



HAL
open science

Computing non-equilibrium trajectories by a deep learning approach

Eric Simonnet

► **To cite this version:**

Eric Simonnet. Computing non-equilibrium trajectories by a deep learning approach. 2022. hal-03854453

HAL Id: hal-03854453

<https://hal.science/hal-03854453>

Preprint submitted on 15 Nov 2022

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

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

Computing non-equilibrium trajectories by a deep learning approach

Eric Simonnet

^aINPHYNI, Université Côte d'Azur et CNRS, UMR 7010, 1361, route des Lucioles, 06560, Valbonne, France

Abstract

Predicting the occurrence of rare and extreme events in complex systems is a well-known problem in non-equilibrium physics. These events can have huge impacts on human societies. New approaches have emerged in the last ten years, which better estimate tail distributions. They often use large deviation concepts without the need to perform heavy direct ensemble simulations. In particular, a well-known approach is to derive a minimum action principle and to find its minimizers.

The analysis of rare reactive events in non-equilibrium systems without detailed balance is notoriously difficult either theoretically and computationally. They are described in the limit of small noise by the Freidlin-Wentzell action. We propose here a new method which minimizes the geometrical action instead using neural networks: it is called *deep gMAM*. It relies on a natural and simple machine-learning formulation of the classical gMAM approach. We give a detailed description of the method as well as many examples. These include bimodal switches in complex stochastic (partial) differential equations, quasi-potential estimates, and extreme events in Burgers turbulence.

Keywords: Freidlin-Wentzell large deviation theory, instantons, geometrical action, gMAM, quasi-potential, Neural Networks

1. Introduction

Understanding rare and extreme events in complex models has become a cornerstone of modern non-equilibrium physics. In many cases, the occurrence of such events has strong impacts on our daily life and cannot be ignored despite their small probability of occurrence. A well-known example is of course climate change and its impact at the regional scale. At the level of biological systems, dramatic phenomena can be tied to very small unexpected changes, in biochemistry, molecular dynamics, with protein bindings/foldings for instance. Remarkably, these events are predictable: they always conspire following the same path associated with some well-defined probability. In fact, small fluctuations either random or deterministic can drive systems to different unknown equilibria provided some energy barrier is reached. This is the Arrhenius law. It is a very simple instance of a large deviation principle (LDP) which expresses the property that a quantity (e.g. a probability) is behaving as $-\epsilon \log \text{Pr} \sim V$ when some parameter ϵ becomes small. Such LDP appears to be much more general than the original Arrhenius law and is not restrictive to the existence of a free energy potential V . In general situations V is then called the *quasi-potential* or rate function. However, understanding non-equilibrium systems which do not have a detailed balance turns out to be a very challenging task, not only theoretically, but computationally as well as.

The natural way to tackle these problems is large deviation theory as it provides a natural and powerful framework for non-equilibrium statistical physics [67]. In the context of metastability, the theory has been developed by Freidlin and Wentzell in the 70s-80s [19] as a large deviation approach to perturbed dynamical systems although these ideas were already known by the physicists with Onsager and Machlup [48]. Noteworthy, there is another path to large deviations which is called the Martin-Siggia-Rose-Jansen-de Dominicis (MSRJD) formalism developed in the 70s ([27] for a

Email address: eric.simonnet@inphyni.cnrs.fr (Eric Simonnet)

presentation). Once a minimum action principle is obtained, it remains to identify the minimizers. This is in general a very difficult optimisation problem as it involves complex quantities related to the nontrivial nonlinear interactions with noise. It turns out that minimizing the Freidlin-Wentzell action directly is also very difficult due to the long time it takes for the fluctuations to build up an optimal path history. It translates mathematically as an optimisation problem in the large time limit which turns out to be ill-conditioned. A key step is therefore to identify an equivalent problem where time has been reparametrized. It gives the so-called *geometrical action* [69]. This has led many authors to consider the so-called gMAM approach bringing essential advantages over the original form. How to minimize this action efficiently using neural networks and machine-learning techniques is the subject of this work.

This work is organized as follow. We first recall fundamental notions of the Freidlin-Wentzell large deviation theory and the concept of geometrical action in Section 2. We then describe in Section 3 how to adapt these problems to the machine-learning context and give a detailed description of the deep gMAM approach and what it brings compared to the classical one. Section 4 gives simple introductory examples of stochastic differential equations (SDEs). It also illustrates as a byproduct a very simple way to compute the quasi-potential for general SDEs. Section 5 considers more challenging examples of stochastic partial differential equations (SPDEs). We then illustrate the deep gMAM method in a very different context of extreme events in Burgers turbulence in Section 6. We conclude in Section 7. Appendix 8.1 discusses a striking example of a deterministic chaotic system and 8.2 provides a simple snippet Julia code of deep gMAM for the SDE case.

2. Definitions and generalities

2.1. Freidlin-Wentzell large deviation theory

We consider the following stochastic (partial) differential equation

$$du = F(u) dt + \sqrt{\epsilon}\sigma(u)dW_t, \quad u \in H \quad (1)$$

where $\epsilon > 0$, W_t is a Wiener process, H is an ad-hoc Hilbert functional space, $u : (0, T) \times \mathcal{D} \rightarrow \mathbb{R}^p : (t, \mathbf{x}) \mapsto u(t, \mathbf{x})$ with $\mathcal{D} \subset \mathbb{R}^d$ is the spatial domain and u is a vector field (u_1, \dots, u_p) . For simplicity, we address the scalar fields only with $p = 1$. The correlation tensor (or diffusion matrix/operator) is

$$\chi = \sigma\sigma^*.$$

The inner product is denoted $\langle \cdot, \cdot \rangle$ and the corresponding norm is $\|u\| = (\langle u, u \rangle)^{\frac{1}{2}}$. We are interested in the probability \Pr to have a time- T trajectory (possibly with $T \rightarrow \infty$) connecting two distinct states a and b when $\epsilon \rightarrow 0$. It is well-known from Freidlin-Wentzell theory [19] that it satisfies the large deviation principle (LDP)

$$\lim_{\epsilon \rightarrow 0} -\epsilon \log \Pr = \inf_{u \in \mathcal{C}} S_T[u], \quad (2)$$

where the set of trajectories starting from a and ending to b at time T are denoted

$$\mathcal{C} \equiv \mathcal{C}_{a,b,T} = \{u : (t, \mathbf{x}) \mapsto u(t, \mathbf{x}), u(0, \mathbf{x}) = a(\mathbf{x}), u(T, \mathbf{x}) = b(\mathbf{x})\}, \quad (3)$$

and S_T is the so-called Freidlin-Wentzell action:

$$S_T[u] = \frac{1}{2} \int_0^T \langle (\dot{u} - F(u), \chi^{-1}(\dot{u} - F(u))) \rangle dt. \quad (4)$$

The argmin of the right-hand side (2) is called in the literature an *instanton path*. In the context where a, b are local attractors for the deterministic dynamics ((1) with $\epsilon = 0$), the (global) argmin solution is also often called a *non-equilibrium transition* or *maximum likelihood path*. The minimum

action viewed as a function of a and b is called the *quasi-potential* [19, 8], the definition in fact involves another minimisation w.r.t. T :

$$V : H \times H \rightarrow \mathbb{R}, (a, b) \mapsto \inf_T \inf_{u \in \mathcal{C}_{a,b,T}} S_T[u]. \quad (5)$$

Strictly speaking, FW theorem applies for finite-dimensional systems only and is very generic with mild conditions, e.g. continuity hypothesis. More specifically, the diffusion matrix χ must be uniformly non-degenerate: $\langle p, \chi p \rangle \geq c \|p\|^2, c > 0$ for a LDP to apply (see [19], chap. 5.3). If not the BV constraints (3) cannot be satisfied.

The case of SPDEs is mathematically much more difficult to handle. A well-known historical example is a rigorous LDP proof for the Ginzburg-Landau 1-D PDE [18]. We thus consider all the following at a formal level only. We will not write explicitly the dependency on \mathbf{x} unless needed. A more general formulation is to consider the minimum action problem as

$$\inf_{u \in \mathcal{C}} \int_0^T L(u, \dot{u}) dt \text{ with } L(u, v) = \sup_p (\langle v, p \rangle - H(u, p)), \quad (6)$$

where H is the Hamiltonian. It can take different forms depending on the class of problems and L is its Legendre transform. The advantage of this formulation is its generality, in particular when the Lagrangian has no closed-form expression (see an example in [19], chap. 5.3 and also [33]). Moreover, one does not need to invert χ . In this work however, we focus on diffusion problems (1) only, in which case the Hamiltonian takes an explicit form

$$H(u, p) = \langle F(u), p \rangle + \frac{1}{2} \langle p, \chi p \rangle. \quad (7)$$

In the following, we propose an equivalent formulation of (6). Instead of considering the Lagrangian (4): $L(u, \dot{u}) = \frac{1}{2} \langle \dot{u} - F(u), \chi^{-1}(\dot{u} - F(u)) \rangle$, we consider the constrained problem:

$$\inf_{u \in \mathcal{C}} \frac{1}{2} \int_0^T \langle p, \chi p \rangle dt, \quad \chi p = \dot{u} - F(u). \quad (8)$$

The general idea of the proposed deep gMAM method, is to solve (8) by a penalisation method. We can therefore handle more complicated situations where χ^{-1} cannot be expressed easily by enriching the cost functional by an additional penalty term. We recall that the notation χp means the operator χ applied to the field p . This formulation is very natural, a similar approach can also be found in [59].

We will provide a highly nontrivial PDE example with a complicated χ in section 6. We next describe the so-called geometrical action.

2.2. Geometrical action

It is often the case that one is interested in minimizing S_T w.r.t. T as well. It is in general equivalent to consider the limit $T \rightarrow \infty$. This situation typically happens in the context of transitions from one equilibrium state (or loosely speaking a local attractor) to another, the transition takes an infinite amount of time to escape from the starting equilibrium. The minimisation problem is therefore

$$\inf_{T > 0} \inf_{u \in \mathcal{C}} \frac{1}{2} \int_0^T \|\dot{u} - F(u)\|_{\chi^{-1}}^2 dt,$$

with $\|f\|_{\chi^{-1}} \equiv \langle f, \chi^{-1} f \rangle$ the induced covariance norm, see (4). In order to avoid the double minimisation problem, it is possible to derive an equivalent formulation which removes the time constraint as shown by [33] (see also [20]). The idea is the following. Using $\|u\|^2 + \|v\|^2 \geq 2\|u\| \|v\|$ with equality when $\|u\| = \|v\|$, one has $\frac{1}{2} \int_0^T \|\dot{u} - F(u)\|_{\chi^{-1}}^2 dt \geq \int_0^T (\|\dot{u}\|_{\chi^{-1}} \|F(u)\|_{\chi^{-1}} - \langle \dot{u}, F(u) \rangle_{\chi^{-1}}) dt$. The remarkable observation is that this inequality is indeed an equality. One can choose a reparametrization of time say $\gamma(t) \geq 0$ such that $\gamma(0) = 0$ and $\gamma(T) = T$, the change of variable $t = \gamma(\tau)$ yields

$\int_0^T (\|\dot{u}\|_{\chi^{-1}} \|F(u)\|_{\chi^{-1}} - \langle \dot{u}, F(u) \rangle_{\chi^{-1}}) dt = \int_0^T (\|u'\|_{\chi^{-1}} \|F(u)\|_{\chi^{-1}} - \langle u', F(u) \rangle_{\chi^{-1}}) d\tau$, with $u' = du/d\tau$. The trick is now to choose τ such that $\|u'\|_{\chi^{-1}} = \|F(u)\|_{\chi^{-1}}$. This is always possible, by taking $\gamma' = \|\dot{u}\|_{\chi^{-1}} / \|F(u)\|_{\chi^{-1}}$. The minimisation problem is equivalent to (changing the notation say $\dot{u} = du/d\tau$)

$$\inf_{u \in \mathcal{G}} \int_0^1 (\|\dot{u}\|_{\chi^{-1}} \|F(u)\|_{\chi^{-1}} - \langle \dot{u}, F(u) \rangle_{\chi^{-1}}) d\tau, \quad \mathcal{G} \equiv \{u : \tau \rightarrow u(\tau), u(0) = a, u(1) = b\}. \quad (9)$$

It is therefore the location on the curve which prevails, the way it is parametrized does not matter. A more rigorous proof can be found in [33].

In practice, when χ has a complicated form, ie when χ^{-1} has no explicit formulation, we just formulate the same problem like in (8) avoiding χ^{-1} , namely solving the constraint problem

$$\inf_{u \in \mathcal{G}} \int_0^1 (\|v\|_{\chi} \|w\|_{\chi} - \langle v, w \rangle_{\chi}) d\tau, \quad \chi v = \dot{u}, \text{ and } \chi w = F(u). \quad (10)$$

and its penalisation (quadratic) formulation:

$$\inf_{u \in \mathcal{G}, v, w} \int_0^1 (\|v\|_{\chi} \|w\|_{\chi} - \langle v, w \rangle_{\chi}) d\tau + \gamma_v \int_0^1 \|\chi v - \dot{u}\|^2 d\tau + \gamma_w \int_0^1 \|\chi w - F(u)\|^2 d\tau. \quad (11)$$

The penalisation parameters $\gamma_v, \gamma_w \gg 1$ must be large enough but at the same time not too large to avoid dealing with an ill-conditioned functional. The norm involved in the geometrical action is the χ -norm, by contrast with the two constraints on (v, w) involving some user-defined one, e.g. an Euclidian weighted norm. The treatment of (10) using (11) is not the only possibility, other strategies are possible, e.g. augmented Lagrangian methods (see [35] in the machine-learning context applied to elliptic and eigenvalue problems, and [59] using some classical approach). One can also consider the minimax original problem (6) using adversarial networks. In the context of machine-learning however, (11) appears to be the simplest strategy giving very good results. This is explained in details in the next section 3.

3. Deep gMAM method

We describe here in more details how to use Neural Networks (NNs) for solving (11). The general idea is to parametrize the argmin solution of (11) by some NNs. In doing so, we already make a choice of working into the physical space-time domain $(t, \mathbf{x}) \in (0, 1) \times \mathcal{D}$. Due to that, convolutional neural networks are not appropriated in this context. It amounts to the fact that one has access to a macroscopic description of the action and takes full advantage of it. The broad philosophy of the proposed approach is therefore the same than for the so-called physics-informed neural networks (PINNs), [52, 64, 72] except that the system to solve is more involved with possibly much more complicated constraints. We will discuss the technical aspects below in the subsections 3.1–3.6 and summarize what it can bring in the last subsection 3.7.

3.1. Neural Network parametrization and cost functional

We consider here one of the simplest neural architecture, the so-called fully-connected feedforward NNs. Let $\mathcal{N} : \mathbb{R}^{d+1} \times \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$, $(t, \mathbf{x}, \boldsymbol{\theta}) \mapsto \mathcal{N}(t, \mathbf{x}; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the vector of n_θ parameters. The function \mathcal{N} is defined as

$$\mathcal{N}(t, \mathbf{x}; \boldsymbol{\theta}) = \mathcal{N}_{\text{out}} \circ \mathcal{N}_L \cdots \mathcal{N}_1 \circ \mathcal{N}_{\text{in}}(t, \mathbf{x}),$$

where \circ is the composition and $\mathcal{N}_k(\mathbf{y}) = \sigma_k(W_k \mathbf{y} + \mathbf{b}_k)$. The parameters $\mathbf{b}_k \in \mathbb{R}^{l_k}$ are called NN biases and the $(l_k \times c_k)$ matrices W_k are called NN weights. The activation functions σ_k must be nonlinear. In the examples discussed later, the dimensions are chosen independent of k with $l_k = c_k = c$. The number of hidden layers L is the depth, whereas c is the network capacity. In the

following, we consider only *swish* activation functions [53]. They correspond to some non-monotonic version of the classical RELU functions $\max(0, x)$, with $\sigma(x) = x/(1 + e^{-x})$. Other choices are of course possible, such as RELU or tanh but in practice the swish nonlinearity gives very good results and at the same time smoother representations. The last layer \mathcal{N}_{out} is in general linear. The depth L and c are important parameters. The expressivity of \mathcal{N} becomes better as L and c increases but at the same time, it becomes harder to train, ie finding ad-hoc optimal values for θ .

