

Research Statement

Analysis and Computation of
Maximum Likelihood Transition Curves

Matthias Heymann, Duke University

April 26, 2010

Introduction. Dynamical systems are often subject to random perturbations. Even when these perturbations have small amplitude, they have a profound impact on the dynamics on the appropriate time scale. For instance, perturbations result in transitions between stable equilibrium points of the unperturbed (deterministic) system that would otherwise be impossible (see Fig. 1). Such transitions are responsible for metastable phenomena observed in nature: regime changes in climate, nucleation events during phase transitions, conformation changes of biomolecules, and bistable behavior in genetic switches are just a few examples among many others.

When the amplitude of the random perturbations is small, the Freidlin-Wentzell theory of large deviations provides the right framework to understand their effects on the dynamics [1]. In a nutshell, this theory builds on the property that rare (unlikely) events, when they occur, do so with high probability by following the pathway that is least unlikely. This makes rare events predictable, in a way that the Freidlin-Wentzell theory quantifies.

The central object in the theory is an action functional whose minimum (subject to appropriate constraints) gives an estimate of the probability and the rate of occurrence of the rare event, and whose minimizer gives the pathway of maximum likelihood by which this event occurs. A key practical question then becomes how to analyze and numerically compute the minimum and minimizer of the Freidlin-Wentzell action functional. This is the main topic of my research.

Summary of my results. In my PhD thesis [2, 3] my advisor Eric Vanden-Eijnden (NYU) and I first reformulated this minimization problem in a way that guarantees the existence of a minimizer in more cases. Based on this formulation we then designed an algorithm for numerically computing the minimizer (i.e. the most likely transition curve). Finally, we applied our algorithm to examples with various types of dynamics, such as stochastic differential equations (SDE), stochastic partial differential equations (SPDE), and continuous-time Markov jump processes. Our applications arose from questions in physics, biology, and mathematical finance.

In my most recent work [4], a 153 page monograph submitted to the *Memoirs of the AMS*, I demonstrate also the analytical advantages of our reformulation: First I developed easy-to-use methods for proving the existence of a minimizer of my functional (whereas the original formulation does not have a minimizer in most cases of practical interest). Then I proved various properties of the minimizer. In fact, my current work applies to a large class of functionals on the space of unparameterized oriented rectifiable curves, and it thus has applications also in other fields, such as the computation of the distance between two points or sets with respect to a Riemannian metric, or of the instanton by which quantum tunnelling occurs (i.e. the minimizer of the Agmon distance).

1 Previous work

Example 1: SDEs. Consider a stochastic dynamical system with small noise in \mathbb{R}^n , such as the stochastic differential equation (SDE)

$$dX_t^\epsilon = b(X_t^\epsilon) dt + \epsilon dW_t, \quad X_{t=0}^\epsilon = x_1, \quad (1)$$

where $(W_t)_{t \geq 0}$ is a Brownian motion, where $\epsilon > 0$ is a small parameter, and where the drift vector field $b \in C^1(\mathbb{R}^n, \mathbb{R}^n)$ has two stable equilibrium points x_1 and x_2 (see Fig. 1). The presence of small noise allows for rare transitions from x_1 to x_2 (such as the green wiggly curve) that would be impossible without the noise, and one is interested in the frequency and the most likely path of these rare transitions, in the limit as $\epsilon \rightarrow 0$. The answer to both is given within the framework of the Wentzell-Freidlin theory [1], the key object being the quasipotential

$$V(x_1, x_2) = \inf_{\substack{T > 0 \\ \chi \in \tilde{C}_{x_1}^{x_2}(0, T)}} \frac{1}{2} \int_0^T |b(\chi(t)) - \dot{\chi}(t)|^2 dt, \quad (2)$$

where $\tilde{C}_{x_1}^{x_2}(0, T)$ denotes the space of all absolutely continuous functions $\chi : [0, T] \rightarrow \mathbb{R}^n$ fulfilling $\chi(0) = x_1$ and $\chi(T) = x_2$. The expected time until the first transition event scales like $\exp(V(x_1, x_2)/\epsilon^2)$ as $\epsilon \rightarrow 0$.

An unpleasant feature of this formulation is that the minimization problem (2) does not have a minimizer (T^*, χ^*) , the main reason being that $\dot{\chi}$ likes to vanish at x_1 and x_2 , and also at the saddle point along the way, and so T likes to be (doubly) infinite. This is a major problem for both analytical and numerical work.




Figure 1: A typical rare transition event: The drift vector field b (given by the black flowlines) has two attractors, and there is one saddle point on the separatrix between their basins of attraction. The green wiggly curve is a typical example of a transition curve, which for small $\epsilon > 0$ can be shown to be close to the minimum action curve (red) with high probability.

Prior state of art: Algorithms. Obtaining reliable information on rare transitions by straightforward simulation of (1) is not possible since (by the very definition of these events) one would have to simulate (1) for a very long time only to obtain a single transition event. Therefore it is necessary to solve the minimization problem (2) instead. At the time I started my PhD, the most prominent approaches for computing $V(x_1, x_2)$ were the following:

(i) In the special case of a gradient drift, i.e. if $b(x) = -\nabla V(x)$ for some potential $V(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, one can show that to minimize (2) one must have $\dot{\chi}(t) \parallel b(\chi(t))$ everywhere along the curve (i.e., the curve of χ is always either parallel or antiparallel to the drift). The string method algorithm [5] exploits this geometric property: Starting from an initial curve (“string”), in every iteration the discretization points first follow the flow b for a short distance and are then equidistantly redistributed along the updated curve. After convergence (i.e. when these two steps cancel each other out) the string will have the desired geometric property.

While the string method is still the fastest method available for SDEs with gradient drift and additive noise, it is based on a geometric property of transition curves that is false in the general case, and thus it is impossible to generalize the string method to arbitrary drift vector fields b , to multiplicative (i.e. state dependent) noise, or to stochastic dynamics other than SDEs, such as the continuous-time Markov jump processes discussed below.

