

Weierstraß-Institut
für Angewandte Analysis und Stochastik
Leibniz-Institut im Forschungsverbund Berlin e. V.

Preprint

ISSN 0946 – 8633

**On the relation between gradient flows and the large-deviation
principle, with applications to Markov chains and diffusion**

Alexander Mielke^{1,2}, Mark A. Peletier³, Michiel Renger¹

submitted: November 12, 2013

¹ Weierstraß-Institut
Mohrenstraße 39
10117 Berlin
Germany
E-Mail: alexander.mielke@wias-berlin.de
michiel.renger@wias-berlin.de

² Institut für Mathematik
Humboldt-Universität zu Berlin
Rudower Chaussee 25
12489 Berlin-Adlershof
Germany

³ Department of Mathematics and Computer Science and
Institute for Complex Molecular Systems (ICMS)
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven
the Netherlands
E-Mail: m.a.peletier@tue.nl

No. 1868

Berlin 2013



2010 *Mathematics Subject Classification.* 35Q82, 35Q84, 49S05, 60F10, 60J25, 60J27.

Key words and phrases. Generalized gradient flows, large deviations, convex analysis, particle systems.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: <http://www.wias-berlin.de/>

Abstract

Motivated by the occurrence in rate functions of time-dependent large-deviation principles, we study a class of non-negative functions \mathcal{L} that induce a flow, given by $\mathcal{L}(\rho_t, \dot{\rho}_t) = 0$. We derive necessary and sufficient conditions for the unique existence of a generalized gradient structure for the induced flow, as well as explicit formulas for the corresponding driving entropy and dissipation functional. In particular, we show how these conditions can be given a probabilistic interpretation when \mathcal{L} is associated to the large deviations of a microscopic particle system. Finally, we illustrate the theory for independent Brownian particles with drift, which leads to the entropy-Wasserstein gradient structure, and for independent Markovian particles on a finite state space, which leads to a previously unknown gradient structure.

1 Introduction

Gradient flows are an important subclass of evolution equations. A gradient-flow structure can be exploited to obtain additional information about the evolution, such as existence and stability of solutions (see e.g. the survey [AGS08]), or to derive limiting equations if some system parameter goes to zero [SS04, AMP⁺12, Mie13a]. Moreover, gradient flows can provide additional physical and analytical insight, such as the maximum dissipation of entropy and energy, or the geometric structure induced by the dissipation distance.

A special subclass is formed by gradient flows with respect to a Wasserstein-type metric. This class was first identified in the seminal work by Jordan, Kinderlehrer, and Otto [JKO98, Ott01] and later shown to encompass a wide range of known and unknown partial differential equations [AGS08, MMS09, PP10, Mie11b, LM12]. For a class of evolution equations defined on discrete spaces such as graphs, various generalizations exist that unite both continuous and discrete spaces into the same structure [Maa11, Mie13b, CHLZ12].

In many of these Wasserstein gradient flows entropy terms are present in the driving functional, which leads to diffusive behavior in the evolution. The notion of entropy is inherently related to microscopic particle systems, through Boltzmann's entropy formula, or through Sanov's large-deviation theorem. This suggests that a gradient flow that is driven by entropy may also be related to a microscopic particle system through its large deviations. In our earlier work we showed such a connection between large deviations and gradient flows, for a number of systems: the standard Fokker-Planck equation [ADPZ11, DLZ12, DLR13], a Fokker-Planck equation with decay [PRV13], a generalized Kramers equation [DPZ13a] and the Vlasov-Fokker-Planck equation [DPZ13b]. The paper [ADPZ12] and the dissertation [Ren13] discuss these and other examples in more detail.

In all these examples, the large-deviation rate functional can somehow be connected to a variational formulation of the limiting, deterministic evolution equation. In some cases this variational formulation is a gradient flow [ADPZ11]; in other cases it can be recognized as a natural generalization of the gradient-flow concept [ADPZ12, DPZ13b]; and in yet other cases it appears to have a different structure altogether [PRV13]. This all raises the question, that we address in this paper,

Does there exist a general connection between large deviations and (generalizations of) gradient flows?

(1.1)

Although this question is framed in terms of large deviations, the only ingredient from large deviations that we will actually use below is the variational form of the large-deviation rate function. To reflect this generality, we will use a more abstract point of view, assuming that some source of understanding (large deviations or otherwise) provides a variational formulation of an evolution equation in terms of what we call an *L-function*.