As stated above, we now replace the main and auxiliary fields (u, v, w) in (11) by their NN parametrizations $u \rightarrow \mathcal{N}_u$, $v \rightarrow \mathcal{N}_v$ and $w \rightarrow \mathcal{N}_w$ with parameters $\theta_u, \theta_v, \theta_w$. The cost functional is then minimized with respect to these parameters: we have thus performed a nonlinear projection onto the neural network space. This is indeed a very general methodology in machine-learning. The cost functional is written as

$$\mathcal{C}[\theta_u, \theta_v, \theta_w] = \gamma_g \mathcal{C}_g[\theta_v, \theta_w] + \gamma_{\text{arc}} \mathcal{C}_{\text{arc}}[\theta_u] + \mathcal{C}_{\text{constraints}} + (\gamma_{\text{bcs}} \mathcal{C}_{\text{BCs}}). \quad (12)$$

The first term corresponds to the geometrical action in its continuous form:

$$\mathcal{C}_g[\theta_v, \theta_w] = \int_0^1 (\|\mathcal{N}_v\|_\chi \|\mathcal{N}_w\|_\chi - \langle \mathcal{N}_v, \mathcal{N}_w \rangle_\chi) d\tau, \quad \|\mathcal{N}\|_\chi \equiv \left(\int_{\mathcal{D}} \mathcal{N}(t, \mathbf{x}) (\chi \mathcal{N})(t, \mathbf{x}) d\mathbf{x} \right)^{\frac{1}{2}}. \quad (13)$$

An important issue is to insure that when $\|\cdot\|_\chi$ is discretized, one still has the properties of a norm, e.g. definite positive and Cauchy-Schwarz inequality. It is in general trivial (e.g. $L^2 \rightarrow l^2$) but it can be more tricky if χ is complicated: a nontrivial case is discussed in subsection 3.6. We now describe the other functionals in the next subsections.

3.2. Boundary value ansatz

The cost functional (12) must take into account the boundary value (BV) problem, namely that $\mathcal{N}_u(0, \mathbf{x}) = a(\mathbf{x})$ and $\mathcal{N}_u(1, \mathbf{x}) = b(\mathbf{x})$ where a, b are given. There are two strategies: either imposing these constraints explicitly by penalisation or considering some ansatz for \mathcal{N}_u which automatically include them. The second choice means that one uses the general ansatz

$$\mathcal{U}(t, \mathbf{x}; \theta_u) = \Lambda_a(t)a(\mathbf{x}) + \Lambda_b(t)b(\mathbf{x}) + \Lambda_u(t)\mathcal{N}_u(t, \mathbf{x}; \theta_u), \quad (14)$$

where $\Lambda_a(0) = 1$, $\Lambda_a(1) = 0$, $\Lambda_b(0) = 0$, $\Lambda_b(1) = 1$, $\Lambda_u(0) = \Lambda_u(1) = 0$. In addition, the zeros of the functions Λ must be only those required at the boundaries $\tau = 0, \tau = 1$.

We then replace \mathcal{N}_u in (12) by \mathcal{U} instead whenever it is explicitly needed. In practice, we use $\Lambda_a(t) = 1 - t$, $\Lambda_b(t) = t$ and $\Lambda_u(t) = t(1 - t)$ mimicking a double Taylor expansion. We do not claim for optimal decision here, as many other choices would work. This one gives very good results in all the situations we have met. It is preferred over the first approach (BV penalization) when it is difficult for the system to relax on a and b .

3.3. Arclength condition

Although the geometrical action does not depend on the time parametrization chosen due to its homogeneity, it is important in practice to restrict the problem. As a matter of fact, in the cost functional landscape, there is an infinity of possible solutions, each having its own parametrization. Fixing an arclength condition, say $\|\dot{u}\| = c, \forall s \in [0, 1]$ is just a matter of convenience. But more importantly, it stabilizes the gradient search preventing the NNs to drift towards ill-conditioned parametrizations, especially in the context of PDEs. The use of a small penalisation parameter $\gamma_{\text{arc}} \ll 1$ is enough to prevent NNs to explore extreme landscape regions. The penalty constraint after straightforward algebra is

$$\mathcal{C}_{\text{arc}}[\theta_u] = \int_0^1 \|\dot{u}\|^2 ds - \left(\int_0^1 \|\dot{u}\| ds \right)^2 \geq 0, \quad (15)$$

where \mathcal{U} takes the form (14). Note that the norm used is user-defined rather than the actual χ or χ^{-1} norm.

3.4. Additional constraints

We now discuss the constraints: the ones associated with the auxiliary fields v, w and external ones which might be needed in some specific cases when looking for minimizers in subspaces of \mathcal{C} . They are, from (11),

$$\mathcal{C}_{v,w} = \gamma_v \int_0^1 \|\chi \mathcal{N}_v - \dot{\mathcal{U}}\|^2 d\tau + \gamma_w \int_0^1 \|\chi \mathcal{N}_w - F(\mathcal{U})\|^2 d\tau. \quad (16)$$

The spatial derivatives involved in F (or possibly in the operator χ) and the derivative $\dot{\mathcal{U}}$ can be obtained either using an efficient automatic differentiation (AD) algorithm or using finite-difference schemes with a very small increment (consistent with the scheme order and the machine precision, see PINNs methods).

Other constraints might be needed. For instance, in some situations (see examples 4.2 and 5.2,5.3), one must fix a total length L for the instanton path, in this case, one uses

$$\mathcal{C}_L[\theta_u] = \gamma_L \left(\int_0^1 \|\dot{\mathcal{U}}\| ds - L \right)^2, \quad (17)$$

where L is fixed by the user. Other typical situations are problems having (discrete) symmetry groups.

One can easily constrain the NNs to satisfy or break some symmetry. We do not write the penalty terms as they are problem-dependent (see an example (34)). This is a very useful technique to detect candidates for global minimizers (see subsection 5.3).

3.5. Boundary conditions (PDEs)

For PDE problems, one must impose some boundary conditions on $]0, 1[\times \partial\mathcal{D}$. Like in subsection 3.2 there are two choices: either penalizing the boundary constraints or finding some ad-hoc ansatz. Both strategies often yield equivalent results unless more precision is required. We discuss here only the case of periodic boundary conditions and propose a new approach which gives very good results. It is an exact approach for C^k -periodic solutions, ie for spatial derivatives up to order k . One of the reasons to stop at some given k is mainly practical: most often constraining C^∞ regularity for NNs is both useless and impossible (unless using very constrained NN architectures [15]). In the examples considered $k = 2$ or $k = 3$ is in fact enough. Diminishing the regularity (C^0 or C^1) is not advised however, having visible impacts on the class of minimizers obtained. The method relies on Hermite interpolation formula. We give here the ansatz formula for $k = 2$ in 1-D only with $\mathcal{D} = (0, 1)$ for illustration.

$$\mathcal{N}_{\text{per}}(\tau, x) = \mathcal{C}_0 + \mathcal{C}_1 x + \frac{1}{2} \mathcal{C}_2 x^2 - (\mathcal{C}_1 + \frac{1}{2} \mathcal{C}_2) x^3 + (3\mathcal{C}_1 + \frac{1}{2} \mathcal{C}_2) x^3 (x-1) - 6\mathcal{C}_1 x^3 (x-1)^2 + x^3 (x-1)^3 \mathcal{N}(\tau, x). \quad (18)$$

Here $\mathcal{C}_{0,1,2} : \tau \mapsto \mathcal{C}_{0,1,2}(\tau)$ are three independent NNs with input dimension equals to $d = 1$. For instance, one can easily check that $\mathcal{N}_{\text{per}}(\tau, 0) = \mathcal{N}_{\text{per}}(\tau, 1) = \mathcal{C}_0(\tau)$ for all τ . Such ansatz must be plugged in the cost functional whenever \mathcal{N}_{per} is needed and the three (small) additional NNs have their own parameters to optimize with the other parameters during the full minimisation descent.

3.6. Monte-Carlo sampling, stochastic and adiabatic gradient descent

We describe here rather standard procedures to solve the minimisation problem (12) as well as useful tips.

3.6.1. Monte-Carlo estimates

The cost functional (11) must be discretized. We follow the classical machine-learning strategy which consists in using simple Monte-Carlo estimates. The basic idea is to draw a given number of random points in space from some ad-hoc probability distribution. The set of points obtained is often called *batch*. There are important reasons for doing so. First, Monte-Carlo integrals do not suffer from the curse of dimensionality so that one can easily handle higher dimensional integrals. Second, it is a very efficient way to deal with the strongly nonconvex behavior of the NN optimisation problem. Due to that, the optimisation descent problem and NN parametrizations must be considered as stochastic rather than deterministic. In our case, there is a particularity owing to the nonquadratic structure of the geometrical action itself. We use here uniformly-distributed tensorial batches:

$$(\tau_i, x_j) \sim \mathcal{U}(0, 1) \times \mathcal{U}(\mathcal{D}), \quad i = 1, \dots, N_\tau, \quad j = 1, \dots, N_x. \quad (19)$$

The integrals are then simple empirical means:

$$\int_0^1 \|f\|_\chi d\tau \approx \frac{\gamma_{\mathcal{D}}}{N_\tau} \sum_{i=1}^{N_\tau} \left(\frac{1}{N_x} \sum_{j=1}^{N_x} f(\tau_i, \mathbf{x}_j)(\chi f)(\tau_i, \mathbf{x}_j) \right)^{\frac{1}{2}}, \quad (20)$$

where $\gamma_{\mathcal{D}}$ is a renormalisation factor (here the square root of the volume $\gamma_{\mathcal{D}} = \sqrt{|\mathcal{D}|}$). The auxiliary constraints (16) can be handled either using the same batches (19) or using spacetime samplings:

$$(\tau_i, \mathbf{x}_i) \sim \mathcal{U}((0, 1) \times \mathcal{D}), \quad i = 1, \dots, N_{\tau x}.$$

Other strategies with better precision might be considered as well, such as cubature formulas and quasi-Monte-Carlo methods e.g. [12].

When monitoring the action values, there are two possible procedures. If the dimension d is large, one can either compute the mean of the action estimate during the stochastic gradient descent (see below) over a sufficiently large number of iterations, or use a single very large batch. In moderate dimensions like in this work $d \leq 3$, it is often better to use Gaussian quadrature points or even Simpson rule on regular spacetime $d+1$ -dimensional grids and to perform a mean estimate over few iterations. We emphasize that very precise action estimates must not be part of the minimisation algorithm itself, it is required in general at the very end.

3.6.2. Stochastic gradient descent

Although widely used in machine-learning, we recall the very basics. The search for optimal NN parameters $\boldsymbol{\theta}$ is performed through a stochastic gradient descent

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \nabla_{\boldsymbol{\theta}} \bar{\mathcal{C}}_N |_{\boldsymbol{\theta}_k}, \quad (21)$$

where $\bar{\mathcal{C}}_N$ is the Monte-Carlo estimate of \mathcal{C} (12) (for instance using formulas (20)) over a total of N random batch points, η is called the learning rate. The gradient of the discretized cost functional $\bar{\mathcal{C}}_N$ is estimated using backpropagation algorithms, more generally, automatic differentiation (AD). These computations must be fast and accurate. The AD algorithmic efficiency is indeed one of the main reason for the success of deep learning methods. The vector of parameters $\boldsymbol{\theta}$ is intrinsically random due to the use of batches and (21) might be interpreted as a random dynamical system. The larger the size of the batch N , the smaller the variance of the estimate $\boldsymbol{\theta}$. More sophisticated first-order gradient descent are often used, in particular the so-called ADAM descent algorithm (see [39]). This is the default choice in this work. The use of second-order deterministic quasi-Newton methods although superior, such as L-BFGS should be taken cautiously: they are costly and should be used in deterministic context instead with careful initialisation procedure. The minimisation can be decomposed into two steps. First, a stochastic gradient descent (ADAM) is performed until the networks become statistically stationary (if it is computationally affordable). Second, another descent can be performed to improve the precision by using a smaller learning rate, larger batches, or switching to deterministic L-BFGS.

3.6.3. Adiabatic descent

An important issue in machine learning is the initialisation of the NN parameters. By default, it is random but requires a flux balance among different layers. A well-known procedure is the Xavier initialisation [24] and variants which are often tightened to the choice of the activation functions. Very often, the convergence to relevant regions of the cost landscape is in fact very slow, or there is no convergence at all. This situation happens either when the optimisation setup is highly penalized (see 3.4 with $\gamma_v, \gamma_w \gg 1$) or when the physical problem itself becomes ill-conditioned in some limit of interest, e.g. small diffusion (see an example in 5.3), or extreme condition (see section 6). A very powerful approach is to slowly perturb the problem from a situation where it is easy to solve to a very stiff configuration. We call this an *adiabatic descent* since only one stochastic descent is performed. We give the simplified algorithmic formulation for clarity. Let μ be an external parameter (a physical one, some penalisation parameter/hyperparameter) and θ the NN parameters to optimize.

- (i) Choose $\mu = \mu_0$ such that the problem is easy to solve, μ_{end} a target problem and $d\mu$ an adiabatic speed. Choose some incremental law $\mu \mapsto S(\mu, d\mu)$, e.g. a linear one $S(\mu, d\mu) = \mu + d\mu$.
- (ii) Perform a stochastic gradient descent for $\mu = \mu_0$ until satisfactory (statistical) convergence.
- (iii) Do $k \rightarrow k + 1$ until $\mu_k = \mu_{\text{end}}$

$$\begin{cases} \mu_k & = S(\mu_{k-1}, d\mu) \\ \theta_{k+1} & = \theta_k - \eta \nabla_{\theta} \bar{\mathcal{C}}_N(\mu_k)|_{\theta_k}, \end{cases}$$

Some tuning is required here especially in the choice of the adiabatic speed $d\mu$ compared to the inertia of the NN considered (typically the training rate): the change of μ must be small enough so that the NN is able to relax to the new configuration. In practice, doing so enable the NN to better learn the slow configuration changes as the weights and bias adjust more efficiently from a near optimal situation. One must not try to perform a converged stochastic descent at each step since it would be too costly. A remarkable aspect is that the NN size can be chosen rather small. This points towards the fact that the minimisation problem is most often the issue rather than the expressivity of the network to capture stiffer solutions. A typical example is shown in subsection 4.1.

The stiffer the configuration the smaller $d\mu$ must be. Therefore, one might often consider non-linear S (e.g. a logarithmic change in the context of large deviations). Finally, we mention that when μ is a physical parameter, it is important to perform an adiabatic search back and forth, ie $\mu_0 \rightarrow \mu_{\text{end}} \rightarrow \mu_0$ in order to detect possible hysteresis phenomena. Note also that this strategy is not very different, although more general, from *scheduling* (and to some extent *early stopping*) which amounts to control the way the learning rate η changes during the descent.

3.7. Comparison with classical gMAM

We now discuss an important aspect: how does deep gMAM compare to the classical gMAM algorithm ? The gMAM algorithm was first introduced in [33] and used many times, e.g. [34, 66, 50]. Although, the minimum geometrical action problem is exactly the same, the way it is solved is very different. For convenience denote $\mathcal{A}[u]$ the problem (9) to minimize. Classical gMAM relies on finding the zeros of the so-called Euler-Lagrange (EL) equations which are formally $\delta\mathcal{A}[u] = \langle \text{EL}[u], \delta u \rangle$.

For clarity, the FW action in its original form would give $\delta\mathcal{A} = \int_0^T \langle \dot{u} - F(u), \delta\dot{u} \rangle dt + \int_0^T \langle \dot{u} - F(u), DF_u \delta u \rangle dt$. It is already clear that in order to obtain some explicit EL, one is forced to integrate by parts, as it is usually done. Here it would be $EL = -\ddot{u} + DF_u \dot{u} + DF_u^*(\dot{u} - F)$ where DF^* is the adjoint operator. The classical gMAM approach therefore consists in solving $EL = 0$ using the explicit form of the EL functional derivative. A simple way for doing this, is to consider a relaxation approach, by introducing a fictitious time t and solving a preconditioned version:

$$\partial_t u = -\gamma_u EL[u]. \quad (22)$$

The discretization is in general in semi-implicit form. The term $\gamma_u = \|F(u)\|/\|\dot{u}\|$ is the ad-hoc preconditioning quantity. There is in addition a technical procedure that we do not detail here, see [33, 34] where one must insure a constant arclength through interpolation every step or so. If not the (deterministic) gradient descent becomes ill-conditioned. An important aspect of this approach is the solution representation itself: $u(\tau, \mathbf{x})$ must be discretized on a structured grid or using a Galerkin basis.

We can identify at least three important advantages of deep gMAM over the classical one.

