 052916 (2014)
26. P. Cvitanovic, R. Artuso, R. Mainieri, G. Tanner, G. Vattay, *Chaos book* (Niels Bohr Institute, Copenhagen, 2016)
27. P. Grassberger, H. Kantz, *Phys. Lett. A* **113**, 167 (1985)
28. P. Grassberger, R. Badii, A. Politi, *J. Stat. Phys.* **51**, 135 (1988)
29. M. Sepúlveda, R. Badii, E. Pollak, *Phys. Rev. Lett.* **63**, 1226 (1989)
30. D. Beigie, A. Leonard, S. Wiggins, *Phys. Rev. Lett.* **70**, 275 (1993)
31. C. Amitrano, R. Berry, *Phys. Rev. Lett.* **68**, 729 (1992)
32. S. Olmi, *Chaos* (Woodbury, NY) **25**, 123125 (2015)
33. G.M. Zaslavsky, *Phys. Rep. Rev. Sect. Phys. Lett.* **371**, 461 (2002)
34. J. Szezech, S. Lopes, R. Viana, *Phys. Lett. A* **335**, 394 (2005)
35. R. Artuso, C. Manchein, *Phys. Rev. E* **80**, 036210 (2009)
36. C. Manchein, M.W. Beims, J.M. Rost, *Chaos* (Woodbury, NY) **22**, 033137 (2012)
37. H.D.I. Abarbanel, R. Brown, M.B. Kennel, *J. Nonlinear Sci.* **1**, 175 (1991)
38. H.D.I. Abarbanel, R. Brown, M.B. Kennel, *J. Nonlinear Sci.* **2**, 343 (1992)
39. C. Beck, F. Schögl, *Thermodynamics of chaotic systems: an introduction* (Cambridge University Press, 1995), ISBN 9780521433679
40. A. Prasad, R. Ramaswamy, *Phys. Rev. E* **60**, 2761 (1999)
41. H. Richter, *Int. J. Parallel Emerg. Distrib. Syst.* **1** (2017)
42. U.H.E. Hansmann, Y. Okamoto, *J. Comput. Chem.* **14**, 1333 (1993)
43. T. Tél, Y.C. Lai, *Phys. Rep.* **460**, 245 (2008)
44. E.G. Altmann, J.S.E. Portela, T. Tél, *Rev. Mod. Phys.* **85**, 869 (2013)
45. M. Sala, J.C. Leitão, E.G. Altmann, *Chaos: Interdiscip. J. Nonlinear Sci.* **26**, 123124 (2016)
46. M. Dhamala, Y.C. Lai, E.J. Kostelich, *Phys. Rev. E: Stat. Nonlinear Soft Matter Phys.* **64**, 056207 (2001)
47. G. Torrie, J. Valleau, *J. Comput. Phys.* **23**, 187 (1977)
48. R.H. Swendsen, J.S. Wang, *Phys. Rev. Lett.* **57**, 2607 (1986)
49. E. Marinari, G. Parisi, *Europhys. Lett.* **19**, 451 (1992)
50. Y. Sugita, Y. Okamoto, *Chem. Phys. Lett.* **314**, 141 (1999)
51. J. Lee, *Phys. Rev. Lett.* **71**, 211 (1993)
52. B.A. Berg, T. Neuhaus, *Phys. Lett. B* **267**, 249 (1991)
53. F. Wang, D.P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001)
54. J. Viana Lopes, M. Costa, J. Lopes dos Santos, R. Toral, *Phys. Rev. E* **74**, 046702 (2006)
55. C.G. Zhou, R.N. Bhatt, *Phys. Rev. E* **72**, 025701 (2005)
56. R.E. Belardinelli, S. Manzi, V.D. Pereyra, *Phys. Rev. E* **78**, 067701 (2008)
57. S. Trebst, D.A. Huse, M. Troyer, *Phys. Rev. E* **70**, 046701 (2004)
58. P. Dayal, S. Trebst, S. Wessel, D. Wurtz, M. Troyer, S. Sabhapandit, S.N. Coppersmith, *Phys. Rev. Lett.* **92**, 097201 (2004)
59. J. Viana Lopes, Ph.D. thesis, Universidade do Porto, 2006
60. M.D. Costa, J.V. Lopes, J.M.B.L. dos Santos, *Europhys. Lett.* **72**, 802 (2005)
61. G.O. Roberts, A. Gelman, W.R. Gilks, *Ann. Appl. Prob.* **7**, 110 (1997)
62. R. Fischer, J.C. Leitão, T.P. Peixoto, E.G. Altmann, *Phys. Rev. Lett.* **115**, 188701 (2015)
63. P. Grassberger, *Phys. Rev. E* **56**, 3682 (1997)
64. C. Grebogi, S. Hammel, J. Yorke, T. Sauer, *Phys. Rev. Lett.* **65**, 1527 (1990)
65. Y. Iba, N. Saito, A. Kitajima, *Ann. Inst. Stat. Math.* **66**, 611 (2014)
66. J. Leitao, *A library to sample chaotic systems*, 2017, <https://github.com/jorgecarleitao/chaospp>
67. M. Grünwald, Ph.D. thesis, University of Wien, 2009
68. C.N. Rowley, T.K. Woo, *J. Chem. Phys.* **131**, 234102 (2009)
69. N. Eidelson, B. Peters, *J. Chem. Phys.* **137**, 094106 (2012)
70. B. Schaefer, S. Mohr, M. Amsler, S. Goedecker, [arXiv:1401.8081](https://arxiv.org/abs/1401.8081), 2014, p. 15
71. J.C. Leitão, Ph.D., Technical University of Dresden, 2016, <http://www.qucosa.de/fileadmin/data/qucosa/documents/20901/main.pdf>
72. S. Gupta, J.C. Leitao, E.G. Altmann, *Phys. Rev. E* **96**, 012201 (2017)