1.1 L-functions

Let \mathcal{Z} be some continuous state space, and let $\mathcal{L}(z, s)$ be a function of the state $z \in \mathcal{Z}$ and tangent vector s at z (these concepts are made more precise in Section 2).

Definition 1.1. *We call \mathcal{L} an L-function if*

- (i) $\mathcal{L} \geq 0$,
- (ii) $\mathcal{L}(z, \cdot)$ is convex for all z ,
- (iii) \mathcal{L} induces an evolution equation $\dot{z}_t = \mathcal{A}_{\mathcal{L}}(z_t)$, in the sense that for all (z, s) ,

$$\mathcal{L}(z, s) = 0 \quad \iff \quad s = \mathcal{A}_{\mathcal{L}}(z) \quad (1.2)$$

(writing $t \mapsto z_t$ for a time-parametrized curve in \mathcal{Z}).

The evolution induced by an L-function \mathcal{L} is a distinguished subclass of the Lagrangian flow, described by the Euler-Lagrange equation

$$D_z \mathcal{L}(z_t, \dot{z}_t) - \frac{d}{dt} (D_s \mathcal{L}(z_t, \dot{z}_t)) = 0, \quad (1.3)$$

where D_z and D_s are derivatives with respect to the first and second variable respectively. For an evolution induced by an L-function, we have that both $D_s \mathcal{L}(z_t, \dot{z}_t) = 0$ and $D_z \mathcal{L}(z_t, \dot{z}_t) = 0$, hence (1.3) is indeed satisfied. Note however, that the induced evolution $\dot{z}_t = \mathcal{A}_{\mathcal{L}}(z_t)$ is a first-order equation with respect to time, whereas (1.3) is of second order and has many more solutions.

L-functions and their induced evolutions are abundant; they not only appear in the large deviations of stochastic processes, but are also well known from Hamiltonian mechanics. In Section 1.4 we describe how such functions arise from large deviations.

Much of the theory in this paper requires only the three defining properties of L-functions. Therefore, in the general setting, we look for a connection between L-functions and gradient flows. At the other side of the connection, we will need will need a generalization of the gradient-flow concept, as we explain in the next section.

1.2 Generalized gradient systems

We first recall the Riemannian description of a gradient flow. For a brief moment, let \mathcal{Z} be a Riemannian manifold with tangent and cotangent spaces $T\mathcal{Z}$ and $T^*\mathcal{Z}$ and symmetric, positive definite metric tensor $G(z) : T_z^*\mathcal{Z} \rightarrow T_z\mathcal{Z}$. A gradient flow of an entropy functional $\mathcal{S} : \mathcal{Z} \rightarrow \mathbb{R}$ in the space \mathcal{Z} is then described by

$$G(z_t)\dot{z}_t = -D\mathcal{S}(z_t) \quad \text{or equivalently} \quad \dot{z}_t = -G(z_t)^{-1}D\mathcal{S}(z_t), \quad (1.4)$$

for all $t \in [0, T]$. This can be alternatively formulated by defining, for $z \in \mathcal{Z}, s \in T_z\mathcal{Z}$ and $\xi \in T_z^*\mathcal{Z}$,

$$\Psi(z, s) := \frac{1}{2}\langle G(z)s, s \rangle \quad \text{and its convex dual} \quad \Psi^*(z, \xi) := \frac{1}{2}\langle \xi, G(z)^{-1}\xi \rangle, \quad (1.5)$$

where $\langle \cdot, \cdot \rangle$ denotes the dual pairing between $T_z^*\mathcal{Z}$ and $T_z\mathcal{Z}$. If $G(z)$ is linear and symmetric, then (1.4) can be written as

$$D_s\Psi(z_t, \dot{z}_t) = -D\mathcal{S}(z_t) \quad \text{or equivalently} \quad \dot{z}_t = D_\xi\Psi^*(z_t, -D\mathcal{S}(z_t)), \quad (1.6)$$

which, by Fenchel-Young theory, is equivalent to requiring

$$\Psi(z_t, \dot{z}_t) + \Psi^*(z_t, -D\mathcal{S}(z_t)) + \langle D\mathcal{S}(z_t), \dot{z}_t \rangle = 0. \quad (1.7)$$

In the more general setting of this paper, we are after formulations of the form (1.7), where \mathcal{Z} can be a general linear space, but Ψ and Ψ^* need not be of the quadratic form (1.5). Such formulations are also called Ψ - Ψ^* -structures, entropy-dissipation inequalities or entropy balances; see [DMT80, DG93, CV90, LS95, MRS08, Mie11a, Mie11b]. Here, Ψ and Ψ^* need to satisfy a number of conditions in order to yield a meaningful generalized formulation; these are summarized in the following definition.