(ii) The Minimum Action Method (MAM) [6] is an inexact algorithm for solving (2) in the general setting (i.e. for general b , and even for multiplicative noise). A large but finite $T > 0$ is held fixed, the integral in (2) is discretized at equidistant time intervals, and a steepest descent method is applied to minimize this discretized version of the integral. The shortcomings of this approach are (a) the error introduced by not varying T , and (b) that the discretization points will accumulate wherever the curve is slow (e.g. near x_1 , x_2 and the saddle point), thus increasing the required number of discretization points unnecessarily. For a given number of discretization points, the error from (a) cannot be made arbitrarily small since increasing T will make the effect in (b) get even worse.

(iii) Shooting methods were also used [7] to solve the boundary value ODE (the Euler-Lagrange equations) obtained by setting the variation of the integral in (2) equal to zero. This approach however quickly becomes impractical in dimensions higher than 2 (and thus in most situations of practical interest).

Geometric reformulation. (See Appendix A for details on the derivation.) The main achievement in my PhD thesis [2, 3] was to combine the exactness and the optimal use of the discretization points in the string method with the generality of the MAM. To do so, we first developed a characterization of the maximum likelihood transition curve that only takes into account the *location* of χ but not its parameterization by time. We rephrased (2) as

$$V(x_1, x_2) = \inf_{\gamma \in \Gamma_{x_1}^{x_2}} S(\gamma), \quad \text{where} \quad S(\gamma) := \int_0^1 (|b(\varphi)| |\varphi'| - \langle b(\varphi), \varphi' \rangle) d\alpha, \quad (3)$$

where $\Gamma_{x_1}^{x_2}$ denotes the space of all (unparameterized) oriented rectifiable curves leading from x_1 to x_2 , and where $\varphi : [0, 1] \rightarrow \mathbb{R}^n$ is any absolutely continuous parameterization of γ . (It is easy to see that $S(\gamma)$ does not depend on the specific choice of φ .)

Since the time-parameterization of γ does not enter this formulation, the above-mentioned main reason for the non-existence of a minimizer in (2) (and thus the reason for both problems with the MAM) has been eliminated. The minimizer γ^* of (3) (the red curve in Fig. 1) – if it exists – can be

interpreted as the maximum likelihood transition *curve*: Rare transition events, once they occur, will with high probability follow γ^* .

In fact, this technique of turning a double minimization problem of the type (2) into a simple minimization problem of the type (3) generalizes also to integrands other than the specific one given in (2), see Appendix A. The integrand in (3) may not always be available explicitly, but this has so far never turned out to be an issue for my analytical or my numerical work.

The gMAM algorithm. Based on this reformulation we then designed an effective algorithm (called the geometric Minimum Action Method, or gMAM) [2, 3, 8, 9, 10] for numerically computing the minimizer γ^* of (3). The algorithm is a blend of the string method and the MAM: An initial curve $\gamma \in \Gamma_{x_1}^{x_2}$ is repeatedly updated by (i) moving it into the direction of steepest decent (i.e. minus the variation of $S(\gamma)$) while keeping its end points fixed, (ii) redistributing the discretization points of the curve equidistantly. This second step, which does not change the action of the curve, not only leads to an optimal use of the chosen number of discretization points, but it was also found to improve stability. Of course, the gMAM algorithm generalizes also to actions $S(\gamma)$ other than the specific one given in (3).

Example 2: Continuous-time Markov jump processes. To get an idea of the versatility of my approach, consider a biological or chemical system in which one is interested in the numbers $X = (X_1, \dots, X_n) \in \mathbb{N}_0^n$ of n different types of molecules within a cell or test tube. Chemical reactions occurring at certain rates $\lambda_i(X)$ cause discrete jumps in X , say of size $e_i \in \mathbb{Z}^n$, and thus such systems are often modelled as a continuous-time Markov jump process.

If the total number N of molecules in the system is large then the concentrations $x = X/N$ follow a jump process with small jump vectors e_i/N and corresponding jump rates that often scale like $N \cdot \nu_i(x)$ as $N \rightarrow \infty$. In this limit, random fluctuations of x become smaller, and the system approaches the limiting ODE

$$\dot{x} = b(x) := \sum_{i=1}^{\#\text{reac}} \nu_i(x) e_i. \quad (4)$$

If b has again two attractors as in Fig. 1, then for large N noise-driven transitions from one attractor to the other become rare. Large deviation theory gives us a formula for $V(x_1, x_2)$ similar to (2) also in this case [11], and my approach again allows us to rewrite it in parameterization-free form as in (3) (with new integrands in both (2) and (3) which are in fact both given only implicitly via some Hamiltonian). We can now use the gMAM algorithm to solve the modified minimization problem (3), and thus to compute the most likely transition curve and the rate at which this rare transition occurs.

In addition, one can use the gMAM to determine which reaction *caused* the rare transition [10]. To understand this, observe that if all reactions occurred at their respective expected rates $N \cdot \nu_i(x)$ then the system would follow (4) and thus stay in its original stable state. Thus, a transition can only happen if at least one reaction occurs at a rate that differs from the expected rate. My work [10] tells us how to find those reactions that are most likely to have atypical reaction rates during the transition. This technique may help design more reliable bistable networks.

One example that fits into this framework is the genetic switch, as e.g. discussed in [2, 3, 7, 10]: Two types of proteins within a cell can bind to the respective other protein’s operator site on the DNA, thus reducing that other protein’s further production. As a consequence, an elevated density

of one type of protein typically reduces the production and thus the density of the other, which in turn causes its own production to increase. This leads to a positive feedback loop with two stable states: one with many proteins of type A and few of type B, and another stable state with many of type B and few of type A.

Finally, let me point out that it is *not* okay in this context to approximate this jump process by an SDE with matching drift and noise covariance matrix: While such an approximation may be justified in a certain sense for questions about the *typical* behavior of the system, here we are dealing with rare and thus *atypical* events. As a consequence, the most likely transition curves in the exact and the “approximated” system can be found to differ significantly from each other in general, further justifying our efforts to adjust the gMAM to jump processes.