- First, the NN parametrization do not suffer from the curse of dimensionality with some outstanding expressivity of complex nonlinearities even in high dimension. Classical Galerkin decompositions are severely constrained by the curse of dimensionality. The minimisation problem (12) can be handled in larger dimension d without changing much the NN size. Classical gMAM might involve very heavy optimisation problems since the size of the grid increases as $O(1/\text{precision}^{d+1})$. The number of NN parameters needed are typically known to behave polynomially with the dimension instead (see e.g. [2]) and is the main reason for the great success of deep learning.
- Second, at the level of EL, one can anticipate that the integration by parts yields differential operators of order $2n$ whenever F contains differential operators of order n . Therefore, one expects that the integration of (22) would then be tightened to more stringent CFL conditions. In the machine-learning context, EL equations are found automatically (AD, backpropagation) in the neural network space, ie without having to integrate by part, see (21).
- Third, the stochastic aspect of the approach brings unexpected flexibility and efficiency. One can easily escape from local minima through stochastic fluctuations. In the deterministic gMAM, being trapped into a local minimum requires an external action. The arclength condition is handled very easily at no cost and in fact the soft condition $\gamma_{\text{arc}} \ll 1$ helps the NN to explore better-adapted parametrizations. Any additional constraints (see section 3.4) are straightforward to implement by simply adding more penalisation terms. This is not the case when dealing with an explicit EL approach, where it is often unclear how to constrain the problem in a simple way.
- Another aspect which is worth to mention, is that the differential operators are in fact exact approximations in the context of NN parametrization, whereas classical discretizations relies on finite-order approximations schemes.

The main drawback of deep gMAM is the precision. Once classical gMAM converges, it has better precision than deep gMAM. The relative error typically encountered is $O(1\%)$. This is in fact a much more general problem in deep learning, it is very difficult, although not impossible to attain the same precision than classical methods. In general, costly procedures are needed to improve this. For instance, one can increase the NN capacity, fine tune the hyperparameters, or make use of quasi-Newton methods. It also includes the preconditioning of the cost functional in the neural network space: this is a well-known and very difficult problem in machine learning. It appears, however, that one is often more interested in the generalisation behavior of NNs than the precision itself, especially when the dimension increases in which case classical approaches become hopeless anyway. It is worth mentioning promising approaches to cope with the precision issue e.g. multi-scaled NNs [7, 42] and domain decomposition [46]. We summarize these remarks in the table below.

	gMAM	Deep gMAM
Discretisation	finite-diff, pseudo-spectral,...	NN parameter θ (exact)
Descent vector	$-\frac{\ F(\mathbf{u})\ }{\ \dot{\mathbf{u}}\ } \times \text{EL, Euler-Lagrange} \equiv \frac{\delta \mathcal{A}_g}{\delta u}$	$-\nabla_{\theta}(\mathcal{A}_g + \mathcal{C}_{\text{arc}} + \dots)$
Descent algor.	deterministic: relaxation, CG, L-BFGS	stochastic: ADAM and variants
Arclength	hard constraint $\ \dot{\mathbf{u}}\ = \text{cst}$ (interpolation)	none (soft constraint)
Dim curse ($d \geq 3$)	Yes	No
Non-convexity	Moderate (\sim continuous action)	Wild (NN par.)
Flexibility	low	high
Precision	high	low

4. Examples: SDEs

Although deep gMAM offers significant advantages in the context of PDEs $d \geq 1$, it also remains highly competitive in the SDEs situation $d = 0$ as well. The examples below are rather illustrative of the different scenarios one can encounter in the context of PDEs. They should not be considered as a finality on their own. We illustrate two important classes of systems (see also Appendix 8.2). The ones where the dynamics has a free-energy potential, the so-called *gradient systems* and the others where there is no potential, ie the *non-gradient systems* where the time-reversal symmetry is broken. In the first case, Arrhenius laws can be derived explicitly. In particular, it is possible to show that the instanton trajectory must cross the saddle points (ie with Morse index one) and orthogonal to the potential isosurfaces. Moreover, the instanton going from a to b has the same geometrical support than the one going from b to a owing to the time-reversal symmetry of the stochastic process. Only the value of the action may differ.

Nongradient systems offer much richer and unexplored behavior. In particular, instanton paths do not need to cross saddle points in the separatrix. One can nevertheless recover an explicit Arrhenius law in the hypothesis of transversality [66]. In fact, we give an example in subsection 4.2 where the separatrix contains no saddle points. The non-equilibrium trajectory therefore must go through a more complicated object playing the role of the saddle.

4.1. Gradient case: the (rugged-)Müller potential

The first example is the well-known Müller potential with two degrees of freedom (d.o.f.) [47]. It is a gradient system $\mathbf{F} = -\nabla V$ with

$$V = \sum_{i=1}^4 A_i \exp \{ \alpha_i (x - X_i)^2 + \beta_i (x - X_i)(y - Y_i) + \gamma_i (y - Y_i)^2 \},$$

where $A = (-200, -100, -170, 15)$, $\alpha = (-1, -1, -6.5, 0.7)$, $\beta = (0, 0, 11, 0.6)$, $\gamma = (-10, -10, -6.5, 0.7)$, $X = (1, 0, -0.5, -1)$, $Y = (0, 0.5, 1.5, 1)$. There are three minima $\mathbf{m}_1 = (-0.5582, 1.4417)$, $\mathbf{m}_2 = (0.6235, 0.0280)$ and $\mathbf{m}_3 = (-0.050, 0.4667)$ and two saddles $\mathbf{s}_1 = (-0.8220, 0.6243)$, $\mathbf{s}_2 = (0.2124, 0.2929)$. We focus here on the transition $\mathbf{a} = \mathbf{m}_1 \rightarrow \mathbf{b} = \mathbf{m}_3$. We call $\mathbf{a} = \mathbf{m}_1$ and $\mathbf{b} = \mathbf{m}_3$. The Arrhenius law gives an action

$$\mathcal{A}_{\mathbf{a} \rightarrow \mathbf{b}} = 2(V(\mathbf{s}_1) - V(\mathbf{a}) + V(\mathbf{s}_2) - V(\mathbf{m}_3)) \approx 229.1$$

We penalize here the BV constraints $u(0) = \mathbf{a}$ and $u(1) = \mathbf{b}$, the other approach described in subsection (3.2) gives identical results. The details given below are just illustrative: in fact any NN configuration would work. In other words, no fine tuning is required here.

We now minimize the geometrical action using two NNs with 4 hidden layers having 10 neurons each, and one for each component (x, y) . One can use a single NN having a 2-dimensional vector input yielding the same result. We impose here a large penalty on the boundary conditions with $\gamma_{bv} = 1000$ and $N_\tau = 512$. Much smaller $N_\tau = O(10)$ would also work. Although not necessary, we also use a learning rate scheduling by decreasing it from 10^{-2} to 10^{-4} in a log linear way over 20000 iterations. We recall that a single iteration is a random shuffle of N batch points uniformly distributed in $(0, 1)$. We obtain the mean value over the last ten iterations $\mathcal{C}_{\mathbf{a} \rightarrow \mathbf{b}} \approx 229.3$, ie a fairly close value to the Arrhenius estimate. Note that closer values can be obtained if the batch size and the NN size is increased. The instanton path which is a minimum energy path in this case compares very well with known results [47, 68].

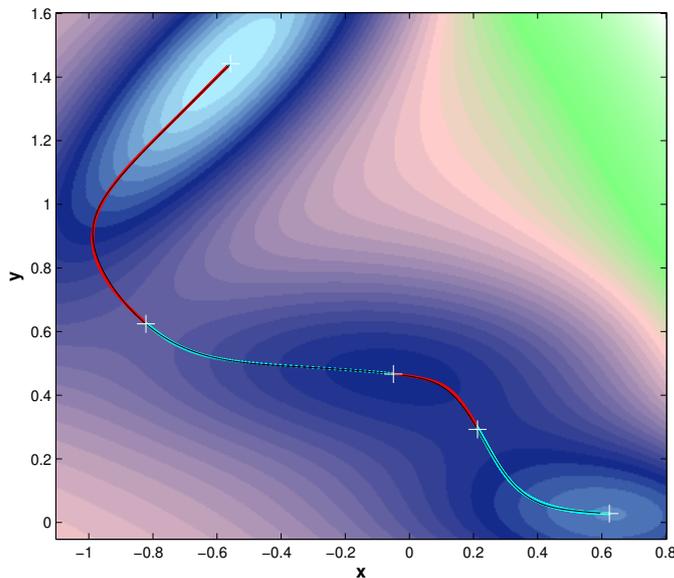


Figure 1: Instanton $\mathbf{a} \rightarrow \mathbf{b}$ obtained using a feedforward NN with 4 hidden layers having 10 neurons in each layers and swish activation function. Superimposed is the solution obtained from the classical gMAM algorithm (black curve). The red parts of the curve correspond to an the action above 0.2 (reactive path) and the complementary parts where the action is below this threshold (relaxation path). The white crosses are the minima and saddles and the Müller potential is the colored background.

We now consider a more challenging situation by considering the so-called "Rugged Müller" version. One adds to the potential a high-frequency perturbation:

$$V_{\text{rugged}} = V + \rho \sin 2\pi n x \sin 2\pi n y.$$

We choose here $\rho = 3$. We are in fact able to handle larger values $\rho \leq 10$ in a similar way. The potential landscape V_{rugged} has now many saddles: there are thus many local action mimimizers and the total number of the action functional minima in fact explodes exponentially with the number of saddles of V_{rugged} . Although $d = 0$ with only two degrees of freedom, the difficulty to capture the correct global minimizer is much more challenging than gradient PDEs having much less saddles. We illustrate the adiabatic descent strategy described in subsection 3.6.3 in Fig. 2. The efficiency of the approach is rather striking, starting from random initial conditions turns out to be very difficult whereas the adiabatic technique in cheap and easy. We believe that the red curve in Fig. 2 is in fact the global energy minimizer, although no known method is able to confirm it is actually the case. The classical gMAM method fails at giving lower-energy states than the ones obtained here. Note that finding low-energy action curves for this example is in fact physically irrelevant and other actions are required (E.Vanden Eijnden, private comm.). We therefore consider this example only as an algorithmic illustration.

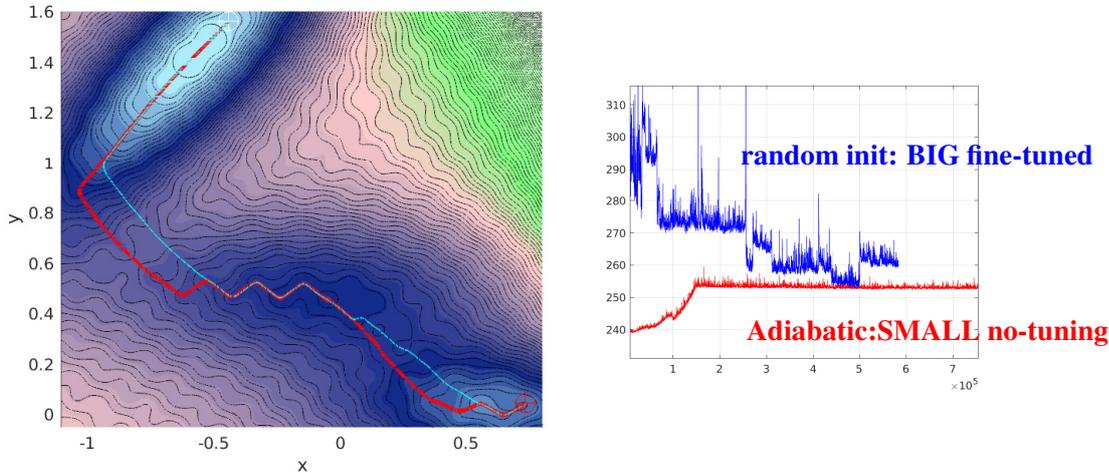


Figure 2: Adiabatic method for ρ : in red we use a medium size NN with 8 hidden layers with 10 neurons capacity and swish activations giving a total of 922 NN parameters and we slowly change ρ from 0 to 3 over $1.5 \cdot 10^5$ iterations (red curve right panel) and then froze $\rho = 3$. In blue, we purposely use a much bigger NN having 20 hidden layers with skip connections giving a total of 11220 parameters and start from a random NN initialisation (blue curve right panel). In this case ρ is fixed $\rho = \rho_{\text{end}} = 3$. In this example, we use formulation (14) and $N_\tau = 256$.

4.2. Nongradient case: hyperbolic limit cycle in the separatrix

We now consider a striking example introduced in [66] where there is no saddle point in the separatrix. It is a 3 degrees of freedom (dof) nongradient system with

$$\begin{cases} F_1(x, y, z) &= -(z+1)(z-2) \frac{x}{(x^4+y^4)^{1/4}} - x - y^3 \\ F_2(x, y, z) &= -(z+1)(z-2) \frac{y}{(x^4+y^4)^{1/4}} + x^3 - y \\ F_3(x, y, z) &= -(z+1)(z-2)z \end{cases} \quad (23)$$

This system has two stable fixed points $\mathbf{a} = (0, 0, -1)$ and $\mathbf{b} = (0, 0, +2)$ and a hyperbolic limit cycle Γ on the separatrix at $z = 0$ satisfying $x^4 + y^4 = 16$ with $\dot{x} = -y^3$ and $\dot{y} = x^3$. This limit cycle is stable on the separatrix $z = 0$ and unstable in the z direction. This system has the so-called transverse (or orthogonal) decomposition

$$\mathbf{F} = -\nabla V + \mathbf{b}, \quad \nabla V \cdot \mathbf{b} = 0, \quad \forall \mathbf{x} \in \mathbb{R}^3 \quad (24)$$

with $\mathbf{b} = (F_1, F_2, 0)$ and $V(x, y, z) = \frac{z^4}{4} - \frac{z^3}{3} - z^2$. Due to that, it is possible to show that the action value for the instanton going from \mathbf{a} to \mathbf{b} follows the classical Arrhenius formula ([66], Theorem 1)

$$\mathcal{A}_{\mathbf{a} \rightarrow \mathbf{b}} = 2(V(\mathbf{x}_s) - V(\mathbf{a})).$$

where \mathbf{x}_s is any point belonging to the limit cycle Γ . In this case the potential depends on z only so that the minimum action is $\mathcal{C}_{\mathbf{a} \rightarrow \mathbf{b}} = \frac{5}{6}$. Note that the path $(0, 0, -1 + 3s)$, $s \in (0, 1)$ connecting directly \mathbf{a} to \mathbf{b} is not a valid solution, although it has the minimal action value $5/6$, this path indeed coincides with the flow singularity $x = y = 0$.

It turns out that it takes an infinite time for the instanton to approach the separatrix. Due to the rotation, it winds around the limit cycle an infinite number of time. There are two consequences. First, the instanton length must be infinite. Second, there are actually an infinite number of instantons having the same action value which differ by a phase. It contrasts with the classical scenario where there is a unique instanton with finite length going through a saddle point. A sketch of this situation is shown in Fig. 3

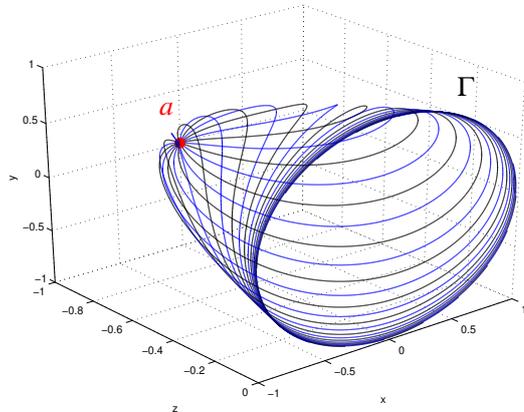


Figure 3: There is a continuum of instantons connecting a to the hyperbolic limit cycle Γ on the separatrix ($z = 0$), all of them have the same action and are of infinite length. This example is extracted from the system $\mathbf{F} = (1 - z^2 - r, 1, z - z^3)$ in cylindrical coordinates (r, θ, z) (see [66]). Only two instantons (blue and black curves) are shown $x(t) = r(t) \cos(t + \varphi)$, $y(t) = r(t) \sin(t + \varphi)$ and $z(t)$ with $r(0) = 0$, $z(0) = -1$ and $r(\infty) = 1, z(\infty) = 0$ and φ is arbitrary.

We now show the results obtained by the deep gMAM algorithm in such a case. Figure 4 corresponds to two different experiments. The first one does not impose any condition on the arclength so that the NNs choose their own parametrization. In the second experiment, we purposely impose a longer arclength for the solution (here $L = 50$ versus $L \approx 17.7$ previously) by adding the constraint (17).

We recover the fact (see Fig. 4) that when we allow for longer arclength instanton, we obtain a solution winding around the hyperbolic limit cycle Γ like in Fig. 3. The obtained values for the action is fairly close to the theoretical one in all the experiments (see Fig. 4). The action value between the two cases are very close: the winding around the limit cycle contributes very little to the total action. Note that in this case, no fine tuning of the hyperparameters is required and smaller networks would work well.