Definition 1.2. *We call the triple $(\mathcal{Z}, \Psi, \mathcal{S})$ a generalized gradient system if $\mathcal{S} : \mathcal{Z} \rightarrow \mathbb{R}, \Psi : T\mathcal{Z} \rightarrow \mathbb{R}$ and for all $z \in \mathcal{Z}$:*

- (i) $\Psi(z, \cdot)$ is convex in the second argument,
- (iia) $\min \Psi(z, \cdot) = 0$, and
- (iib) $\Psi(z, 0) = 0$.

In addition, Ψ is called symmetric if $\Psi(z, s) = \Psi(z, -s)$ for all $(z, s) \in T\mathcal{Z}$.

A curve $z : [0, T] \rightarrow \mathcal{Z}$ satisfying (1.7) is then called a flow induced by the gradient system $(\mathcal{Z}, \Psi, \mathcal{S})$.

To explain the motivation behind this definition, observe that Condition (i) guarantees the duality between Ψ and its dual Ψ^* , i.e.

$$\Psi^*(z, \xi) = \sup_{s \in T_z \mathcal{Z}} \langle \xi, s \rangle - \Psi(z, s), \quad \text{and} \quad \Psi(z, s) = \sup_{\xi \in T_z^* \mathcal{Z}} \langle \xi, s \rangle - \Psi^*(z, \xi), \quad (1.8)$$

so that formulations (1.6) and (1.7) are still equivalent.

With this duality we have the following equivalences for Conditions (iia) and (iib)

$$\min \Psi(z, \cdot) = 0 \quad \iff \quad \Psi^*(z, 0) = 0, \quad (1.9a)$$

$$\Psi(z, 0) = 0 \quad \iff \quad \min \Psi^*(z, \cdot) = 0, \quad (1.9b)$$

so that both $\Psi \geq 0$ and $\Psi^* \geq 0$. This is a central part of the definition; it implies that for curves satisfying (1.7), we have $\frac{d}{dt} \mathcal{S}(z_t) = \langle D\mathcal{S}(z_t), \dot{z}_t \rangle \leq 0$. Moreover, we see that Conditions (iia) and (iib) together imply that the dissipation $\Psi(z, s)$ per unit of time is minimal when the generalized velocity s is zero, and that stationary points of \mathcal{S} are stationary points of the evolution; this can be seen from writing the evolution as $\dot{z}_t = D_\xi \Psi^*(z_t, -D\mathcal{S}(z_t))$.

Finally, concerning the symmetry of Ψ : this property is natural in many systems, in which forward and backward velocities lead to the same dissipation (exceptions arise e.g. in plasticity [Mie03, Sect. 3]) or friction [RSV13, Eqn. (2.2)]. Note that Ψ is symmetric if and only if Ψ^* is symmetric.

The non-negativity, duality and symmetry of Ψ and Ψ^* will play a central role in this paper.

1.3 From L-functions to generalized gradient flows

The question (1.1) can now be made more concrete. For each case that we studied, we showed that there is a deeper connection between the L-function and the gradient flow: not only does the zero-level set of \mathcal{L} generate the evolution $\dot{z}_t = \mathcal{A}(z_t)$ (i.e. (1.2)), but in fact the whole function \mathcal{L} can be *identified* with the gradient-flow structure for the evolutions, in some exact or approximate sense. This is remarkable, since the function \mathcal{L} contains much more information than its zero-level set does, and this additional information does not seem relevant for the evolution $\dot{z}_t = \mathcal{A}(z_t)$. Inspired by these examples, we make Question (1.1) more precise as

For a given L-function \mathcal{L} , can we find a generalized gradient system $(\mathcal{Z}, \Psi, \mathcal{S})$ that induces the same evolution as \mathcal{L} , and somehow preserves the structure provided by \mathcal{L} ?