Existence and properties of minimum action curves – Overview. (See Appendix B for details on my techniques.) While the parameterization-free formulation (3) is more feasible to have a minimum action curve, i.e. a curve $\gamma^* \in \Gamma_{x_1}^{x_2}$ such that

$$S(\gamma^*) = \inf_{\gamma \in \Gamma_{x_1}^{x_2}} S(\gamma),$$

a rigorous proof of the existence of such a minimizer γ^* is still necessary. This is not only of theoretical importance: If no minimizer exists then naturally the gMAM algorithm will fail to find one, but the user may (for lack of better knowledge) blame numerical problems instead. In this situation, solid knowledge about when and why minimizers fail to exist can spare the user from further desperate efforts to find the minimizer by tweaking the parameters of the algorithm.

For similar reasons, any analytically obtained knowledge about properties of γ^* is helpful also in practice: The user can use such knowledge either (i) to gain confidence in his numerically obtained curve by checking if it has these properties as well, or (ii) to speed up his search for a minimizer by restricting it to only curves with these properties.

My most recent work [4], a 153 page monograph completed in 04/2010, is an extensive study of these two problems. First I developed tools (based on the flowline diagram of b) with which one can determine whether a minimizer exists. Although parts of my proofs are long and technical, my criteria are easy to understand and easy to use in practice (see Appendix B for details). Figures 3-4 on pages 14-15 show a few exemplary systems in \mathbb{R}^2 to which my criteria can successfully be applied. Then I proved some properties of γ^* , in particular I showed that its first and last hitting point of the separatrix between x_1 and x_2 must be points where b vanishes. (In the example of Fig. 1 the first and last hitting points coincide, but this need not be the case in general.)

In fact, my existence proof applies to a large class of parameterization-free “geometric” action functionals $S(\gamma)$: A geometric action S is a mapping that assigns to every unparameterized rectifiable oriented curve γ in \mathbb{R}^n a number $S(\gamma) \in [0, \infty)$. It is defined via a curve integral

$$S(\gamma) := \int_{\gamma} \ell(z, dz) := \int_0^1 \ell(\varphi, \varphi') d\alpha, \quad (5)$$

where $\varphi : [0, 1] \rightarrow \mathbb{R}^n$ is any absolutely continuous parametrization of γ , and where the local action $\ell \in C(\mathbb{R}^n \times \mathbb{R}^n, [0, \infty))$ must have the properties

- (i) $\forall x, y \in \mathbb{R}^n \ \forall c \geq 0 : \ell(x, cy) = c\ell(x, y)$,
- (ii) for every fixed $x \in \mathbb{R}^n$ the function $\ell(x, \cdot)$ is convex.

While (i) guarantees that the second integral in (5) is independent of the particular choice of φ , (ii) is necessary to ensure that S is lower semi-continuous in a certain sense. The notion of the drift vector field b is then extended also to general geometric actions (see Appendix B).

Examples outside of the context of large deviation theory include: (i) the trivial example $\ell(x, y) = |y|$, in which case $S(\gamma)$ is just the Euclidean length of γ , or more generally, $\ell(x, y) = |y|_{g_x}$ for any Riemannian metric g ; (ii) $\ell(x, y) = U(x)|y|$ for some potential $U: \mathbb{R}^n \rightarrow [0, \infty)$ which may vanish at some isolated points (the corresponding action is the Agmon distance used to determine the instanton by which quantum tunnelling occurs).

Prior state of art: Existence theory. While the existence of minimum action curves has likely already been addressed elsewhere for the special case $\ell(x, y) = |y|_{g_x}$ (where the proof is comparably easy), to the best of my knowledge nobody has ever investigated a class of geometric action functionals as general as the one treated in my work [4]. First, note that we do *not* require that $\ell(x, -y) = \ell(x, y)$, and thus S need not be invariant under changing the direction of a curve, i.e. we may have $S(\gamma) \neq S(-\gamma)$ in general.

More importantly, the local action $\ell(x, y)$ may vanish in certain “drift directions” $y \neq 0$. E.g., in the example of the large deviation geometric action S for SDEs in (3), given by $\ell(x, y) = |b(x)||y| - \langle b(x), y \rangle$, we have $\ell(x, y) = 0$ for $y = b(x)$, and the same can be shown in the case of a continuous-time Markov chain, with $b(x)$ given by (4). In fact, this property is generic to large deviation geometric actions: For stochastic dynamics, the drift vector field b given by the zero-noise limit $\dot{x} = b(x)$ is the direction which the system can follow without the aid of the noise (as $\epsilon \rightarrow 0$), and so any curve segment that follows the drift direction has zero cost.

It is just this last property of the local action that makes my proofs in [4] so much more difficult. Therefore, mastering these difficulties and providing general and easy-to-use existence criteria that also apply to large deviation geometric actions is what I consider the most important contribution of my work [4].

2 Future work

My previous results have opened up two lines of research that I am going to pursue – one in pure and one in applied mathematics:

Pure mathematics

(i) *Dissemination of results.* In its current math-heavy form, my work [4] is likely not going to reach the mathematically less-trained part of the applied community. As explained above however, a solid understanding of when and why minimum action curves may fail to exist can be invaluable to those who use algorithms for finding them. Therefore, after publishing [4] in one piece I will have to target scientists from various fields, by writing a sequence of shorter, application-oriented follow-up papers that focus only on certain aspects of my theory (instead of aiming for maximal generality) and that refer to [4] for the rigorous proofs.

(ii) *Extension of my existence theory.* I will extend my existence theory developed in [4] to deal also with certain currently untractable constellations. Two examples are the following (see the end of Appendix B for some additional comments):

(ii.1) *Limit cycles.* The minimization problem of finding the minimum action curve from some attractor of the drift vector field to a limit cycle around it (see Fig. 2a) is “solved” by a spiral that




Figure 2: Two systems to which my criteria cannot be applied: (a) an attractor enclosed by a limit cycle; (b) an attractor enclosed by a closed chain of flowlines.

only approaches the limit cycle asymptotically. However, the current framework in [4] does not allow for such solutions and therefore proves that “no minimizing curve (of finite length) exists.” Luckily, my theory already provides all the tools necessary to deal with this special case with a little extra effort. This mostly risk-free project will be the subject of a future paper. On the other hand, the equally important problem of reaching a limit cycle in dimensions higher than two may involve currently unpredictable difficulties, since then transition curves are not confined to a compact region (such as the disk in Fig. 2a).