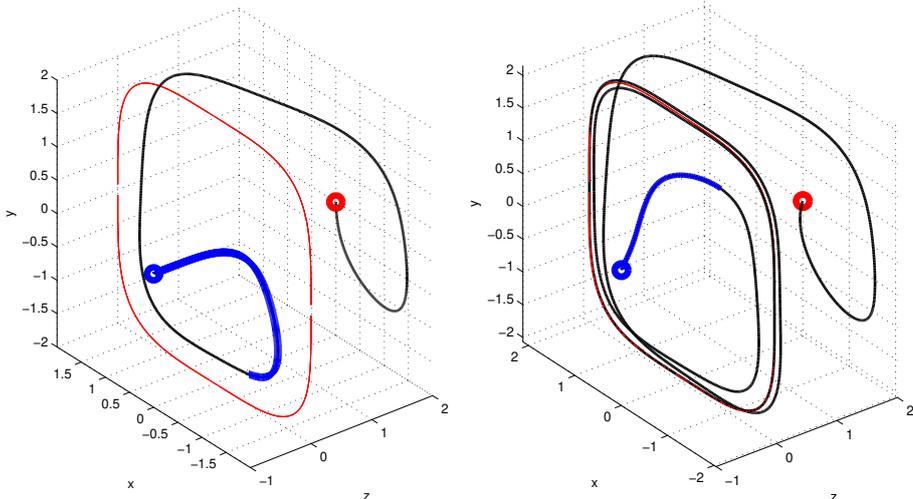


Figure 4: Instantons connecting \mathbf{a} (blue dot) to \mathbf{b} (red dot) for the flow (23) using a NN with no arclength condition (left panel) and an imposed arclength $L = 50$ (right panel). The left panel gives an arclength $L \approx 17.7$. The hyperbolic limit cycle is shown in red in the $z = 0$ separatrix hyperplane. The thick blue part of the curve is where the action is largest. The total action is $\mathcal{A}_{\mathbf{a} \rightarrow \mathbf{b}} \approx 0.838$ (left panel) and ≈ 0.837 (right panel) (to be compared with 0.83333). NN size is 14 hidden layers with 10 neurons for each layer using swish activation functions. The descent algorithm is ADAM with learning parameter $\alpha = 10^{-3}$ and batch size $N = 512$. The penalty parameter $\gamma_L = 1$ for the right-panel experiment and $\gamma_{bv} = 10$ for both.

4.3. SDE quasi-potential: Maier-Stein system

We illustrate here as a byproduct of the method an efficient and simple way to compute the quasi-potential (see definition (5)). Consider an ODE having N d.o.fs $\dot{u} = F(u)$, for simplicity we assume $\chi = Id$. We want to compute (slices of) the quasi-potential V in some p -dimensional subspace $b \in \mathcal{D}_b \subset \mathbb{R}^N$ and for a fixed:

$$V : \mathbb{R}^N \rightarrow \mathbb{R}, b \mapsto V(a, b) = \inf_{u \in \mathcal{G}_{a,b}} \int_0^1 (|\dot{u}| \|F(u)\| - \langle \dot{u}, F(u) \rangle) d\tau,$$

where $\mathcal{G}_{a,b} = \{u(\tau) \in \mathbb{R}^N, u(0) = a, u(1) = b\}$. The domain \mathcal{D}_b can be anything. When $p = N$, one considers the full quasi-potential and when $p < N$ only lower-dimensional slices of V . Deep gMAM offers a very simple way to handle this problem by considering a larger network:

$$\mathcal{U}(\tau, b) : \mathbb{R}^{p+1} \rightarrow \mathbb{R}^N, p \leq N$$

and minimizing the sum over b of all geometrical actions:

$$\mathcal{C}[\mathcal{U}] = \int_{\mathcal{D}_b} db \int_0^1 (|\dot{\mathcal{U}}| \|F(\mathcal{U})\| - \langle \dot{\mathcal{U}}, F(\mathcal{U}) \rangle) d\tau. \quad (25)$$

This means that the neural network is encoding all the instantons at once and provides an estimate of the quasi-potential by summing over τ , ie evaluating the geometrical action at prescribed points once the minimisation has converged. Classical approaches e.g. [8] rely on solving the stationary Hamilton-Jacobi PDE: $2\nabla V \cdot F + |\nabla V|^2 = 0, V(a) = 0$. It is a notoriously difficult problem with multiple (viscosity) solutions and requires careful numerical strategies. In our case, there is no particular treatment. It moreover offers the advantage to be a purely *local* approach: there is no constraints on \mathcal{D}_b . It is only limited by the NN capacity to represent high-dimensional ending conditions b .

In practice, it is better considering a BV ansatz (14) so that the procedure is transparent and does not need to monitor some BV penalisations. One needs to generate batches in the space $(0, 1) \times \mathcal{D}_b$. In principle, one must build tensorial batches like in (19) of size $N_\tau \times N_b$. We show below that it is perfectly viable to consider $N_\tau = 1$ with consequence that the arclength condition (15) is by definition zero. We therefore remove the arclength condition in the example below. The example is a well-known testbed for nongradient SDEs called the Maier-Stein system [45].

$$F(u) = (u_1 - u_1^3 - 10u_1u_2^2, -(1 + u_1^2)u_2).$$

We fix $a = (-1, 0)$ and $p = N$. The corresponding snippet Julia code is given in the Appendix 8.2. Note that the quasi-potential V is nondifferentiable along the x axis which makes the convergence more difficult. In general, one has nonuniform convergence of deep gMAM when the domain is large. When the domain is small, the NN imposes on the contrary a spatially coherent global structure. It means that the various instantons tend to share similar geometrical paths at the qualitative level. It is therefore a good practice to consider large enough NNs in order to detect possible nondifferentiable structures correctly.

5. Examples: SPDEs

The deep gMAM algorithm offers an even more interesting alternative in the PDE case where the trajectory now lives in an infinite-dimensional functional space instead. A striking property is that one can use the same NNs used for SDE problems to handle SPDEs problem: only the input dimension $d + 1$ must be modified. It thus suggests that NNs are very efficient to detect low-dimensional manifolds in very large-dimensional spaces. We revisit some well-known examples of increasing difficulty starting from $d = 1$ (gradient PDE) to $d = 3$ (nongradient PDE).

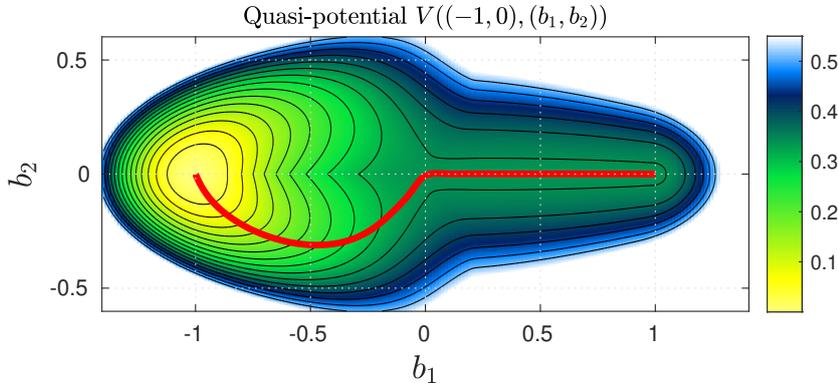


Figure 5: Quasi-potential for the Maier-Stein system [45] using the snippet code 8.2 with 10 hidden layers (15 neurons/layer, swish), with $N_b = 1000$, $N_\tau = 1$. Training rate is 10^{-3} . The red curve is one of the two instantons connecting $(-1, 0)$ to $(+1, 0)$ (the other is found by the symmetry $b_2 \rightarrow -b_2$).

5.1. Gradient case: Ginzburg-Landau

The Allen-Cahn-Ginzburg-Landau (ACGL) equation is very well documented, starting from the seminal work of [18]. It has been used many times to test ideas for computing non-equilibrium transitions due to its relative simplicity, e.g. in [16, 55]. The d -dimensional version with Dirichlet boundary conditions reads

$$\partial_t u = F(u), \text{ with } F(u) \equiv \nu \Delta u - P'(u), \quad u = 0 \in \partial \mathcal{D} \quad (26)$$

in the domain $\mathcal{D} = (0, 1)^d$ and with

$$P(u) = \frac{1}{4}(1 - u^2)^2, \quad P'(u) = u^3 - u.$$

This PDE is in fact a gradient system which can be written as $\partial_t u = -\frac{\delta V}{\delta u}$, with the free energy potential

$$V(u) = \frac{1}{2} \int_{\mathcal{D}} (\nu |\nabla u|^2 + 2P(u)) \, d\mathbf{x}. \quad (27)$$

We use here an approximated formula of the stationary solutions, namely we use

$$a(\mathbf{x}) = \prod_{l=1}^d \tanh \frac{x_l}{\sqrt{2\nu}} \tanh \frac{1-x_l}{\sqrt{2\nu}}, \quad \text{and } b(\mathbf{x}) = -a(\mathbf{x}).$$

This approximation is in fact very close to the true solution up to exponentially small correction terms. We want to compute the most probable transitions between the two metastable states a and b . This problem has several group symmetries: the symmetric group $\mathbf{x} \mapsto \sigma(\mathbf{x})$ where σ is a permutation, the reflection group $\mathbf{x} \mapsto (-1)^\sigma \mathbf{x}$ or the transformation mapping $u \rightarrow -u$. We thus expect that any solution is in fact associated with several (at least $2d!$ possibly identical) states and that many symmetry-breaking bifurcations can occur. Typically, the trivial state $u_{s,0} = 0$ is invariant by all symmetry groups, and the states a, b are invariant by all the permutations and reflections. Depending on the value of the diffusion parameter ν , the transitions $a \rightarrow b$ can differ. For instance, a first symmetry-breaking bifurcation occurs at a critical value ν^* for which a new family of critical points appears. In dimension $d = 1$, there are two of them which are antisymmetric w.r.t. the axis $x = \frac{1}{2}$. They behave as $u_{s,1} \sim \pm \tanh(x - \frac{1}{2})/\sqrt{2\epsilon}$ near $x = \frac{1}{2}$. In d dimension, we expect $2d$ saddle states each corresponding to an antisymmetric solution w.r.t. to the d hyperplane $x_l = \frac{1}{2}$. We omit the exact discussion on these group symmetries and focus first on $d = 1$. It can be shown that a cascade of bifurcations occurs at the critical values

$$\nu_n = \frac{1}{(n+1)^2 \pi^2},$$

where pair of states $\pm u_{s,n}$ appear and correspond to fronts having n zero in the interval $]0, 1[$. Since (26) has a gradient structure, one can use the Arrhenius formula:

$$\mathcal{A}_{a \rightarrow b}^{(n)} = 2(V(u_{s,n}) - V(a)). \quad (28)$$

Very precise estimates can be obtained provided ν is small enough. For $n = 1$, $V(u_{s,1}) = V(0) = \frac{1}{4}$. We can estimate also $V(a) \sim \frac{2\sqrt{2}}{3}\sqrt{\nu}$ using the fact that the solutions are ± 1 with boundary-layer corrections. An easy computation gives

$$\text{flip transition: } \mathcal{A}_{a \rightarrow b}^{(0)} \sim \frac{1}{2} - \frac{4}{3}\sqrt{2\nu} \quad (29)$$

The expression for $n > 1$ are different, one can remark that both a and the critical points have the same behavior at the boundaries so that their contribution cancel out. It only remains to compute the contributions of the zeros only which are $2n$ times those from the boundaries. It comes

$$n \text{ front transition: } \mathcal{A}_{a \rightarrow b}^{(n)} \sim \frac{4}{3}n\sqrt{2\nu}, \quad n \geq 1, \quad \nu \leq \nu_n. \quad (30)$$

The consequence is that the global minimizer can be easily distinguished, if $\nu > \frac{1}{4\pi^2}$, it is the flip solution and if $\nu < \frac{1}{4\pi^2}$ it is the $n = 1$ front transition. All other solutions are local minima having larger action.

We now test the deep gMAM algorithm. The results are shown in Fig. 6 and are in close agreement with the theoretical ones. Only the 1-front solution is the true global minimum having smaller action. We observe the typical reactive/relaxation parts for both solutions: positive action before reaching the separatrix and zero-action when it relaxes to the target solution b . Remarkably, we observe that for the front transition, the region around the saddle is very flat, confirming the analysis done in [55]: for finite-amplitude noise one indeed observes random-walk diffusion due to the landscape flatness.

5.2. Nongradient case: $d = 1$ Sheared Ginzburg-Landau

We now consider an example very similar to the previous one, except for two aspects: 1) it has an additional transport term $\partial_x u$ which makes it a nongradient version of the 1-D ACGL 2) we consider periodic boundary conditions. It is also taken from [66] and it corresponds to the same scenario than in 4.2 but in the context of PDE: there is an hyperbolic limit cycle inside the separatrix so that transitions can go through it rather than going through a saddle point. The PDE reads

$$\partial_t u = F(u), \quad F(u) = \nu \partial_{xx}^2 u + u - u^3 + 0.1 \partial_x u, \quad (31)$$

in the periodic domain $\mathcal{D} = (0, 1)$. It appears that this equation also satisfies the transverse decomposition (24) with the same potential (27) than for the ACGL equation. The orthogonal term b in (24) is simply $b(u) = 0.1 \partial_x u$. Any fixed points of the ACGL equation (ie with $c = 0$) are in fact limit cycles of (31), since $u(x, t) = v(x + 0.1t)$ satisfies $\partial_t u = F(u)$ ([66]). Due to that, one can obtain precise theoretical estimates of the possible transitions by using the same Arrhenius law (28) where $u_{s,n}$ are now periodic solutions, ie they satisfy

$$\nu \partial_{xx}^2 u + u - u^3 = 0, \quad u(x + 1) = u(x).$$

It turns out that the solutions are the same than before except that the solutions with n odd must be removed since they cannot fulfill the periodic condition. They can be studied in more details by interpreting space as time and analyzing the nonlinear oscillator $\nu \ddot{u} = u^3 - u$ with the constraint $u(t + 1) = u(t)$ (see e.g. a detailed explanation in [66]). We can therefore use the same formula (30) but for $n = 2p$, $p \geq 1$. In particular the mode with smallest action is now the mode with $n = 2$.

In summary, we have a classical saddle-point transition going through the trivial saddle $u_{s,0} = 0$ and the transitions going through hyperbolic limit cycles:

$$\text{flip transition: } \mathcal{A}_{a \rightarrow b}^{(0)} = \frac{1}{2}, \quad \text{LC } 2p\text{-transition: } \mathcal{A}_{a \rightarrow b}^{(p)} \sim \frac{8}{3}p\sqrt{2\nu}. \quad (32)$$

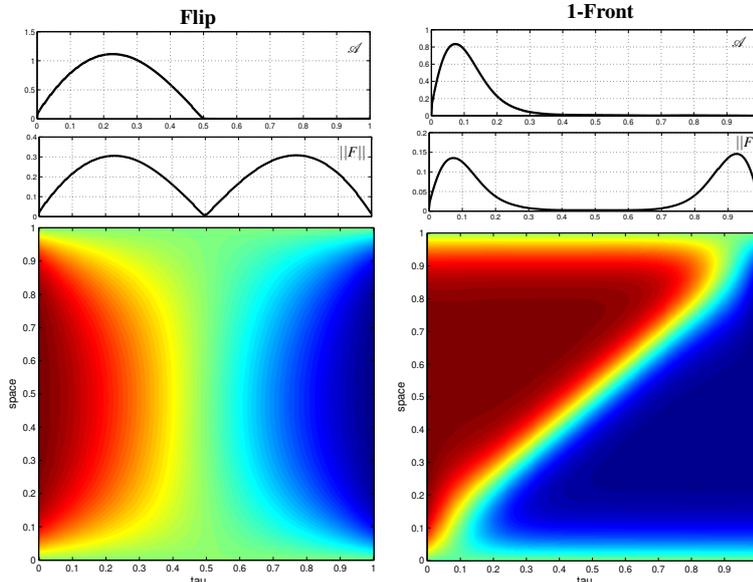


Figure 6: Two transitions obtained using 5 hidden layers with 10 neurons each with swish activations for $\nu = 5 \cdot 10^{-3}$: the flip transition (left) and the 1-front transition (right). The mean action obtained over the last 100 iterations is $\mathcal{A}_{a \rightarrow b}^{(0)} \approx 0.36$ (left) and $\mathcal{A}_{a \rightarrow b}^{(1)} \approx 0.13$ (right) to be compared with the theoretical estimates 0.366 and 0.133 resp. The value of the action and the spatial norm $\|F\|$ are shown as a function of the normalized arclength (ie computed only for visualisation) in the panels above. The penalty parameters are $\gamma_{bv} = \gamma_{bcs} = 10$, and $\gamma_{arc} = 0.1$. The batch size is $N_\tau \times N_x = 16 \times 64$ and ADAM learning rate $\alpha = 10^{-3}$. The two solutions have different total length: $L_{flip} \approx 1.8$ and $L_{1-front} \approx 3.1$.