Observe that, by Fenchel's inequality, (1.7) is always non-negative. Therefore, (1.7) describes a principle similar to (1.2). In Section 2 we exploit this similarity by constructing, for a large class of L-functions, a generalized gradient system $(\mathcal{Z}, \Psi, \mathcal{S})$ such that

$$\mathcal{L}(z_t, \dot{z}_t) = \Psi(z_t, \dot{z}_t) + \Psi^*(z_t, -D\mathcal{S}(z_t)) + \langle D\mathcal{S}(z_t), \dot{z}_t \rangle. \quad (1.10)$$




Figure 1: integrability condition.




Figure 2: time-symmetry condition.

Once we find such a relation, we can conclude that

- the evolution induced by \mathcal{L} is the generalized gradient flow induced by the system $(\mathcal{Z}, \Psi, \mathcal{S})$;
- the cost \mathcal{L} to deviate from the optimal curve is the same as the ‘gradient-system defect’, described by the right-hand side of (1.10);
- in this sense, the structure of the generalized gradient system is consistent with \mathcal{L} —in fact it is *equivalent* to \mathcal{L} , in the sense that \mathcal{L} can be constructed out of $(\mathcal{Z}, \Psi, \mathcal{S})$, and vice versa.

In our first main result, Theorem 2.3, we will see the surprising fact that the existence of a generalized gradient system for which (1.10) holds is equivalent to the ‘integrability condition’

$$\text{there exists } \mathcal{S} : \mathcal{Z} \rightarrow \mathbb{R} \text{ such that } D_s \mathcal{L}(z, 0) = D\mathcal{S}(z) \text{ for all } z. \quad (1.11)$$

In addition, we will see that the symmetry of the resulting potentials Ψ and Ψ^* is equivalent to the ‘time-symmetry condition’

$$\mathcal{L}(z, s) - \mathcal{L}(z, -s) = 2\langle D\mathcal{S}(z), s \rangle \quad \text{for all } z, s. \quad (1.12)$$

The two conditions (1.11) and (1.12) are illustrated graphically in Figures 1 and 2. For L-functions from large deviations of empirical processes, as we explain in Section 1.4 below, we will interpret both conditions in terms of properties of the underlying stochastic process (Section 3.2).

1.4 L-functions from large deviations

In Section 3 we study a specific class of L-functions, namely those that arise from large deviations of empirical processes, as we will now explain.

Let $X_1(t), X_2(t), \dots$ be a sequence of independent Markov processes in a state space Ω , that all have the same linear generator $\mathcal{Q} : \text{Dom } \mathcal{Q} \rightarrow C_b(\Omega)$ with $\text{Dom } \mathcal{Q} \subset C_b(\Omega)$, where $C_b(\Omega)$ denotes the space of bounded, continuous functions on Ω . Then, under suitable initial conditions, the empirical process

$$\rho^{(n)} : t \mapsto \frac{1}{n} \sum_{k=1}^n \delta_{X_k(t)}, \quad t \in [0, T] \quad (1.13)$$

converges almost surely as $n \rightarrow \infty$ to the deterministic evolution $\rho : [0, T] \rightarrow \mathcal{P}(\Omega)$ satisfying [Dud89, Th. 11.4.1]

$$\frac{d}{dt} \int_{\Omega} \phi(x) \rho_t(dx) = \int_{\Omega} (\mathcal{Q}\phi)(x) \rho_t(dx) \quad \text{for all } \phi \in \text{Dom } \mathcal{Q},$$

which we abbreviate as

$$\dot{\rho}_t = \mathcal{Q}^T \rho_t. \quad (1.14)$$

In many cases, the rate of convergence of $\rho^{(n)}$ to the deterministic evolution (1.14) is characterized by a large-deviation principle of the form (see [FK06, Def. 1.1] for the rigorous definition)

$$\text{Prob}\left(\left(\rho_t^{(n)}\right)_{t=0}^T \approx \rho\right) \underset{n \rightarrow \infty}{\sim} e^{-n(I_0(\rho_0) + I_T(\rho))}, \quad I_T(\rho) := \int_0^T \mathcal{L}(\rho_t, \dot{\rho}_t) dt, \quad (1.15)$$