(ii.2) Closed chains of flowlines. A problem that is yet to be understood is the one of reaching a chain of flowlines that connect four saddle points, as illustrated in Fig. 2b. With the technique used for (ii.1), it is not hard to show that a “minimizer” exists also here, but it is not clear whether this minimizer takes the form of a minimizing spiral again, or of a curve of finite length which (as shown in [4]) would have to end in one of the four saddle points. The main issue here is the non-locality of this problem: No matter how close to the target set we pick our initial point, the minimizing curve may lead far away from it (to one of the saddle points). Dealing with this non-locality will require entirely new techniques, the development of which would be a worthy topic for a doctoral thesis.

Applied mathematics

(iii) Applications. I will apply the gMAM algorithm to more problems arising in practice, such as in biology, chemistry or physics. Two examples of specific collaborations that have already been initiated are the following:

(iii.1) Dynamics of electric field domains in semiconductor nanostructures. In collaboration with Prof. Stephen Teitworth (Duke University, Physics Department) I will investigate a certain semiconductor nanostructure consisting of 20 alternating layers of two different materials (modelled by a 20-dimensional SDE). Under applied voltage, such a structure exhibits an electric field distribution that is highly non-uniform in space. Depending on how the voltage was reached, the electric field distribution may be in one of several metastable configurations. The gMAM can be used to find the rate of noise-driven transition between two such states, and we will be able to compare the outcome of the gMAM algorithm with actual measurements of Prof. Teitworth’s experiments.

(iii.2) *Transitions in a perturbed Ginzburg-Landau SPDE.* Another ongoing collaboration (with Prof. Eric Vanden-Eijnden, NYU) concerns the analysis of nucleation events in an SPDE system with shear flow. Consider an SPDE for a function $\phi(x, y, t) : [0, 1] \times [0, 1] \times [0, \infty) \rightarrow \mathbb{R}$ of the type

$$\dot{\phi} = \kappa \Delta \phi - \nabla V_{c_1}(\phi) + c_2 \sin(2\pi y) \partial_x \phi + \epsilon \eta, \quad (6)$$

where we assume periodic boundary conditions in $(x, y) \in [0, 1]^2$, $\Delta = \partial_x^2 + \partial_y^2$ is the Laplacian, $\kappa > 0$ is the diffusion coefficient, $V : \mathbb{R} \rightarrow \mathbb{R}$ is a potential function, and $\eta(x, y, t)$ is spatio-temporal white-noise with covariance $\langle \eta(x, y, t), \eta(x', y', t') \rangle = \delta(x - x') \delta(y - y') \delta(t - t')$. For $c_2 = 0$, the equation (6) is the steepest-descent dynamics on the Ginzburg-Landau energy functional $E_{c_1}(\phi) = \int_{[0, 1]^2} [\frac{1}{2} |\nabla \phi|^2 + V_{c_1}(\phi)] dx dy$. The additional term $c_2 \sin(2\pi y) \partial_x \phi$ mimics the presence of a nonlinear shear flow with shear rate c_2 in the x -direction. In [8] we adjusted the gMAM algorithm to this SPDE setting: In the specific c_1 -independent case $V(\phi) = \frac{1}{4}(1 - \phi^2)^2$ [12, 13, 14] we investigated the transitions between the two metastable states $\phi_1(x, y) \equiv -1$ and $\phi_2(x, y) \equiv 1$. We now want to investigate the effect of breaking the symmetry of the potential V by adding a tilt (with parameter c_1).

(iv) *Optimization of control parameters.* In addition, we can use the gMAM to optimize the transition rate with respect to one or more parameters, such as the parameters c_1 and c_2 in the example (iii.2) above. To do so, one can at each iteration of the gMAM add the updating step $\vec{c}_{k+1} = \vec{c}_k - \delta \cdot \nabla_{\vec{c}} S(\gamma_k, \vec{c}_k)$ for some small stepsize $\delta > 0$. See [8] for a proof of concept in the SPDE case (where we only minimize over c_2). Doing the same for jump processes may require some extra effort since the integrand of S is not available explicitly.

Work for the more distant future. In the long run it will be important to adjust both my existence theory and the gMAM algorithm to situations in which the geometric action $S(\gamma)$ is defined only for curves γ that are contained in some fixed *manifold* in \mathbb{R}^n . (This situation can occur when not every component in the stochastic system is subject to noise, in which case the stochastic process may be living on a manifold.) However, at present it is more important to fully illuminate and understand every aspect of my current setup before I move on to this next level.

3 Conclusions

A new analytic approach (the geometric reformulation) has led to a versatile algorithm (gMAM) for computing the frequencies and pathways of rare transition events in stochastic dynamical systems subject to small noise. Motivated by the geometric large deviation actions, a theoretical framework for proving the existence or non-existence of minimum action curves has been developed from scratch.

As future work, besides disseminating my previous results, several explicit projects are suggested to push forward the theoretical framework, to apply the gMAM to several problems arising in practice, and to further develop the gMAM and adjust it to the specific needs arising from applications. Additional projects will naturally develop as I begin interacting with scientists from various departments of my new employer.

Appendix

*Additional detailed information
on some of my techniques and results*

A Geometric reformulation – Details

To understand the workings of the geometric reformulation, suppose one is asked to find the minimizer of the real-valued function $f(x, y) := x^2 + e^{-y}$ among all $(x, y) \in \mathbb{R}^2$, and that the use of the symbol ∞ is forbidden. Saying that “there is no minimizer (x^*, y^*) ” is correct but unsatisfactory, since we see that in a way we must have $x^* = 0$. This can be made rigorous by proving that $\hat{f}(x) := \inf_{y \in \mathbb{R}} f(x, y) = x^2$ and then stating that $x^* = 0$ is a minimizer of \hat{f} , so that $\hat{f}(x^*) = \inf_x \hat{f}(x) = \inf_{x, y} f(x, y)$. In other words, one must identify the component of (x, y) that is responsible for the non-existence of a minimizer (x^*, y^*) and carry out the minimization over that component analytically. For more complicated functions, it may require some smart change of variables to make both these steps work.