Note that for the flip transition, the solution is in fact explicit, it takes the form $u(s, x) = -1 + 2s$, $s \in (0, 1)$ and it is easy to check that the geometrical action takes the value $\frac{1}{2}$ in this case.

We now describe the numerical results. They are obtained for $\nu = 5 \cdot 10^{-3}$. First of all, the trivial flip transition is easily obtained (not shown) and show the same type of behavior than in Fig. 6 (left) although now the solution is periodic in space. We are interested in the global minimizer instead whose action is in this case $8/3\sqrt{2\nu} \approx 0.266$ and from the discussion above, there is in fact a continuum of instantons of infinite length going through an hyperbolic limit cycle, like in Fig. 4. We thus conduct two different experiments where we impose an instanton length $L = 10$ and $L = 20$. The results shown in Fig. 7 strongly agree with the theoretical predictions. For instance, the norm $\|F\|$ does not vanish and is a very good indicator of the type of transition. This can be used for instance to decide whether one must impose some total length condition or not, since the consequence of such a behavior is the nonuniqueness of the action global minimizers. The larger L the closer the solution is to the actual infinite-length instanton. We find that the short length solution overestimates the (mean) action value (Fig. 7, left), whereas the longer one gives a smaller value more consistent with the theoretical value (Fig. 7, right).

5.3. Nongradient case: $d = 2$ Sheared Ginzburg-Landau

This problem corresponds to a 2-D sheared version of the Ginzburg-Landau equation originally studied in [9]. Action minimizers have already been investigated in [34] and [66] using the classical gMAM algorithm and can serve as useful benchmarks. This example gathers all the difficulties together but now in spatial dimension $d = 2$: it is a nongradient PDE having both hyperbolic limit cycles and saddle fixed points in the separatrix. It is moreover believed (see [66]) that this PDE cannot has the transverse decomposition (24). Indeed, an example can be found where two different instantons having different values of the action cross the same saddle fixed point, therefore contradicting Arrhenius formula (28). The equation reads

$$\partial_t u = F(u), \quad F(u) = \nu \Delta u + u - u^3 + 0.1 \sin 2\pi y \partial_x u, \quad (33)$$

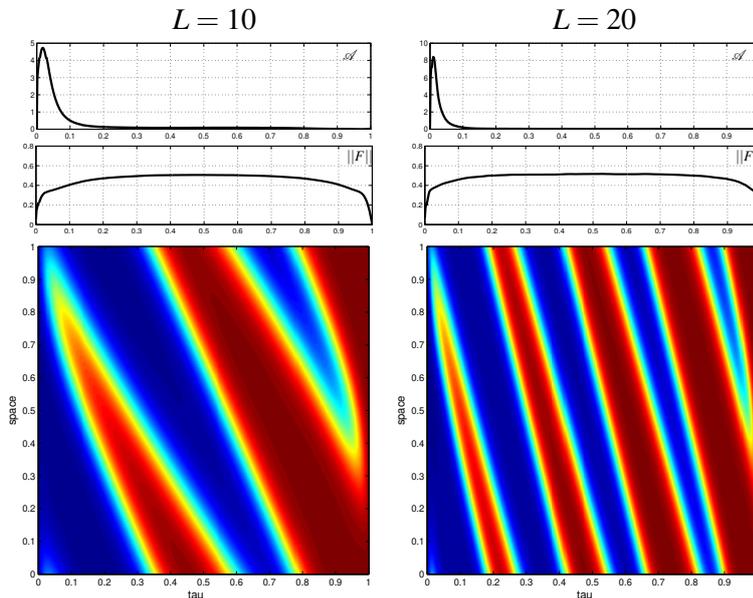


Figure 7: Instantons of different total length L and going through the $p = 1$ hyperbolic limit cycle (see text). The NN has 7 hidden layers with 10 neurons/layer with swish activations for $\nu = 5 \cdot 10^{-3}$. The value of the action and the spatial norm $\|F\|$ are shown as a function of the normalized arclength in the panels above. Action values are $\mathcal{A}_{a \rightarrow b} \approx 0.32$ (left), and $\mathcal{A}_{a \rightarrow b} \approx 0.28$ (right).

with $a = -1$, $b = +1$ as before, in the periodic domain $\mathcal{D} = (0, 1)^2$. The set of action minimizers contains in particular the solutions which do not depend on x , ie the ACGL 1-D solutions. A global minimizer candidate is therefore the one corresponding to (32) with $p = 1$ and going through the corresponding saddle-fixed point. It is shown in Fig. 8.

There is another solution with lower action which breaks the translational symmetry w.r.t. x . It was first detected in [34] and is shown in Fig. 9. To obtain this solution easily, we start from random initial conditions but use a “symmetry breaker” term by adding the penalty constraint in (12):

$$\gamma_{\text{sym}} \left(\int_0^1 \|\partial_x \mathcal{U}\|^2 d\tau - \text{cst} \right)^2. \quad (34)$$

The NN is forced to find solutions which are not invariant by x . The penalty coefficient γ_{sym} is then set to zero to verify that the solution obtained is stable for the NN representation.

In addition, deep gMAM is also able to detect solutions going through some hyperbolic limit cycles in the separatrix and obtained by imposing a total length (17) $L = 10$. There are shown in Fig. 10. We do not explore further the action functional space, since it is very rich. It illustrates that deep gMAM is a very powerful tool for detecting various types of local minimizers.

We conclude this part, by illustrating the use of adiabatic descent 3.6.3. We would like to study the Arrhenius law for Eqs. (33) as a function of the diffusion coefficient ν . Once the solution shown in Fig. 9 is computed, we can apply the procedure described in (3.6.3) using ν . The boundary layer size decreases like $\sqrt{\nu}$, it is therefore better to use a logarithmic change, e.g. $\nu_k = \max(\nu_{\text{end}}, \nu_0(\nu_{\text{end}}/\nu_0)^{k/N})$ with N large, we use here $\nu_0 = 9 \cdot 10^{-3}$ and $\nu_{\text{end}} = 10^{-3}$. The result is shown in Fig. 11 using the same NN used in Fig. 8. It is possible to decrease even further ν and obtain more precise curves by using NNs with larger capacity.

5.4. Nongradient case: $d = 3$ Sheared Ginzburg-Landau

The last example illustrates the fact that higher-dimensional PDEs can be considered as well. In many instances, one wishes to study the action functional space in higher dimensions, say $d = O(10)$

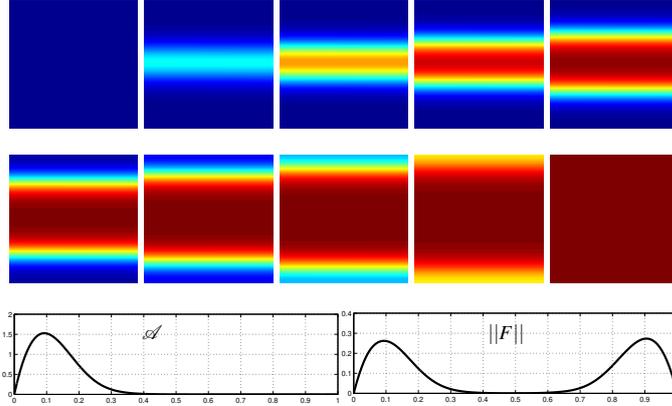


Figure 8: A local x -invariant minimizer (nucleation) for $\nu = 5 \cdot 10^{-3}$ and $c = 0.1$ for Eq.(33). Ten snapshots at equally distance are shown with color ranging from -1 (blue) to +1 (red). The NN has 7 hidden layers with 10 neurons with swish activations. The value of the action and the spatial norm $\|F\|$ are shown as a function of the normalized arclength in the panels below. The saddle corresponds approximately to the panel (2,1) in the snapshot sequence. The total length is $L \approx 2.9$ and the action $\mathcal{A}_{a \rightarrow b} \approx 0.269$ (to be compared to 0.266, see formula (32)).

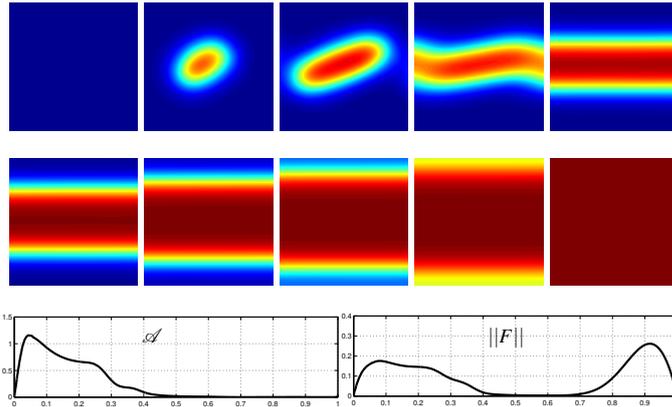
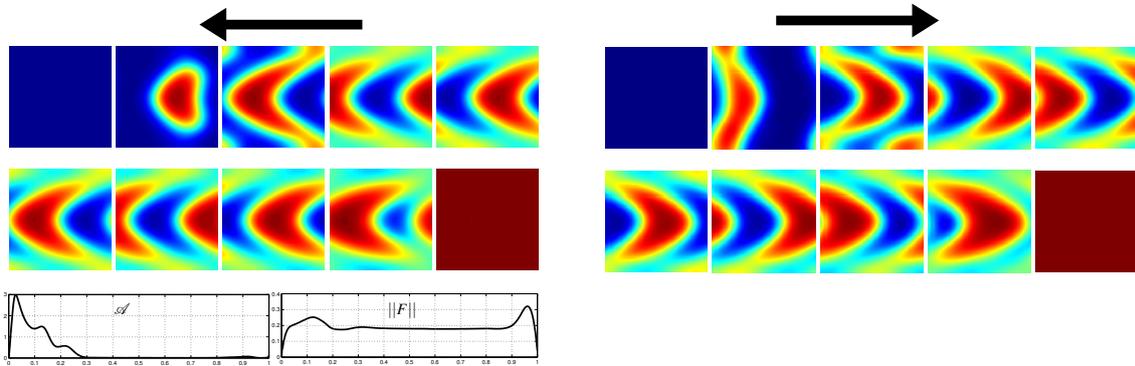


Figure 9: Same as in Fig. 8. The nucleation now breaks the x -invariance and is believed to be the instanton global minimizer (to be compared with [34]). The total length is $L \approx 3.3$ and the action $\mathcal{A}_{a \rightarrow b} \approx 0.247$.



$$\mathbf{L = 10, \mathcal{A} = 0.35}$$

Figure 10: Two different transitions going through two different hyperbolic limit cycles. Same setup than Fig. 8. They seem to be related to the limit cycle obtained in [66] Fig. 13 by a symmetry-breaking mechanism.

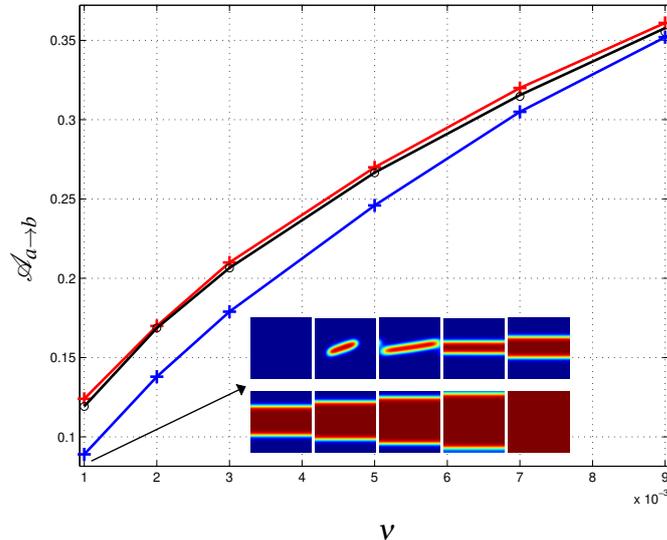


Figure 11: Dependence of the Arrhenius law with the viscosity ν for the x -invariant solutions of Fig. 8 (red curve) compared with the theoretical formula (32) (black curve). The instanton solutions which break the x invariance with smaller action are also shown (blue curve) together with a panel at very low viscosity $\nu = 10^{-3}$ showing the oumuamua typical shape.

(we recall that it is the dimension of the physical space here). The case $d = 3$ is often encountered in practice, say in hydrodynamics, but should not be considered as an upper limit of what can be achieved. In fact, one can also increase the input dimension by adding problem parameters inside the NN parametrization and compute directly a whole family of instantons. This approach is not used here but an explicit example is given in Appendix 8.2. A very rough estimate of the typical size of the minimisation problem using classical methods when $d = 3$ is at best $O(10M)$ dofs, since the fields are four-dimensional ($d + 1$). This is a straightforward consequence of the curse of dimensionality. In deep gMAM the cost is reduced drastically, the experiment shown in Fig. 12 use 1662 parameters only. The explanation for such an impressive difference, is in fact that the transition paths have a rather simple structure embedded in a very high-dimensional (phase) space. In other words, the geometrical structures live in some reduced space of finite and small dimension. Nevertheless, the ability for simple NNs to detect complicated embedding is simply outstanding. Figure 12 shows a typical codim-1 nucleation for the 3-D ACGL PDE (33) using the shear $0.1 \sin 2\pi y (\partial_x u + \partial_z u)$ instead. Note that we were not able to find codim-2 nucleations, in this case: all codim-2 candidates indeed are relaxing on codim-1 nucleations.

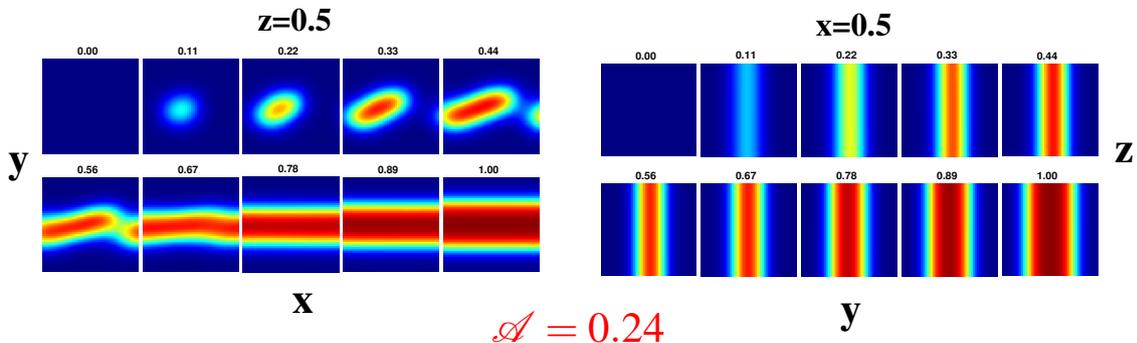


Figure 12: Codim-1 nucleation obtained using a NN with 7 hidden layers, 10 neurons/layer and tensorial batches $N_\tau \times N_x = 16 \times 32$ (see (19)) and two symmetry-breakers in x and y (see (34)).

6. Extreme events in Burgers 1-D

The last example is very different from the previous ones. The problems before are relying on FW large deviation estimates of the transition probability between two bimodal states a and b . The small large deviation parameter is the noise amplitude and one is looking at *rare events*.

In this section, we rather consider some large deviation estimates of the occurrence of extreme events in Burgers fluid turbulence. The large deviation parameter is now related to the notion of extreme events instead, here the pointwise gradient amplitude of the solution denoted α in the following. Moreover, the problem now involves a more complicated covariance operator χ so that one must use (16). Moreover, the ending state b is now replaced by a set of states characterized by an observable.