where I_0 characterizes the stochastic fluctuations in the initial datum $\rho_0^{(n)}$, and I_T characterizes the fluctuations in the dynamics of $t \mapsto \rho_t^{(n)}$. The rate functional $I_0 + I_T$ can be interpreted as the probabilistic cost, or the ‘unlikeliness,’ to deviate from the deterministic evolution. This interpretation implies formally that \mathcal{L} is a non-negative function that attains 0 if and only if (1.14) holds. Therefore, \mathcal{L} induces an evolution $\dot{\rho}_t = \mathcal{A}(\rho_t)$ as in (1.2), if we define \mathcal{A} as the adjoint of \mathcal{Q} , acting on probability measures. Indeed, such \mathcal{L} can be expected to be an L-function, as we discuss in more detail in Section 3.1.

L-functions that arise from large deviations of empirical processes have a very specific structure; this structure can be used to refine the theory of Section 2. In particular, in our second main result, Theorem 3.3, we will see that for such L-functions, the integrability condition (1.11) and time symmetry (1.12) are directly related to the detailed balance condition of the underlying stochastic process.

Throughout this work we avoid the term *reversible Markov processes*, which has its origin in the microscopic time reversibility of the paths. In the theory of Markov processes, it is well-known (and we use this in the proof of Theorem 3.3) that this property is equivalent to the detailed balance condition. The point is that the usage of the term *reversibility* is quite differently in thermodynamics: the acronym GENERIC for General Equations for Non-Equilibrium Reversible Irreversible Coupling is used for combining Hamiltonian systems (called reversible systems) and gradient systems (called irreversible systems), see e.g. [Mie11a, DPZ13b]. Since “reversible Markov processes” can be described in terms of dissipative gradient systems, we rather use the term “time symmetry”, because the microscopic time symmetry is reflected in the symmetry of the dissipation potentials.

1.5 Applications

After we develop the general theory in Section 2 and 3, we apply it to two central examples that arise in the large deviations of an empirical process, as explained above.

Our first example, treated in Section 4.1, concerns a Markov process on a finite state space $\Omega = \{1, \dots, J\}$, with generator matrix Q . We consider the empirical process in $\mathcal{P}(\{1, \dots, J\})$ which converges to the deterministic linear system of first-order equations,

$$\dot{\rho}_t = Q^T \rho_t, \quad \text{in } \mathcal{P}(\{1, \dots, J\}) \times [0, T]. \quad (1.16)$$

We prove in the appendix that the corresponding large deviations L-function is given by

$$\mathcal{L}(\rho, s) = \sup_{\xi \in \mathbb{R}^J} \xi \cdot s - \mathcal{H}(\rho, \xi), \quad \text{where} \quad \mathcal{H}(\rho, \xi) = \sum_{i,j=1}^J \rho_i Q_{ij} (e^{\xi_j - \xi_i} - 1).$$

From this we construct a generalized gradient system inducing (1.16) by the methods developed in this paper.

Our current research was largely motivated by such Markov evolutions, since both [Maa11, CHLZ12] and one of us [Mie13b] serendipitously discovered a (quadratic) entropy-driven gradient system for these equations, where the metric can be seen as a finite-space counterpart of the Wasserstein metric. In fact, the latter gradient structure appears as a special case of the general gradient structure of nonlinear chemical reactions studied in [Mie11b, MM13]. Interestingly, it turns out that the generalized gradient system that we derive in this paper is fundamentally different than the quadratic structure proposed in [Maa11, CHLZ12, Mie13b].

In the second application we study the empirical process of independent Brownian particles in $\Omega = \mathbb{R}^d$ with a drift given by a general force field F . The empirical process then converges almost surely to a solution of the drift-diffusion equation,

$$\dot{\rho}_t = \Delta \rho_t + \operatorname{div}(\rho_t F), \quad \text{in } \mathcal{P}(\mathbb{R}^d) \times [0, T]. \quad (1.17)$$

The corresponding large deviation L-function is then given by

$$\mathcal{L}(\rho, s) = \frac{1}{4} \|s - \Delta \rho - \operatorname{div}(\rho F)\|_{H^{-1}(\rho)}^2.$$

Whenever F is itself a gradient, we again apply the theory from this paper to derive a connection with a gradient system for (1.17), and the resulting gradient system is the well-known Wasserstein gradient flow [JKO98, AGS08]. This connection between large deviations and the Wasserstein gradient system is certainly not new; it is investigated in [ADPZ11, ADPZ12] but was probably understood in the probability community well before. We include it here to show how the theory of this paper is a generalization of this result.