Now to minimize an action $S_T(\chi) := \int_0^T L(\chi, \dot{\chi}) dt$ over all (T, χ) , observe that every time-parameterized curve $\chi: [0, T] \rightarrow \mathbb{R}^n$ can be represented by a triple (γ, T, β) consisting of a curve γ , the total time T needed to traverse it, and a non-decreasing time-parameterization $\beta \in C([0, T], [0, 1])$. Explicitly, we have $\chi(t) = \varphi_\gamma(\beta(t))$ for $\forall t \in [0, T]$, where $\varphi_\gamma: [0, 1] \rightarrow \mathbb{R}^n$ is the arclength parameterization of γ . (This bijection $(T, \chi) \leftrightarrow (\gamma, T, \beta)$ is our “change of variables.”) Since we identified the time-parameterization (i.e. the components (T, β)) as the reason for the non-existence of minimizers, we follow the above strategy to minimize over (T, β) analytically. In the SDE case with S_T given by the integral in (2) we find that

$$\begin{aligned} S(\gamma) &:= \inf_{\substack{T > 0 \\ \beta: [0, T] \rightarrow [0, 1]}} S_T(\varphi_\gamma \circ \beta) = \inf_{T, \beta} \frac{1}{2} \int_0^T |b(\varphi_\gamma(\beta(t))) - \varphi'_\gamma(\beta(t))\beta'(t)|^2 dt \\ &= \inf_{T, \beta} \frac{1}{2} \int_0^1 |b(\varphi_\gamma(\tilde{\beta})) - \varphi'_\gamma(\tilde{\beta}) \cdot \beta'|^2 \frac{d\tilde{\beta}}{\beta'} \Big|_{\beta' = \beta'(t(\tilde{\beta}))} \\ &= \inf_{T, \beta} \frac{1}{2} \int_0^1 \left(\frac{|b(\varphi_\gamma)|^2}{\beta'} - 2\langle b(\varphi_\gamma), \varphi'_\gamma \rangle + |\varphi'_\gamma|^2 \beta' \right) d\tilde{\beta} \Big|_{\beta' = \beta'(t(\tilde{\beta}))}. \quad (7) \end{aligned}$$

Minimizing the last integrand with respect to β' (pointwise for every $\tilde{\beta}$), we find that we must have $\beta' = |b(\varphi_\gamma)|/|\varphi'_\gamma|$, and plugging this into (7) we arrive at the representation (3) of $S(\gamma)$, which (just as \hat{f} above) fulfills $\inf_\gamma S(\gamma) = \inf_{\gamma, T, \beta} S_T(\varphi_\gamma \circ \beta) = \inf_{T, \chi} S_T(\chi)$.

The same strategy works for general integrands L of S_T , but β' and thus the integrand $\ell(x, y)$ of S may not always be available explicitly. Of course, making the above steps rigorous requires some extra work, see [2, 3].

B Existence of minimum action curves – Details

I will now give some details on the criteria developed in [4]. To keep things short, this appendix is not aiming for maximal generality, and in particular it will state my results as they hold for large deviation geometric actions. For simplicity I will also omit some minor technical details. A first result which is relatively easy to obtain is the following.

Lemma 1. *If there exists a minimizing sequence $(\gamma_n)_{n \in \mathbb{N}}$ of (3) that has uniformly bounded curve lengths, then there exists a minimizer $\gamma^* \in \Gamma_{x_1}^{x_2}$.*

Proof. If we parameterize the curves γ_n by arclength, i.e. with absolutely continuous functions $\varphi_n : [0, 1] \rightarrow \mathbb{R}^n$ fulfilling $|\varphi'_n| \equiv \text{length}(\gamma_n)$ a.e. on $[0, 1]$, then we can use Arzelà-Ascoli's theorem to show that there exists a subsequence $(\varphi_{n_k})_{k \in \mathbb{N}}$ that converges uniformly to some limit φ^* . It is not hard to show that this limit is absolutely continuous as well and thus a parameterization of a rectifiable curve γ^* . The lower-semicontinuity of S then implies that γ^* is a minimizer. \square

In practice however, this criterion alone is of little use since minimizing sequences are not at our direct disposal and can thus be hard to control. Instead, one would rather like to have some higher-level criteria that are based on some explicitly available key ingredient of the local action ℓ . What could this key ingredient be?

The proofs of these higher-level criteria work by constructing a minimizing sequence with uniformly bounded curve lengths, so that Lemma 1 implies the existence of a minimizer. Now recall that $\ell(x, y)$ vanishes whenever y aligns with $b(x)$, which allows for long curves with zero action (the flowlines). Depending on the flowline diagram of b , this may force minimizing sequences $(\gamma_n)_{n \in \mathbb{N}}$ to have *unbounded* curve lengths, and thus the flowline diagram of b is the key ingredient of my criteria. For general geometric actions, we define the drift as follows.

Definition 1. *A vector field $b \in C^1(\mathbb{R}^n, \mathbb{R}^n)$ is called a drift of S if*

$$\forall \text{ compact } K \subset \mathbb{R}^n \exists c > 0 \forall x \in K \forall y \in \mathbb{R}^n : \quad \ell(x, y) \geq c(|b(x)||y| - \langle b(x), y \rangle), \quad (8)$$

where $\ell(x, y)$ is the local action of S .

We observe that wherever $b(x) \neq 0$, the drift $b(x)$ is the only candidate for a direction $y \neq 0$ in which $\ell(x, y)$ may vanish.