Burgers equations are very useful for studying turbulence as they share similarities with the 3-D Navier-Stokes equations, in particular they both have a direct energy cascade and exhibit *intermittency*. This phenomenon is related to the anomalous scaling in the velocity increments. In more details the mean quantities $S_n = \langle |u(x+h) - u(x-h)|^n \rangle$ called structure functions scale like h^{ξ_n} with some ξ_n smaller than the typical Kolmogorov dimension estimate $n/3$. Closely related to these quantities are the velocity gradients which is the focus here. A detailed description of this problem can be found in [27]. The stochastically-forced 1-D Burgers reads after proper rescaling of α

$$\partial_t u - F(u) = \sigma \eta, \quad F(u) = -u \partial_x u + \partial_{xx}^2 u, \quad (35)$$

with some additive noise having correlation in space and white in time

$$\langle \eta(t, x) \eta(t', x') \rangle = \delta(t - t') \chi(x - x'),$$

and with small amplitude in the limit of extreme events $\alpha_0 \gg 1$ (see Remark 1 in [58] and Appendix 8.3)

$$\sigma = \left(\frac{\chi_0}{\nu} \right)^{\frac{1}{2}} \frac{1}{\alpha_0}.$$

One has therefore a coincidence between the limit of small noise amplitude and the limit of extreme events. Traditionally, one chooses

$$\chi(x) = (1 - x^2) e^{-\frac{x^2}{2}} = -\partial_{xx}^2 (e^{-\frac{x^2}{2}}), \quad (36)$$

in order to force the large scales at length $L = 1$ (and imposes the solution to have zero mean). The small scales are thus essentially governed by the Burgers dynamics through the direct energy cascade. We are interested in the probability to observe a given velocity gradient at a particular region of space time, say $t = 0, x = 0$:

$$\Pr((\partial_x u)(0, 0) = \alpha), \quad (37)$$

and its behavior when $|\alpha| \rightarrow \infty$. The phenomenology is very different depending on the sign of α . As explained in [1, 27], for $\alpha > 0$ (right tail), the dynamics essentially follow the characteristics which damp the gradients. The noise must compensate this nonlinear effect in a significant way so that one expects an important deviation from Gaussianity. It can be shown [31] that in this case $\log \Pr \sim -c_\chi |\alpha|^3$. The case $\alpha < 0$ (left tail) is more complicated and this is the situation we are interested in. In this case, one has the appearance of shocks regularized by the viscosity ν . As $|\alpha|$ becomes large, the noise must compensate the viscosity effects preventing the blow up of the solution. A LDP can be shown to exist in this regime, ie in the viscous tail of the distribution, with a scaling behavior which can be described by the instanton of the FW action. In such a case, the scaling is known to behave as [1]

$$\log \Pr \sim -c_{L, \nu, \chi} |\alpha|^{\frac{3}{2}}. \quad (38)$$

We now describe the minimum action FW problem. Although the following is technical, it illustrates well how to adapt a complicated problem to a setup which works. Due to the presence of the covariance term, the norms involve the convolution operator with χ :

$$\|f\|_\chi = \langle F, \chi \star F \rangle, \quad \chi \star F = f.$$

The geometrical action is now (10) with the above norm. The initial state a is zero and there is no final state b , it is replaced instead by the pointwise gradient constraint:

$$(\partial_x u)(1, 0) = \alpha.$$

The main technical difficulty is to handle the convolution operator properly, ie one must solve $\chi \star v = \dot{u}$ and $\chi \star w = F(u)$ and use the auxiliary fields v, w in the geometrical action (see Eqs 11,16). It turns out that an efficient way to do it is to use the Fourier space (continuous) representation, with the unitary conventions $\hat{f}(\xi) = (2\pi)^{-\frac{1}{2}} \int_{\mathbb{R}} f(x) e^{-i\xi x} dx$, $f(x) = (2\pi)^{-\frac{1}{2}} \int_{\mathbb{R}} \hat{f}(\xi) e^{i\xi x} d\xi$. In this case, we know that (Plancherel)

$$\|f\|_{\chi} = \langle F, \chi \star F \rangle = \langle \hat{F}, \hat{\chi} \hat{F} \rangle = \|\hat{F}\|_{\hat{\chi}}.$$

where the last term, is a simple weighted L_2 norm with weight

$$\hat{\chi}(\xi) = \xi^2 e^{-\frac{\xi^2}{2}}.$$

The geometrical action is thus

$$\mathcal{A}_g = \int_0^1 (\|\hat{v}\|_{\hat{\chi}} \|\hat{w}\|_{\hat{\chi}} - \langle \hat{v}, \hat{w} \rangle_{\hat{\chi}}) d\tau,$$

with the two constraints

$$\hat{\chi} \hat{v} - \hat{u} = 0, \quad \hat{\chi} \hat{w} - \widehat{F(u)} = 0.$$

The final technical difficulty is to estimate correctly the Fourier transforms \hat{u} and $\widehat{F(u)}$ using ad-hoc rescaling. One can in principle truncate the spatial integrals by relying of the fast decay of u : $u(\tau, x) \rightarrow 0$, $x \rightarrow \infty$. We do instead a nonlinear transformation by mapping $\mathbb{R} \rightarrow (-1, 1)$ using the transformation

$$Y : x \mapsto \tanh \frac{x}{l}, \quad l > 0.$$

We give the change of variables for convenience, it is $\int_{\mathbb{R}} F(u) e^{-i\xi x} dx = \int_{-1}^{+1} (-u \partial_Y u + \nu l^{-1} ((1 - Y^2) \partial_Y^2 u - 2Y \partial_Y u) \exp(-i\xi l \tanh Y) dY$. Since we are interested in the case $|\alpha| \rightarrow \infty$, we also need to rescale the problem. Although for moderately large $|\alpha|$ it is not needed, it is better to do so to avoid possible NN saturations when the gradients become too large. We therefore use a NN ansatz following (14) except for the ending condition and with an additional scaling factor. The ansatz is, using the same notation u in the Y coordinate,

$$u(\tau, Y) = l^{-1} |\alpha|^\gamma (1 - Y^2) \tau \mathcal{U}(\tau, Y), \quad \tau \in [0, 1], \quad Y \in]-1, 1[.$$

where \mathcal{U} is now a feedforward NN. The term $1 - Y^2$ not only enables one to automatically satisfy the boundary conditions at $x \rightarrow \pm\infty$ (ie $Y \rightarrow \pm 1$) but also avoid asymptotical pathological behavior in the NN representation due to the term \hat{u} . The scaling $l^{-1} |\alpha|^\gamma$ is here to compensate the very large amplitude of the instantons. It can be removed from the equations by rescaling the auxiliary fields \hat{v}, \hat{w} as well:

$$\hat{v} = l^{-1} |\alpha|^\gamma \hat{V}, \quad \hat{w} = l^{-1} |\alpha|^\gamma \hat{W}.$$

Finally, the penalisation constraints for the auxiliary fields \hat{V}, \hat{W} are after simplifications

$$\mathcal{C}_v = \int_0^1 \int_0^{\xi_c} \left| \hat{\chi} \hat{V}(\xi) - \int_{-1}^{+1} \partial_\tau (\tau \mathcal{U}(\tau, Y)) \exp((-i\xi l \tanh Y) dY \right|^2 d\xi d\tau$$

and

$$\mathcal{C}_w = \int_0^1 \int_0^{\xi_c} \left| \hat{\chi} \hat{W}(\xi) - \int_{-1}^{+1} \mathcal{F}(\mathcal{U}) \exp((-i\xi l \tanh Y) dY \right|^2 d\xi d\tau,$$

with

$$\mathcal{F}(\mathcal{U}) = -|\alpha|^\gamma \mathcal{U}_s \partial_Y \mathcal{U}_s + \nu l^{-1} ((1 - Y^2) \partial_{YY}^2 \mathcal{U}_s - 2Y \partial_Y \mathcal{U}_s), \quad \mathcal{U}_s = (1 - Y^2) \tau \mathcal{U}(\tau, Y).$$

The integrals in Fourier space are half-integrals using a truncated frequency ξ_c since u is real. The geometrical action is

$$\mathcal{A}_g = l^{-2} |\alpha|^{2\gamma} \int_0^1 (\|\hat{V}\|_{\hat{\chi}} \|\hat{W}\|_{\hat{\chi}} - \langle \hat{V}, \hat{W} \rangle_{\hat{\chi}}) d\tau, \quad (39)$$

where $\langle \hat{V}, \hat{W} \rangle_{\hat{\chi}} = \int_0^{\xi_c} \Re\{\hat{V} \overline{\hat{W}}\} \hat{\chi} d\xi$. The choice of γ is in principle arbitrary, but according to (38), it is convenient to choose $\gamma = \frac{3}{4}$ so that the integral in (39) must asymptotically behave independently of α . We thus call this term the *compensated geometrical action*. The gradient constraint in the new Y coordinate after simplifications reads

$$\partial_Y \mathcal{U}(1, 0) = l\alpha |\alpha|^{-\gamma}.$$

This constraint is handled through a straightforward quadratic penalisation with a large penalty coefficient. Due to the rescaling used, the right-hand side term scales like $|\alpha|^{1/4}$ so that large gradient constraints can be used. We use a value of $l = 25$ in applications, a too small value would yield a bad representation of the u asymptotical tails $|x| \rightarrow \infty$ by strongly zooming into the interval $|x| \ll 1$. This setting is likely not optimal (e.g. a better tuning would use $l = l(\alpha)$) but it works well in practice for α in the range $O(10^2)$ – $O(10^3)$. The auxiliary fields V, W are replaced by two NNs of dimension output 2, one for the real part and one for the imaginary part. One can also use four NNs instead. The batches are of two kinds, the ones in physical space $(t, x) \in (0, 1) \times (-1, 1)$ and the others in the Fourier space $(t, \xi) \in (0, 1) \times (0, \xi_c)$ using uniform distributions.

We show first the instanton path obtained for a small value of α together with its spatial profile for $\tau = 1$. Its space-time structure compares very well with known results, e.g. Fig. 2 in [26]. It has zero spatial mean as expected and shares some similarities with the derivative of the kink solution in Burgers although with heavier tails and a change of sign visible at $|x| \approx \pm 5$. The increase of the distance between the two extrema for $\tau < 1$ is also typical of the space-time instantons found in the literature. We now show the estimate of the left tail of the distribution by changing the amplitude

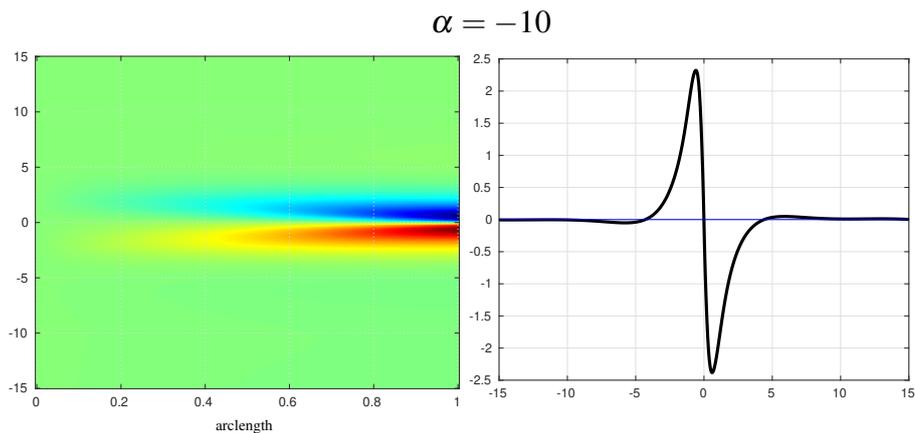


Figure 13: Instanton solution obtained by deep gMAM using a NN with 14 hidden layers, 10 neurons/layer and swish activations with batches $N_\tau = 15$, $N_Y = 500$, $N_\xi = 500$, training rate 10^{-4} , $l = 10$ and $\xi_c = 30$.

of the gradient α to larger values using the setup described above. The method strikingly recovers the correct exponent $\frac{3}{2}$. We do not use the adiabatic approach here but rather a poor-man version using initial conditions from previously computed instantons and freeze α in the estimates. The geometrical action estimates are the mean over the last 50 000 iterations. The setup proposed here is likely not optimal due to the strongly ill-conditioned problem with about 5% relative fluctuations

from the mean action. The result shows that the power-law estimate has in fact a smaller exponent in the range $|\alpha| = O(10^2)$ in accordance with the 1.15 exponent found in [28, 25]. This can be seen in the monotonic decrease of the compensated action in the insert of Fig. 14. However, in our case, we are able to explore an additional decade with very clear convergence to $\frac{3}{2}$ consistent with the theory [1].

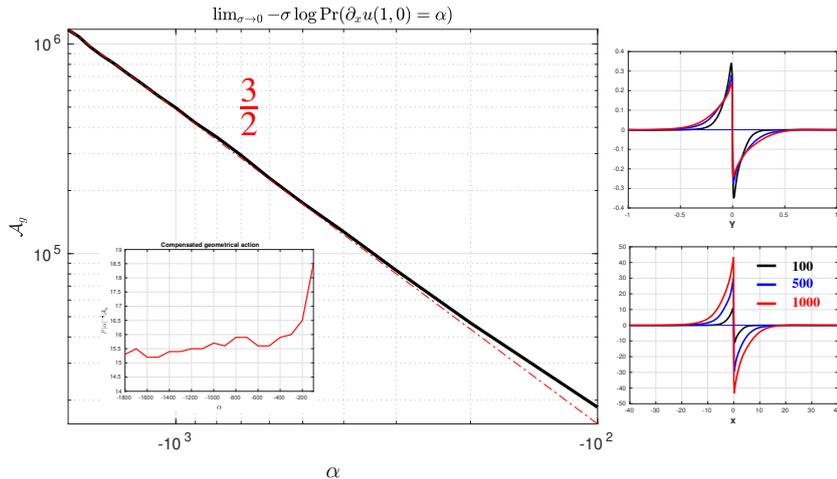


Figure 14: Power law for the viscous left tail in Burgers shocks. The insert is what is actually computed, ie the compensated action as a function of α and the loglog figure is (39) (the exact $\frac{3}{2}$ power law is represented by a dash-dot red line). The right panels show the solutions $u(1, Y)$ for different values of α . Upper-right panel are the renormalized curves in Y coordinate and lower-right panel shows the true ones in x coordinate. The penalisation parameters are $(\gamma_g, \gamma_{\text{arc}}, \gamma_\alpha, \gamma_v, \gamma_w) = (5, 1, 200, 1000, 1000)$ and $l = 25$, other parameters like in Fig. 13.

An important remark is that these approaches can be extended to higher-dimensional problems: the use of the continuous Fourier representation, including the change of variables, together with the Monte-Carlo sampling strategy make this possible. This would be impossible if one attempts to project the fields on a discrete Fourier basis, or mimicking some FFT representation or any other Galerkin projection, since one would face very heavy optimisation problems (curse of dimensionality at $d \geq 3$).

7. Conclusion

The deep gMAM method is a natural formulation of the classical gMAM method in the context of machine learning. It is a penalisation approach which offers great flexibility at the cost of less precision than classical methods. However, its outstanding property is its expressivity and its scaling with the dimension, the method already starts to shine for spatial dimension $d = 2$. The reason lies on the well-known property for deep networks to overcome the curse of dimensionality. We have shown various examples, from classical bimodal switches in SDEs and Ginzburg-Landau PDEs to genuine nongradient, nontransverse situations. One can easily obtain general Arrhenius laws into a wide variety of contexts. Furthermore, the deep gMAM approach is able to compute extreme events as well: it is illustrated in a challenging example of Burgers turbulence. We are able to recover the correct power law exponent $\frac{3}{2}$ for the left probability tail at very small numerical cost.

We discuss below first an important alternative to the FW minimum action principle and second some general perspectives.

7.1. Other approaches: rare event algorithms

An important alternative class of methods to estimate rare events are the so-called importance sampling [36] and importance splitting methods. The splitting approach is particularly illustrative of methods which do not exploit some minimum action principle at all. It relies on smart sampling

strategies to provide events in the tail of the distribution. One can already identify an important advantage compared to minimum action methods: they can be used in a wider variety of situations since no LDP is required. In particular, these approaches can consider finite-amplitude noise rather than the zero noise limit. This would be indeed an ideal situation. Unfortunately, there are some fundamental issues which cannot be ignored. The case of importance splitting methods is particularly interesting to this respect. They originally come from an idea of J.Von Neumann in Los Alamos and were called *Von Neumann splitting* [37, 57]. A revival of this approach appeared much later with the work of mathematicians in probability theory [49, 10] and are now called adaptive multilevel splitting (AMS) or rare events or even genetic algorithms. The idea is to draw an ensemble of short simulations and select only the "most interesting" events as new initial conditions: this is often referred to as a selection-mutation step owing to some Darwinian selection principle in the model. The main issue is to define what "most interesting" means. These methods therefore rely on a so-called *reaction coordinate* to decide when some event is better than another. The mutation step is also strongly tightened to the amplitude of the noise in the model as one must bring diversity to obtain relevant statistical ensembles. A too small noise amplitude (or a nonchaotic deterministic system) would yield fast species extinction. The method in these cases are failing to obtain relevant ensembles associated to the tail distribution unless one uses prohibitively large-sized ensembles (the number of clones). These methods have thus many issues in the small noise-amplitude limit. The main reason takes its root in the choice of a good reaction coordinate. The better the reaction coordinate the less variance in the estimates with nearly Gaussian behavior [32, 61] (see also [6] for the general case). The less noise, the more sensitive is the choice of the reaction coordinate [56, 6]. It turns out that the optimal reaction coordinate is known, it is called the committor function and is a solution of the backward Focker-Planck (Kolmogorov) equation (BFP) [17]. Since for a system with N d.o.f.s the corresponding BFP equations is a N -dimensional PDEs, it is therefore a very challenging if not an impossible task unless $N = O(10-100)$ at most (for models in fluid dynamics, $N = O(10^6-7)$ and more). In fact, being able to compute a pointwise estimate of the committor is just what these rare events algorithms are trying to do. There are lots of works devoted to this question. In the machine-learning context, one can indeed try to compute the committor function in some reduced space of collective variables [3], (see also [38]). A more involved approach denoted loop-AMS has been proposed recently which amounts to identify a better reaction coordinate in a dynamical way [44]. Although the results are encouraging, they also face the issue of the small noise-amplitude limit: the committor function in phase space becomes nearly singular with Heaviside-type of behavior. The corresponding exponentially flat regions are indeed reflecting the large deviation scaling of the probability to estimate. One can also consider spectral theory [62, 70] to investigate the BFP spectral gap and leading spectrum but again the singular small-amplitude limit [21] cannot be handled easily. Nevertheless, AMS approaches using non-optimal reaction coordinates with large enough noise-level show spectacular results which cannot be ignored: recent examples include heat waves [51], 1-D Ginzburg-Landau [55] and quasigeostrophic 2-D turbulence [5, 63], transition to turbulence [54].