1.6 Relation with earlier work

Several examples of this connection have been discussed before, such as in [ADPZ11, ADPZ12, Ren13, PRV13, DLR13]. In some of this earlier work in this topic, we studied particle systems of the form (1.13), (1.14) and (1.15), but for a *fixed time step* $t > 0$, which yields a ‘conditional’ large-deviation principle (see [PRV13, Th. 17], [Léo07, Prop. 3.2]), formally denoted as:

$$\operatorname{Prob}(\rho_t^{(n)} \approx \rho_t \mid \rho_0^{(n)} \approx \rho_0) \underset{n \rightarrow \infty}{\sim} \exp(-n \hat{I}_t(\rho_t \mid \rho_0)), \quad (1.18)$$

For the Fokker-Planck equation, we could then prove that, in the sense of a special Γ -convergence [ADPZ11, DLR13]:

$$\frac{1}{2} \hat{I}_t(\rho_t \mid \rho_0) \sim \mathcal{S}(\rho_t) - \mathcal{S}(\rho_0) + \frac{1}{2t} d^2(\rho_0, \rho_t) \quad \text{as } t \rightarrow 0, \quad (1.19)$$

where \mathcal{S} is the free energy with respect to the Lebesgue measure and d is the Wasserstein metric. In the right-hand side of (1.19) one recognizes the discrete-time approximation of the Wasserstein gradient flow of the free energy, as found in the original paper [JKO98].

Let us also mention here the relation between our time-reversibility condition (1.12) and the concept of entropy production as described in [Mae99, MRV00]. For the empirical measure (1.13) of a fixed number of particles n , they define a notion of entropy $\text{Ent}^{(n)}$ through the relation

$$\text{Ent}^{(n)}(\rho_T) - \text{Ent}^{(n)}(\rho_0) = \log \frac{d\mathbb{P}_{[0,T]}^{(n)}}{d\mathbb{P}_{[T,0]}^{(n)}}(\rho), \quad (1.20)$$

where $\mathbb{P}_{[0,T]}^{(n)}$ is the probability on paths of the empirical measure, and $\mathbb{P}_{[T,0]}^{(n)}$ is the probability on time-reversed paths. This expression presumes a property of $\mathbb{P}_{[0,T]}^{(n)}$, which is that the right-hand side can be written as a function of the initial and final points of ρ only, a property akin to a vector field being conservative.

If we scale this identity with $1/n$, and call $2\mathcal{S}$ the limit of this entropy (formally $2\mathcal{S} := \lim_{n \rightarrow \infty} \frac{1}{n} \text{Ent}^{(n)}$), then the large deviation principle (1.15) gives, again formally,

$$\begin{aligned} 2\mathcal{S}(\rho_T) - 2\mathcal{S}(\rho_0) &= \lim_{n \rightarrow \infty} \frac{1}{n} \text{Ent}^{(n)}(\rho_T) - \lim_{n \rightarrow \infty} \frac{1}{n} \text{Ent}^{(n)}(\rho_0) \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{[0,T]}^{(n)}(\rho) - \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{[T,0]}^{(n)}(\rho) = \int_0^T \mathcal{L}(\rho_t, \dot{\rho}_t) dt - \int_0^T \mathcal{L}(\rho_t, -\dot{\rho}_t) dt. \end{aligned}$$

This is exactly the time-integrated version of our time-symmetry condition (1.12), which we show to be equivalent to detailed balance (Section 3). This confirms the relationship between the assumption underlying (1.20) and detailed balance that was also discussed in [MRV00].

Finally, we like to note that in practical applications, often one does not have a truly Riemannian manifold. In the case of the Wasserstein space, the linear structure of the space of Borel measures can be used to define derivatives, the tangent space is still well-defined [AGS08], and at least formally, a generalized gradient system in the sense of Section 1.2 can be described. On general metric spaces, however, this may fail. It is then possible to define a purely metric formulation of generalized gradient systems, for which we refer to [RMS08].