Since the right-hand side of (8) is a constant multiple of the local large deviation geometric action of the SDE (1), we see that the vector field b in (1) is clearly a drift also in this generalized sense (take $c = 1$). Furthermore, one can show that the drift b of an SDE with multiplicative noise is a drift also in this generalized sense, as is the function b given by (4) in the case of a Markov jump process. Note that the drift is not unique in general (e.g., $b \equiv 0$ is always a drift). But as it turns out, drift vector fields with only as few roots as necessary give us the strongest existence results, and in that sense our natural choice for the drift in the large deviation setup is in fact optimal. Next we introduce a local existence property.

Definition 2. *We say that a point $x \in \mathbb{R}^n$ has local minimizers if*

$$\exists r, \eta > 0 \forall x_1, x_2 \in \bar{B}_r(x) \exists \gamma^* \in \Gamma_{x_1}^{x_2} : \quad S(\gamma^*) = \inf_{\gamma \in \Gamma_{x_1}^{x_2}} S(\gamma) \quad \text{and} \quad \text{length}(\gamma^*) \leq \eta. \quad (9)$$

In other words, for start and end points chosen *near* x there exists a minimizer γ^* with controllable length. Using this property, one can use a compactness argument to prove the following global theorem (i.e. for x_1 and x_2 *not* close to each other).

Theorem 1 (Existence Theorem). *Let $x_1, x_2 \in \mathbb{R}^n$, let $K \subset \mathbb{R}^n$ be a compact set consisting only of points that have local minimizers, and let us assume that the corresponding minimization problem has a minimizing sequence $(\gamma_n)_{n \in \mathbb{N}}$ such that $\gamma_n \subset K$ for $\forall n \in \mathbb{N}$. Then there exists a minimizer.*

Proof. Cover K by finitely many (say N) open balls $B_{r_x}(x)$ obtained from (9). Then cut each γ_n into at most N pieces whose start and end point is contained in the same ball, and replace each

segment of γ_n by the corresponding minimizer given by (9). The newly obtained curves $\tilde{\gamma}_n$ fulfill $S(\tilde{\gamma}_n) \leq S(\gamma_n)$ and thus form a new minimizing sequence, and since their lengths are uniformly bounded by the sum of the bounds from (9) for each ball, Lemma 1 implies the existence of a minimizer. \square

The decisive advantage of Theorem 1 over Lemma 1 is that the bounded-length-condition of the minimizing sequence is replaced by the weaker condition that $\gamma_n \subset K$ for $\forall n \in \mathbb{N}$. This in turn boils down to the following simple estimate.

Lemma 2. *Suppose there exists a $\gamma_0 \in \Gamma_{x_1}^{x_2}$ with $\gamma_0 \subset K$ such that every curve $\gamma \in \Gamma_{x_1}^{x_2}$ that leaves K along its way has an action $S(\gamma) \geq S(\gamma_0)$. Then there exists a minimizing sequence $(\gamma_n)_{n \in \mathbb{N}}$ with $\gamma_n \subset K$ for $\forall n \in \mathbb{N}$.*

Proof. Let $(\gamma_n)_{n \in \mathbb{N}}$ be any minimizing sequence. If we replace every curve γ_n that is not entirely contained in K by γ_0 then we only reduce the action. Thus we obtain a new minimizing sequence that is now entirely contained in K . \square

For “contracting” flowline diagrams like the ones depicted in Figures 3-4 (where all flowlines far outside lead to some region in the center), Lemma 2 is rather intuitive to see: Pick any fixed $\gamma_0 \in \Gamma_{x_1}^{x_2}$, say the straight connection line from x_1 to x_2 , then for sufficiently large K , any curve that leaves K will have to move against the direction of the flow for so long that its action is larger than $S(\gamma_0)$.

This leaves us with the question how one can effectively prove that a given point has local minimizers. In [4], this is only where my hard work starts, and in a total of about 70 pages I proved the criteria listed in the Proposition below. Luckily, actually applying these criteria in practice is very easy, and it requires no knowledge of the proof and no technical estimates whatsoever. Instead, all one needs is a good understanding of the flowline diagram of a good choice of a drift b .

The key tool in two of the three criteria below is the following.

Definition 3. *A set $M \subset \mathbb{R}^n$ is called an admissible manifold if $\exists f \in C(\mathbb{R}^n, \mathbb{R})$ such that*

- (i) $M = f^{-1}(\{0\})$,
- (ii) M is compact,
- (iii) f is C^1 in a neighborhood of M , and
- (iv) for $\forall x \in M$ we have $\langle \nabla f(x), b(x) \rangle > 0$.

In other words, M is an $(n - 1)$ -dimensional compact C^1 -manifold with the property that the flowlines of b cross M always in the same direction (“in” or “out”). See the red lines in Figures 3-4 for examples. Our main criteria for finding points with local minimizers are the following:

Proposition 1. (i) *Let $x \in \mathbb{R}^n$ be such that $b(x) = 0$ and that all the eigenvalues of the matrix $\nabla b(x)$ have negative (positive) real parts. Then every point in its basin of attraction (repulsion) has local minimizers.*

(ii) *If M is an admissible manifold then every point on the flowlines emanating from M has local minimizers.*

(iii) *Let $x \in \mathbb{R}^n$ with $b(x) = 0$ be a saddle point such that all the eigenvalues of the matrix $\nabla b(x)$ have non-zero real parts, let W_s and W_u denote its global stable and unstable manifolds, respectively, and suppose that there exist admissible manifolds M_1, \dots, M_m such that every flowline in W_s and W_u intersects at least one of these admissible manifolds. Then x has local minimizers.*

Assuming that for every point x with $b(x) = 0$ the eigenvalues of the matrix $\nabla b(x)$ have non-zero real parts, these criteria are enough to show for each of the examples given in Fig. 3-4 that in fact *every* point in \mathbb{R}^2 has local minimizers.

Consider for example the flowline diagram in Fig. 3a: The points in the two basins of attraction have local minimizers by Proposition 1 (i). The three red lines are admissible manifolds, and we observe that every flowline on the stable and the unstable manifold of the saddle point (blue) eventually intersects one of them. Proposition 1 (ii) thus implies that every point on these flowlines has local minimizers, and Proposition 1 (iii) implies that the saddle point itself has local minimizers as well. Similar arguments can be used for the remaining examples in Fig. 3-4, with the admissible manifolds shown in those figures.