For these reasons, if the noise level in the model is too small and an effective action is known (here the FW action), it is much more reliable to consider minimum action methods in general. Both approaches are complementary however: they are optimal in two almost separated regimes.

7.2. Perspectives

The next immediate perspective is to consider situations where there is no access to an explicit Lagrangian, see Eqs (6). In this case, one must solve the minimax problem using adversarial networks by some stochastic descent-ascent algorithms. Although the general ideas of the deep gMAM method are the same, it requires a more specific and careful treatment with ad-hoc examples.

In general, one would like to extend the domain of validity of the proposed approach. First, an important theoretical question is to extend the method beyond the FW regime. Having access to an explicit macroscopic action is not always possible, even worst the existence of a LDP is not guaranteed at all. An illuminating example is the 2-D quasigeostrophic barotropic equation stochastically forced at the small scales. It is used as a model for mid-latitude atmospheric and oceanic jets through an

inverse turbulent cascade (zonostrophy, see [63]). In such a case, the FW action related to bimodal switches is not relevant anymore [41]. One needs an effective action which can be identified by advanced kinetic analysis [4, 65]. In this scenario, the use of rare events algorithms such as AMS (see subsection 7.1) seems unavoidable and gives good results as soon as the noise amplitude is large enough [5].

Another situation is the important case of chaotic deterministic systems. Very often, there is no LDP and the mean waiting time is in fact replaced by some power law [29]. We discuss in the Appendix 8.1 a striking example of a simple chaotic system for reversal dynamo which has both. The use of deep gMAM reveals that the FW small action minimizers are strongly connected if not equal to the exact preferred transition path. We thus believe that deep gMAM can be extended in much more general situations, much beyond the LDP hypothesis. It also suggests that a marriage between splitting methods and gMAM approaches is indeed possible.

Acknowledgments

The author thanks T.Grafke, E.Vanden-Eijnden, R.Zakine and M.D.Chekroun for stimulating discussions and support.

8. Appendix

8.1. Bimodality: minimizing the FW zero action paths

In many cases, the addition of noise in a deterministic system – often for modelisation reasons – is necessary to observe transitions between two attractors a and b . One calls them noise-induced since the purely deterministic system is not able by itself to trigger transitions between the attracting states. In the context of chaotic deterministic systems, the situation can be very different. Chaos can often triggers transitions between a and b through intermittency alone: the transitions are no more noise-induced. As soon as, a timescale separation is present, these transitions are readily interpreted as rare events. However, their probabilities of occurrence do not satisfy some LDP. The example is a simple model of dynamo reversal having three d.o.fs [23, 22]:

$$\dot{x} = \mu x - yz, \dot{y} = -\nu y + xz, \dot{z} = -z + xy + \Gamma. \quad (40)$$

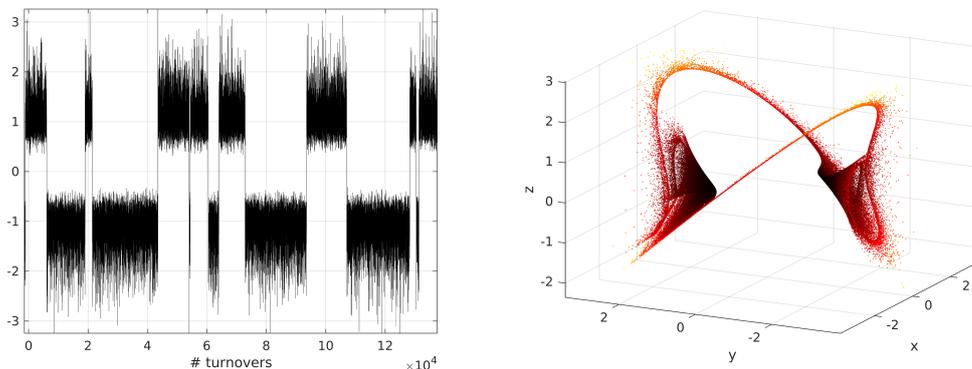


Figure 15: Bimodality in system (40): timeseries (left panel) versus attractor in phase space (right panel) using $\mu = 0.1192805$, $\nu = 0.1$, $\Gamma = 0.9$.

Such phenomenon is well-known in chaos theory and is called *crisis-induced intermittency*, see the seminal work of Grebogi, Ott and York [30]. It occurs when two disconnected basins of attraction start to overlap at a given critical value of the system parameters. In chaotic systems having discrete group symmetries such scenario is the rule rather than the exception [13]. It appears that the mean

waiting time τ , ie the mean time interval between two transitions is scaling like a power law rather than exponentially:

$$\tau \sim (\mu_c - \mu)^{-\gamma}, \quad \mu < \mu_c$$

with $\gamma \sim \frac{1}{2}$ in the dynamo model (40) [22]. In fact, the critical exponent depends on the type of crisis tangency [29]. With this in mind, one can ask about the status of the FW action (see (4) with $\chi = Id$). Since no noise is actually needed to observe a transition, the transition paths are purely deterministic. The action along these paths is therefore identically zero.

Before the crisis bifurcation ($\mu > \mu_c$), it is clear however that the transitions must be noise-induced and FW theorem does apply. Let us look in more details. The system has one critical point $O = (0, 0, \Gamma)$ and two stable oscillatory points at $S_{\pm} \equiv (\pm(\nu + \Gamma(\nu/\mu)^{\frac{1}{2}})^{\frac{1}{2}}, -\pm(\nu + \Gamma(\nu/\mu)^{\frac{1}{2}})^{\frac{1}{2}}, -(\mu\nu)^{\frac{1}{2}})$. An important particularity is that O has Morse index 2 with an oscillatory unstable manifold. In this case, the instanton path does not go through O but crosses the separatrix through a nontrivial point. This contrasts with the Lorenz63 system [43] where transition paths do cross the separatrix through the saddle [71]. It shows that for nongradient systems, critical points in the separatrix are not guaranteed to play a role. The nature of these points in the separatrix is the subject of many theoretical researches, they are often called *noisy precursor of crisis* [60, 40, 14]. The instanton path is shown in Fig. 16 using deep gMAM. It agrees with direct ensemble simulations of (40) with a small additive Gaussian noise (not shown).

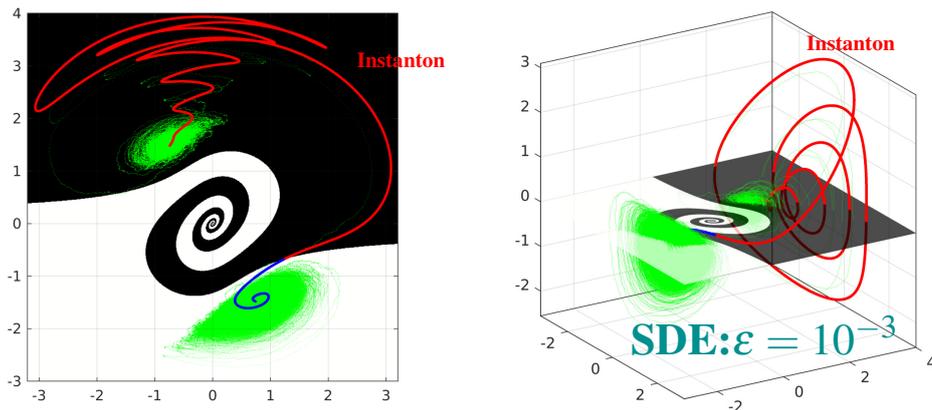


Figure 16: Instanton path computed by deep gMAM in the FW regime ($\mu > \mu_c$). The black and white spiraling regions show the oscillatory separatrix around O and the green curves are stochastic direct simulations with small noise ϵ .

We now consider the regime where FW do not apply, ie in the crisis regime for $\mu < \mu_c$. We naively apply deep gMAM on various initial conditions on the attractor a . The result is striking and is shown in Fig. 17. First the instanton paths not surprisingly depend on the initial condition, second and more importantly, many of them do agree with the deterministic transition by concentrating around it (red-blue curves in Fig. 17). The green curve is an outlier with a transition path which differs from the exact solution. However, this outlier is indeed correct up to a scaling factor: there is some $g > 0$ such that $g\Gamma$ is in the neighborhood of the deterministic transition.

In principle it would be interesting to relax the BV constraints using instead a whole set of conditions in the neighborhood of the attracting points S_{\pm} . For instance, by penalizing the characteristic functions $\chi_{B(S_{\pm}, \rho)}$ where $\chi_{\mathcal{D}} = 1$ if $\mathbf{x} \in \mathcal{D}$ and zero otherwise, like in section 6. In fact, AMS also uses a similar idea by sampling the equilibrium measure on a codim-1 surface surrounding the attractors [11]. It appears however that the correct approach can be more subtil, it is discussed in [40].

One of interpretation for the validity of FW transition paths in the near zero-action regime is mostly related to the use of the geometrical action itself and should not come as a surprise. In-

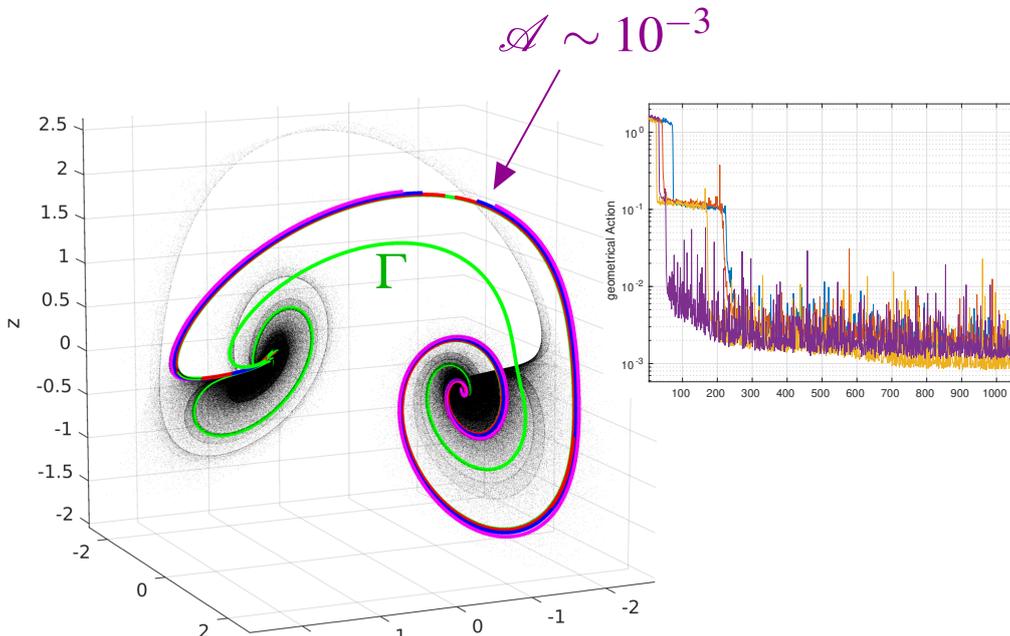


Figure 17: Near zero-action instantons paths with different initial conditions a using deep gMAM: trajectories in phase space (left panel) and action during the stochastic gradient descent. The red-purple curves concentrate on the deterministic transition whereas the outlier green curve is correct up to a scaling factor.

stead of integrating the system for a very long time, one is able to detect trajectories with time reparametrization by minimisation of the geometrical action. It therefore bypasses the long residence time (and accumulation points) for the state to visit the attractors. We believe that such phenomenon should be analyzed in more details on more complicated systems as it is potentially a very strong alternative to rare events algorithms even when no LDP exists. In such scenario, one has no more access to probability estimates since the LDP interpretation has disappeared but nevertheless the correct transitions are captured. Probability estimates must therefore be obtained by other approaches like AMS (see subsection 7.1).

8.2. A simple Julia code with comments

We illustrate here the great flexibility, simplicity and efficiency of neural networks in general when coupled with user-friendly modern languages (Julia, PyTorch). We provide here a snippet Julia code for computing the quasi-potential in subsection 4.3: it is short, simple and can be easily modified and extended. We do not write the monitoring and display subroutines as they are user-dependent. It is important for instance to test different hyperparameters (NN size, training rate, batch size, activation functions, descent algorithm, etc...). One can also set `din = 1` for computing only one instanton. In this case, it is interesting to add the arclength condition (15). One must simply add $\text{mean}(A2) - (\text{mean}(A))^2$ to the returned quantity in the function loss, with some (small) penalty weight. Note that when $N_\tau > 1$, the arclength penalty is no more additive. One can also explore different ODEs, e.g. Lorenz63, or even setting `din = 5` for computing $(a, b) \mapsto V(a, b)$. A more difficult exercise is then to switch to the PDE case.

We give here another simple example of a 2 d.o.fs family of SDEs depending on some parameter σ such that it is gradient when $\sigma = 1$ and Hamiltonian when $\sigma = 0$. Let $u = (u_1, u_2)$ and

$$F(u, \sigma) = \sigma G + (1 - \sigma)G^\perp, G = (-u_1^3 + u_1 - 0.3u_1^2u_2 - 0.6u_1^2, -u_2 - 0.1u_1^3). \quad (41)$$

By definition this system has the transverse decomposition (24) and classical Arrhenius laws can be obtained explicitly [66]. It has two stable points and a saddle point at the origin. When $\sigma < 1$, the geometrical support of the transition $a \rightarrow b$ differs from $b \rightarrow a$: it is called the figure-eight scenario. The

transitions $a \rightarrow b$ must all cross the saddle $(0,0)$ and then relax on deterministic zero-action paths spiraling towards b (and outwards from a). The cost functional is $\mathcal{C}[u] = \int_0^1 d\sigma \int_0^1 d\tau A_g(u(\tau, \sigma))$. Figure 18 gives the typical result obtained for the two transitions $a \rightarrow b$ and $b \rightarrow a$. The code is essentially the same code than below with very few modifications.

```

~~~~~ Compute the Maier-Stein quasi-potential b->V(a,b) for a=(-1,0)
using Flux,Distributions
#-----
N = 2 # Maier-Stein has only N=2 dofs
din = N+1 # first dim is tau dim 2,...,N are for b
ee = vcat(1e-6,zeros(N)) # fd tau increment [eps,0,...,0]
a = [-1;0] # ODE fixed pts a (N=2 only)
cap = 15 # NN capacity neurons/layer
act = swish # activation function
tr = 1e-3 # training rate
NT = 1000 # batch size
Nmax = 50000 # max numb of descent iterations
opt = Flux.ADAM(tr) # we use ADAM descent with training rate tr
#-----
# We define a feedforward NN with input dimension N+1 and output dimension N for (u1,...,uN)
# it has 2 hidden layers having cap neurons/layer using swish activations
# This NN is likely too small -> one must increase the # of hidden layers

NN = Chain(Dense(din,cap,act),Dense(cap,cap,act),Dense(cap,cap,act),Dense(cap,N))
ps = Flux.params(NN) # define the NN parameters