2 From L-functions to generalized gradient flows

In this section we will see exactly under which conditions a generalized gradient system can be constructed out of a given L-function. To stress the generality of our work, we now switch to a more general setting, not necessarily related to large deviations. Throughout this section, \mathcal{Z} denotes a space with well-defined (linear) tangent vector bundle $T\mathcal{Z}$ and cotangent vector bundle $T^*\mathcal{Z}$ that are in dual pairing with each other via a bilinear form $\langle \cdot, \cdot \rangle$. All Fréchet derivatives are defined through this dual pairing: if $\mathcal{S} : \mathcal{Z} \rightarrow \mathbb{R}$ and $\Psi : T\mathcal{Z} \rightarrow \mathbb{R}$ then $D\mathcal{S}(z), D_s\Psi(z, s) \in T_z^*\mathcal{Z}$, and if $\Psi^* : T^*\mathcal{Z} \rightarrow \mathbb{R}$ then

$D_\xi \Psi^*(z, \xi) \in T_z \mathcal{Z}$. To focus on the main ideas, we assume that all functionals are everywhere finite and differentiable (see Remark 2.4).

2.1 The relationship between \mathcal{L} and the generalized gradient system

Assume we are given an L-function $\mathcal{L} : T\mathcal{Z} \rightarrow \mathbb{R}$. In what follows, it will be convenient to consider its convex dual \mathcal{H} , defined through the relations

$$\mathcal{H}(z, \xi) = \sup_{s \in T_z \mathcal{Z}} \langle \xi, s \rangle - \mathcal{L}(z, s), \quad \text{and} \quad \mathcal{L}(z, s) = \sup_{\xi \in T_z^* \mathcal{Z}} \langle \xi, s \rangle - \mathcal{H}(z, \xi). \quad (2.1)$$

The following lemma is central in all that follows, despite its short and simple proof. In this construction, we will allow any covector field $\mathcal{V}(z)$, not necessarily of the form $D\mathcal{S}(z)$.

Lemma 2.1. *Let \mathcal{L} and \mathcal{H} be related through (2.1), and $z \mapsto \mathcal{V}(z) \in T_z^* \mathcal{Z}$ be any covector field. Moreover, let Ψ and Ψ^* be related through (1.8) with $\Psi^*(\cdot, 0) = 0$. Then \mathcal{L} can be written as*

$$\mathcal{L}(z, s) = \Psi(z, s) + \Psi^*(z, -\mathcal{V}(z)) + \langle \mathcal{V}(z), s \rangle \quad \text{for all } (z, s) \in T\mathcal{Z}, \quad (2.2)$$

if and only if $\Psi^* = \Psi_{\mathcal{L}, \mathcal{V}}^*$ where

$$\Psi_{\mathcal{L}, \mathcal{V}}^*(z, \xi) := \mathcal{H}(z, \mathcal{V}(z) + \xi) - \mathcal{H}(z, \mathcal{V}(z)). \quad (2.3)$$

Proof. First assume that $\Psi^* = \Psi_{\mathcal{L}, \mathcal{V}}^*$. We calculate

$$\begin{aligned} \Psi_{\mathcal{L}, \mathcal{V}}(z, s) &= \sup_{\xi} \langle \xi, s \rangle - \Psi_{\mathcal{L}, \mathcal{V}}^*(z, \xi) \\ &= \sup_{\xi} \langle \xi, s \rangle - \mathcal{H}(z, \mathcal{V}(z) + \xi) + \mathcal{H}(z, \mathcal{V}(z)) \\ &= \sup_{\xi} \langle \xi - \mathcal{V}(z), s \rangle - \mathcal{H}(z, \xi) + \mathcal{H}(z, \mathcal{V}(z)) \\ &= \mathcal{L}(z, s) - \langle \mathcal{V}(z), s \rangle + \mathcal{H}(z, \mathcal{V}(z)). \end{aligned} \quad (2.4)$$

Since $\mathcal{H}(z, \mathcal{V}(z)) = -\Psi_{\mathcal{L}, \mathcal{V}}^*(z, -\mathcal{V}(z))$ by (2.3), we obtain (2.2).