Observe that our criteria are insufficient to show that the points on the limit cycle in Fig. 2a have local minimizers: To use Proposition 1 (ii), we would have to find an admissible manifold that intersects the limit cycle, but we can easily convince ourselves that this is not possible. (Indeed, any loop that intersects the limit cycle once must intersect it at least a second time such that the drift flows in at one intersection and out at the other.) As a result, we cannot use our existence theorem.

There is a good reason for this: For large deviation geometric actions, these points can be shown not to have local minimizers, and one can prove that in fact no minimizer (of finite length) exists [4]. On the other hand, for the open problem in Fig. 2b our criteria fail as well, but it is not clear at this point if they should, i.e. whether these points have local minimizers or not, and whether or not a minimizer exists.

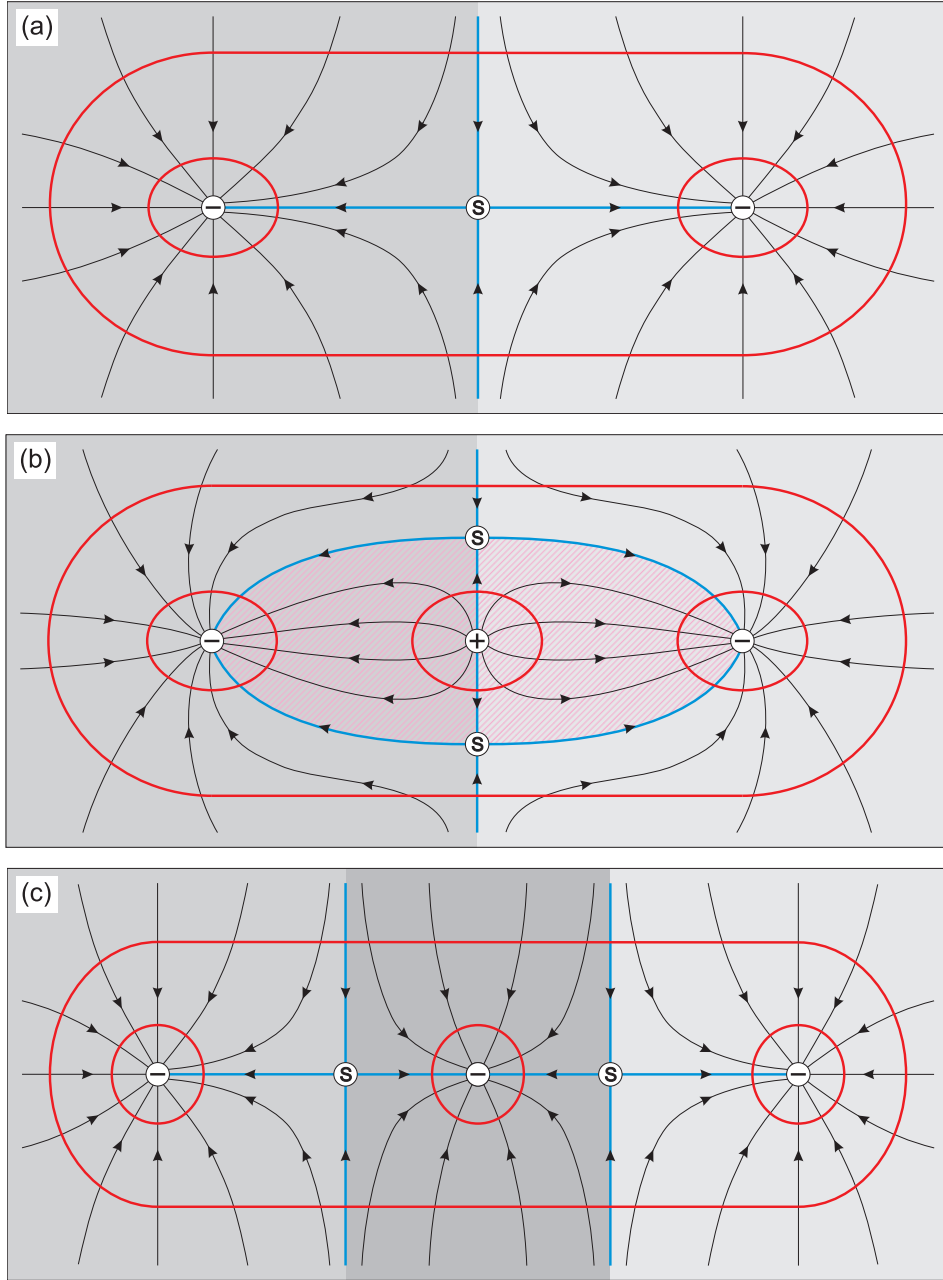


Figure 3: Three exemplary systems to which my criteria can successfully be applied. The black and blue lines are the flowlines of the drift vector field b , the red lines are admissible manifolds (as introduced in Appendix B). The roots of b are denoted by the symbols \ominus (attractor), \oplus (repeller) and \textcircled{S} (saddle point). The basins of attraction are colored in different shades of gray, the basins of repulsion are drawn in pink stripes.