#-----
# the BV ansatz: U(0) = a, U(1) = b, see Eq.(13)
function U(r)
    tau = r[1:1,:]
    b = r[2:din,:]
    mtau = 1.0 ./ tau
    return a*mtau + b.*tau + tau.*mtau.*NN(r)
end
#-----
# Maier-Stein system with N = 2 d.o.fs
function RHS(u)
    u1 = u[1:1,:]
    u2 = u[2:2,:]
    F1 = u1 - u1.^3 -10*u1.*u2.^2
    F2 = -(1.0 ./ u1.^2).*u2
    return [F1;F2]
end
#-----
function loss(r) # r is a random set of pts of size (din,NT)
    dotU = (U(r.+ee)-U(r.-ee))/(2*ee[1]) # or using Zygote AD
    FU = RHS(U(r))
    A2 = sum(abs2,dotU,dims=1) # L2 norm squared ||dotU||^2(tau,b)
    A = sqrt.(A2) # L2 norm ||dotU||(tau,b)
    B = sqrt.(sum(abs2,FU,dims=1)) # L2 norm ||F(U)||(tau,b)
    AB = sum(dotU.*FU,dims=1) # Hilbert product <dotU,F(U)>(tau,b)
    ACTION = mean(A.*B) - mean(AB) # sum over b and tau: see Eq.(25)
    return ACTION
end
#-----
# The stochastic gradient descent using ADAM to find all the instantons
for k in 1:Nmax
    batch = rand(din,NT) # generate a total of NT random pts for (tau,b) in (0,1)^3
    gs = gradient(ps) do # Zygote AD of the cost functional
        loss(batch) # Zygote AD of the cost functional
    end # Zygote AD of the cost functional
    Flux.update!(opt,ps,gs) # NN parameters update using ADAM
end
#-----

```

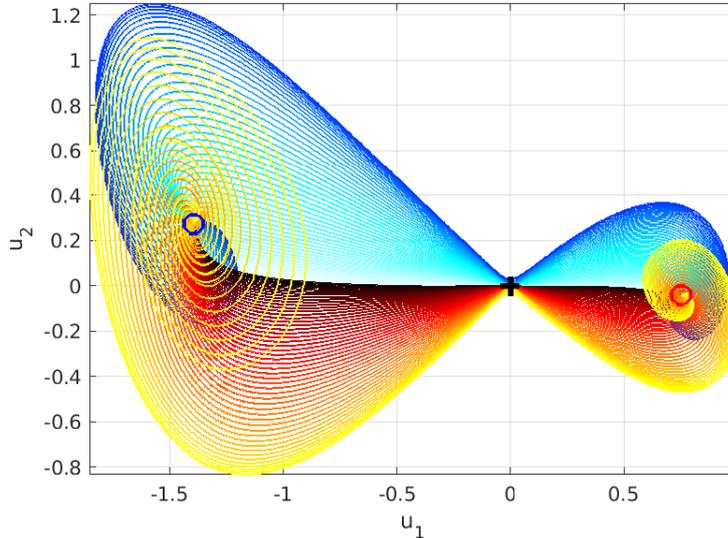


Figure 18: Family of instantons for $a \rightarrow b$ (red colors) and $b \rightarrow a$ (blue colors). They all cross the saddle (black cross) and are spiraling outwards or towards a (left) and b (right) in the Hamiltonian regime when $\sigma \rightarrow 0$. The dark-colored instantons correspond to the gradient regime ($\sigma \sim 1$). Note that in this example the transitions $b \rightarrow a$ are much more likely than $a \rightarrow b$ and can be quantified precisely by monitoring the action values.

8.3. Burgers rescaling

We set $u = Uu'$, $x = Lx'$, $t = Tt'$. The noise term therefore scales as $\eta = (\chi_0^{\frac{1}{2}} T^{-\frac{1}{2}}) \eta'$. The extreme event condition reads $(U/L)(\partial_{x'} u')(1, 0) = \alpha_0 \alpha'$ and we take $\alpha_0 = U/L$. Choosing $U = L/T$, $T = L^2/\nu$ gives $L = (\nu/\alpha_0)^{\frac{1}{2}}$ and

$$\partial_t u + u \partial_x u - \partial_{xx}^2 u = \sigma \eta \text{ with } \sigma = \left(\frac{\chi_0}{\nu} \right)^{\frac{1}{2}} \frac{1}{\alpha_0} \xrightarrow{\alpha_0 \rightarrow \infty} 0$$

One can therefore use the FW formalism in this context [28, 58]. When ν is small, one expects that the asymptotical regime is attained for much larger α_0 . This is self-consistent with the idea that atypical extreme events must be much more extreme when the Reynolds is large than the ones in the small Reynolds number regime.

References

- [1] E. Balkovsky, G. Falkovich, I. Kolokolov, and V. Lebedev. Intermittency of burgers' turbulence. *Phys. Rev. Lett.*, 78:1452–5, 1997.
- [2] Ch. Beck, S. Becker, P. Cheridito, A. Jentzen, and A. Neufeld. Deep splitting method for parabolic PDEs. *SIAM J. Sci. Comp.*, 43(5):A3135–A3154, 2021.
- [3] Z. Belkacemi, G. Paraskevi, T. Lelièvre, and G. Stoltz. Chasing collective variables using autoencoders and biased trajectories. *arXiv:2104.11061v2*, 2022.
- [4] F. Bouchet, C. Nardini, and T. Tangarife. Kinetic theory of jet dynamics in the stochastic barotropic and 2D Navier-Stokes equations. *J. Stat. Phys.*, 153(4):572–625, 2013.
- [5] F. Bouchet, J. Rolland, and E. Simonnet. Rare event algorithm links transitions in turbulent flows with activated nucleations. *Phys. Rev. Lett.*, 122(7):074502, 2019.
- [6] C.E. Bréhier, M. Gazeau, L. Goudenège, T. Lelièvre, and M. Rousset. Unbiasedness of some generalized adaptive multilevel splitting algorithms. *Annals of Appl. Prob.*, 26(6):3559–3601, 2016.
- [7] W. Cai and Zhi-Qin J. Xu. Multi-scale deep neural networks for solving high dimensional PDEs. *arXiv:1910.11710v1*, 2019.
- [8] M.K. Cameron. Finding the quasipotential for nongradient SDEs. *Physica D: Nonlinear Phenomena*, 241(18):1532–1550, 2012.
- [9] A. Cavagna, A.J. Bray, and Rui D.M. Travasso. Ohta-Jasnow-Kawasaki approximation for nonconserved coarsening under shear. *Phys. Rev. E*, 62:4702–4719, Oct 2000.

- [10] F. Cérou and A. Guyader. Adaptive Multilevel Splitting for Rare Events Analysis. *Stoch. Anal. and Appl.*, 25:417–443, 2007.
- [11] F. Cérou, A. Guyader, T. Lelièvre, and D. Pommier. A multiple replica approach to simulate reactive trajectories. *J.Chem.Phys.*, 134(054108), 2011.
- [12] J. Chen, R. Du, P. Li, and L. Lyu. Quasi-monte carlo sampling for solving partial differential equations by deep neural networks. *Numer. Math. Theor. Meth. Appl.*, 2020.
- [13] P. Chossat and M. Golubitsky. Symmetry-increasing bifurcation of chaotic attractors. *Physica D*, 88:423–436, 1988.
- [14] J. Demaeyer and P. Gaspard. A trace formula for activated escape in noisy maps. *J. Stat. Mech.: Theor. and Exp.*, 2013(10):P10026, 2013.
- [15] S. Dong and N. Ni. A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks. *J. Comp. Phys.*, 435:110242, 2021.
- [16] W. E, W. Ren, and E. Vanden-Eijnden. Minimum action method for the study of rare events. *Comm. Pure Appl. Math.*, 57:1–20, 2004.
- [17] W. E and E. Vanden-Eijnden. Towards a theory of transition paths. *J. Stat. Phys.*, 123:503–523, 2006.
- [18] W.G. Faris and G. Jona-Lasinio. Large fluctuations for a nonlinear heat equation with noise. *J.Phys. A: Math. Gen.*, 15, 1982.
- [19] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems, volume 260 of Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, New York, 1984.
- [20] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems, 2nd ed., volume 260 of Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, New York, 1998.
- [21] P. Gaspard. *Chaos, Scattering and Statistical mechanics*. Number 9. Cambridge University Press, 2005.
- [22] C. Gissinger. A new deterministic model for chaotic reversals. *Eur. Phys. J. B*, 85(137), 2012.
- [23] C. Gissinger, E. Dormy, and S. Fauve. Morphology of field reversals in turbulent dynamos. *Euro. Phys. Lett.*, 90:49001, 2010.
- [24] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *Aistats*, 9:249–256, 2010.
- [25] T. Gotoh. Probability density dunctions in steady-state Burgers turbulence. *Phys. Fluids*, 11:2143–2148, 1999.
- [26] T. Grafke, R. Grauer, and T. Schäfer. Instanton filtering for the stochastic Burgers equation. *J. Phys. A: Math. Theor.*, 46:062002, 2013.
- [27] T. Grafke, R. Grauer, and T. Schäfer. The instanton method and its numerical implementation in fluid mechanics. *J. Phys. A: Math. Theor.*, 48:333001, 2015.
- [28] T. Grafke, R. Grauer, T. Schäfer, and E. Vanden-Eijnden. Relevance of instantons in Burgers turbulence. *EPL*, 109(3):34003, 2015.
- [29] C. Grebogi, E. Ott, F. Romeira, and J.A. York. *Phys. Rev. A*, 36:5365–5380, 1987.
- [30] C. Grebogi, E. Ott, and J.A. Yorke. Crises, sudden changes in chaotic attractors, and transient chaos. *Physica D*, 7:181–200, 1983.
- [31] V. Gurarie and A. Migdal. Instantons in the burgers equation. *Phys. Rev. E*, 54:4908–14, 1996.
- [32] A. Guyader, N. Hengartner, and E. Matzner-Løber. Simulation of extreme quantiles and extreme probabilities. *Appl. Math. Optim.*, 64:171–196, 2011.
- [33] M. Heymann and E. Vanden-Eijnden. The geometric minimum action method: A least action principle on the space of curves. *Comm. Pure Appl. Maths*, 61(8):1052–1117, 2008.
- [34] M. Heymann and E. Vanden-Eijnden. Pathways of maximum likelihood for rare events in nonequilibrium systems: Application to nucleation in the presence of shear. *Phys. Rev. Lett.*, 100:140601, Apr 2008.
- [35] J. Huang, H. Wang, and T. Zhou. An augmented lagrangian deep learning method for variational problems with essential boundary conditions. *Commun. Comput. Phys.*, 31:966–986, 2022.
- [36] Y. Jiawei, H. Touchette, and G.M. Rotskoff. Learning nonequilibrium control forces to characterize dynamical phase transitions. *Phys. Rev. E*, 105:024115, Feb 2022.
- [37] H. Kahn and T. Harris. Estimation of particle transmission by random sampling. *Natl. Bur. Stand. Appl. Math. Ser.*, 12:27–30, 1951.
- [38] Y. Khoo, J. Lu, and L. Ying. Solving for high-dimensional committor functions using artificial neural networks. *Research in the Mathematical Sciences*, 6(1):1–13, 2019.
- [39] D. P. Kingma and J. Ba. ADAM: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [40] S. Kraut and C. Grebogi. Escaping from nonhyperbolic chaotic attractors. *Phys. Rev. Lett.*, 92(23):234101, 2004.
- [41] J. Laurie and F. Bouchet. Computation of rare transitions in the barotropic quasi-geostrophic equations. *New J. of Phys.*, 17:015009, 2015.
- [42] Z. Liu, W. Cai, and Zhi-Qin J. Xu. Multi-scale deep neural network (mscalednn) for solving poisson-boltzmann equation in complex domains. *arXiv:2007.11207v3*, 2020.
- [43] E.N. Lorenz. Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20(2):130–141, 1963.
- [44] D. Lucente, J. Rolland, C. Herbert, and F. Bouchet. Coupling rare event algorithms with data-based learned committor functions using the analogue markov chain. *arXiv:2110.05050v3*, 2022.
- [45] R.S. Maier and D.L. Stein. A scaling theory of bifurcations in the symmetric weak-noise escape problem. *J. Stat. Phys.*, 83(3–4):291357, 1996.
- [46] B. Moseley, A. Markham, and T. Nissen-Meyer. Finite basis physics- informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations. *arXiv:2107.07871*, 2021.
- [47] K. Müller. Reaction paths on multidimensional energy hypersurfaces. *Angew. Chem.*, 92, 1980.
- [48] L. Onsager and S. Machlup. Fluctuations and irreversible processes: Ii. systems with kinetic energy. *Phys. Rev.*,

- 91:1512–5, 1953.
- [49] P. Del Moral. *Feynman-Kac formulae, Genealogical and interacting particle systems with applications*. Probability and Applications, Springer-Verlag, New York, 2004.
- [50] G. Poppe and T. Schäfer. Computation of minimum action paths of the stochastic nonlinear Schrödinger equation with dissipation. *J. Phys. A: Math. Theor.*, 51 335102:1–17, 2018.
- [51] F. Ragone, J. Wouters, and F. Bouchet. Computation of extreme heat waves in climate models using a large deviation algorithm. *PNAS*, 115(1):24–29, 2018.
- [52] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comp. Phys.*, 378:686–707, 2019.
- [53] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Swish: a self-gated activation function. *arXiv: Neural and Evolutionary Computing*, 2017.
- [54] J. Rolland. Collapse of transitional wall turbulence captured using a rare events algorithm. *J. Fluid Mech.*, 931 A22:503–523, 2022.
- [55] J. Rolland, F. Bouchet, and E. Simonnet. Computing transition rates for the 1-D stochastic Ginzburg–Landau–Allen–Cahn equation for finite-amplitude noise with a rare event algorithm. *J. Stat. Phys.*, 162:277–311, 2016.
- [56] J. Rolland and E. Simonnet. Statistical behavior of adaptive multilevel splitting algorithms in simple models. *J. Comp. Phys.*, 283:541–558, 2015.
- [57] M. Rosenbluth and A. Rosenbluth. Monte Carlo calculation of the average extension of molecular chains. *J. Chem. Phys.*, 23(2):356–359, 1955.
- [58] T. Schorlepp, T. Grafke, and R. Grauer. Gel’fand–Yaglom type equations for calculating fluctuations around instantons in stochastic systems. *J. Phys. A: Math. Theor.*, 54:235003, 2021.
- [59] T. Schorlepp, T. Grafke, S. May, and R. Grauer. Spontaneous symmetry breaking for extreme vorticity and strain in the 3d navier-stokes equations. *Phil Trans Roy Soc A*, 380(2226), 2022.
- [60] A.N. Silchenko, S. Beri, D.G. Luchinsky, and P.V. McClintock. Fluctuational transitions through a fractal basin boundary. *Phys. Rev. Lett.*, 91(17):174104, 2003.
- [61] E. Simonnet. Combinatorial analysis of the adaptive last particle method. *Stats and Comp.*, 26:211–230, 2016.
- [62] E. Simonnet and M.D. Chekroun. Deep spectral computations in linear and nonlinear diffusion problems. *arXiv:2207.03166v1*, 2022.
- [63] E. Simonnet, J. Rolland, and F. Bouchet. Multistability and rare spontaneous transitions in barotropic beta-plane turbulence. *J. Atmos. Sci.*, 78(6):1889–1911, 2021.
- [64] J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *J. Comp. Phys.*, 375:1339–1364, 2018.
- [65] T. Tangarife. *Kinetic theory and large deviations for the dynamics of geophysical flows*. PhD thesis, ENS de Lyon, ENS Lyon, <https://tel.archives-ouvertes.fr/tel-01241523>, 11 2015.
- [66] M. Tao. Hyperbolic periodic orbits in nongradient systems and small-noise-induced metastable transitions. *Physica D: Nonlinear Phenomena*, 363:1–17, 2018.
- [67] H. Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, 478(1):1–69, 2009.
- [68] E. Vanden-Eijnden and M. Heymann. The geometric minimum action method for computing minimum energy paths. *J. Chem. Phys.*, 128:061103, 2008.
- [69] E. Vanden-Eijnden and M. Heymann. The geometric minimum action method: A least action principle on the space of curves. *Comm. on Pure and Applied Math.*, 61(8):1052–1117, 2008.
- [70] W. Zhang, T. Li, and C. Schütte. Solving eigenvalue PDEs of metastable diffusion processes using artificial neural networks. *J. Comput. Phys.*, 465:111377, 2022.
- [71] X. Zhou and W. E. Study of noise-induced transitions in the Lorenz system using the minimum action method. *Commun. Math. Sci.*, 8:341–355, 2010.
- [72] K. Zubov, Z. McCarthy, Y. Ma, F. Calisto, V. Pagliarino, S. Azeglio, L. Bottero, E. Luján, V. Sulzer, A. Barambe, N. Vinchhi, K. Balakrishnan, D. Upadhyay, and C. Rackauckas. NeuralPDE: Automating Physics-Informed Neural Networks (PINNs) with Error Approximations. *arXiv:2107.09443v1*, 2021.