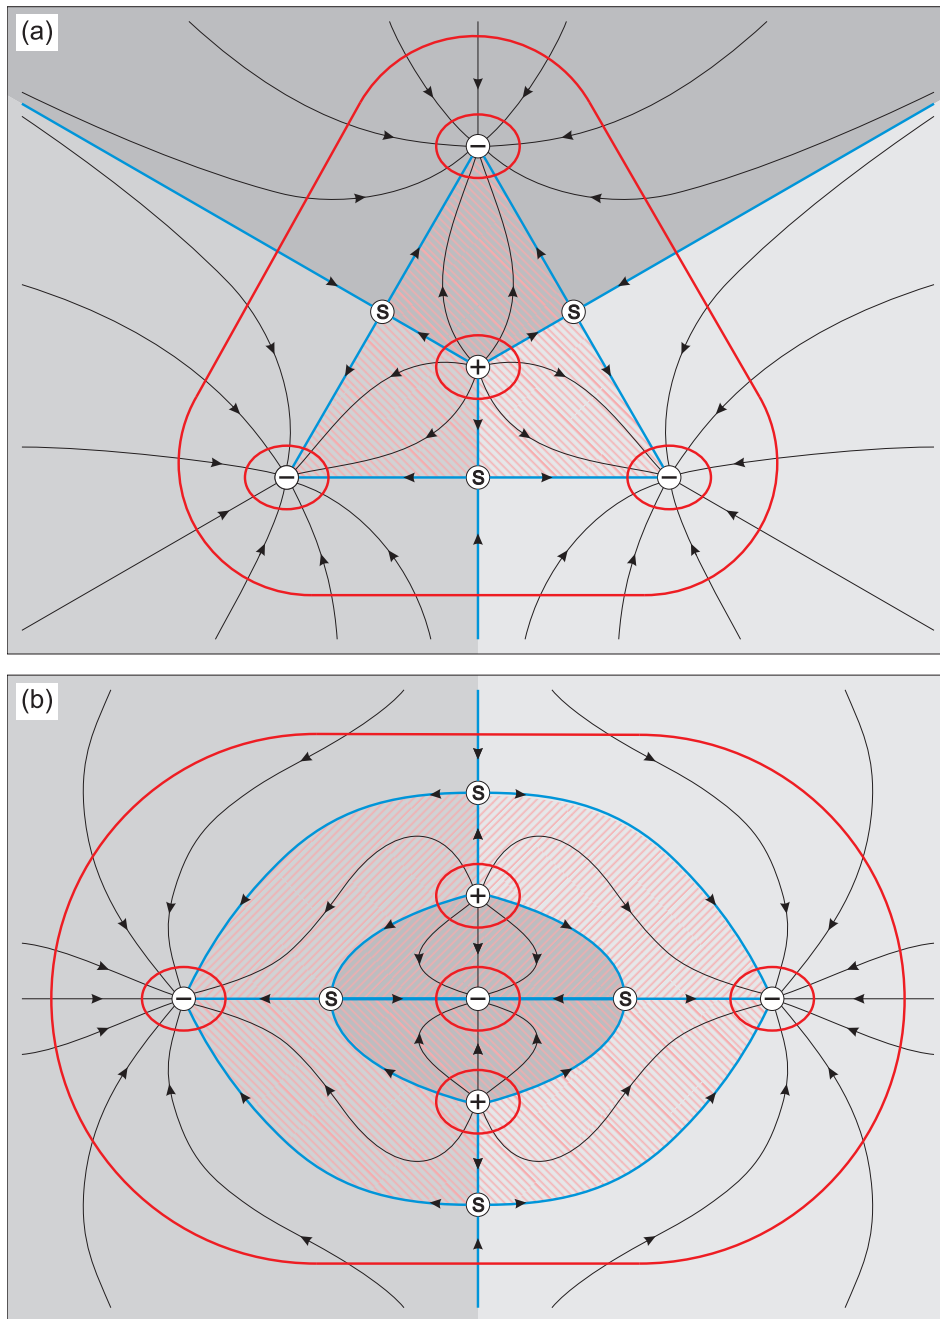


Figure 4: Two more exemplary systems (with three attractors each) to which my criteria can successfully be applied.

References

- [1] M. I. Freidlin and A. D. Wentzell, *Random perturbations of dynamical systems*, vol. 260 of *Grundlehren der Mathematischen Wissenschaften*. Springer, New York, second ed., 1998.
- [2] M. Heymann, *The geometric minimum action method: A least action principle on the space of curves*. PhD thesis, New York University, July 2007.
- [3] M. Heymann and E. Vanden-Eijnden, “The geometric minimum action method: A least action principle on the space of curves,” *Communications in Pure and Applied Mathematics*, vol. 61, pp. 1052–1117, August 2008.
- [4] M. Heymann, “Existence and properties of minimum action curves for degenerate Finsler metrics,” *submitted to the Memoirs of the AMS*, 04/26/2010.
- [5] W. E, W. Ren, and E. Vanden-Eijnden, “String method for the study of rare events,” *Physical Review B*, vol. 66, p. 052301, 2002.
- [6] W. E, W. Ren, and E. Vanden-Eijnden, “A minimum action method for the study of rare events,” *Communications in Pure and Applied Mathematics*, vol. 52, pp. 637–656, 2004.
- [7] D. M. Roma, R. A. O’Flanagan, A. E. Ruckenstein, and A. M. Sengupta, “Optimal path to epigenetic switching,” *Physical Review E*, vol. 71, p. 011902, 2005.
- [8] M. Heymann and E. Vanden-Eijnden, “Pathways of maximum likelihood for rare events in nonequilibrium systems – Application to nucleation in the presence of shear,” *Physical Review Letters*, vol. 100, no. 14, p. 140601, 2007.
- [9] E. Vanden-Eijnden and M. Heymann, “The geometric minimum action method for computing minimum energy paths,” *Journal of Chemical Physics*, vol. 128, p. 061103, February 2008.
- [10] M. Heymann and E. Vanden-Eijnden, “The sources of rare transitions in continuous-time Markov jump processes,” *In preparation*, 2009.
- [11] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis – Queues, Communication, and Computing*. Chapman & Hall, 1995.
- [12] C. Jarzynski, “Nonequilibrium Equality for Free Energy Differences,” *Physical Review Letters*, vol. 78, pp. 2690–2693, 1997.
- [13] J. Kurchan, “Fluctuation Theorem for stochastic dynamics,” *Journal of Physics A*, vol. 31, pp. 3719–3729, 1998.
- [14] G. E. Crooks, “Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences,” *Physical Review E*, vol. 60, pp. 2721–2726, 1999.

Click on the links to download the papers. (Be patient as downloading some papers can take 1-2 minutes.) For brief summaries, visit www.matthiasheyman.de (click on Mathematics, then Papers & Publications).