For the other direction, assume that (2.2) holds. The dual of (2.2) is then given by

$$\begin{aligned} \mathcal{H}(z, \xi) &= \sup_s \langle \xi, s \rangle - \mathcal{L}(z, s) \\ &= \sup_s \langle \xi - \mathcal{V}(z), s \rangle - \Psi(z, s) - \Psi^*(z, -\mathcal{V}(z)) \\ &= \Psi^*(z, \xi - \mathcal{V}(z)) - \Psi^*(z, -\mathcal{V}(z)). \end{aligned}$$

The last term can be found by substituting $\xi = \mathcal{V}(z)$:

$$\mathcal{H}(z, \mathcal{V}(z)) = -\Psi^*(z, -\mathcal{V}(z)),$$

since $\Psi^*(z, 0) = 0$ by assumption. We conclude that Ψ^* is of the form (2.3). \square

The surprising aspect of this lemma is that *any* L-function \mathcal{L} can be written in the form (2.2), for *any* covector field \mathcal{V} . Therefore, this property contains no information in itself. Convenient choices of the covector field \mathcal{V} might be determined by the question whether (a) \mathcal{V} is a derivative of some functional \mathcal{S} , and (b) whether the corresponding Ψ and Ψ^* are both non-negative (see (1.9a) and (1.9b))—since the combination of both leads to monotonicity of \mathcal{S} . We first address property (b) in Proposition 2.2 below, and we simultaneously address the question when Ψ is symmetric.

Proposition 2.2. *Let an L-function $\mathcal{L} : T\mathcal{Z} \rightarrow \mathbb{R}$ and a covector field \mathcal{V} be given, and let $\Psi_{\mathcal{L},\mathcal{V}}^*$ and its dual $\Psi_{\mathcal{L},\mathcal{V}}$ be constructed according to Lemma 2.1. Then*

(i) $\Psi_{\mathcal{L},\mathcal{V}}, \Psi_{\mathcal{L},\mathcal{V}}^* \geq 0$ if and only if $\mathcal{V}(z) = \mathcal{V}_{\mathcal{L}}(z) := D_s\mathcal{L}(z, 0)$ for all $z \in \mathcal{Z}$. In that case, $\Psi_{\mathcal{L},\mathcal{V}}$ can be written as

$$\Psi_{\mathcal{L}}(z, s) := \mathcal{L}(z, s) - \mathcal{L}(z, 0) - \langle \mathcal{V}_{\mathcal{L}}(z), s \rangle. \quad (2.5)$$

(ii) Assume that (i) is satisfied. Then the following statements are equivalent:

■ $\Psi_{\mathcal{L}}$ is symmetric (Definition 1.2),

■ $\Psi_{\mathcal{L}}^*$ is symmetric, i.e. $\Psi^*(z, \xi) = \Psi^*(z, -\xi) \quad \forall (z, \xi) \in T^*\mathcal{Z}$,

■ $\mathcal{L}(z, s) - \mathcal{L}(z, -s) = 2\langle \mathcal{V}_{\mathcal{L}}(z), s \rangle \quad \forall (z, s) \in T\mathcal{Z}$, ‘time symmetry’ (2.6)

■ $\mathcal{H}(z, \mathcal{V}_{\mathcal{L}}(z) - \xi) = \mathcal{H}(z, \mathcal{V}_{\mathcal{L}}(z) + \xi) \quad \forall (z, \xi) \in T^*\mathcal{Z}$. (2.7)

Proof. To prove (i), we first assume that $\mathcal{V}(z) = \mathcal{V}_{\mathcal{L}}(z) = D_s\mathcal{L}(z, 0)$. By duality, $\mathcal{V}_{\mathcal{L}}(z)$ is a minimizer of $\mathcal{H}(z, \cdot)$, so that (2.3) gives

$$\Psi_{\mathcal{L},\mathcal{V}}^*(z, \xi) = \mathcal{H}(z, \mathcal{V}_{\mathcal{L}}(z) + \xi) - \mathcal{H}(z, \mathcal{V}_{\mathcal{L}}(z)) \geq 0,$$

and clearly $\Psi_{\mathcal{L},\mathcal{V}}^*(z, 0) = 0$. By (1.9a) this implies that $\Psi_{\mathcal{L},\mathcal{V}}, \Psi_{\mathcal{L},\mathcal{V}}^* \geq 0$. Conversely, if $\Psi_{\mathcal{L},\mathcal{V}} \geq 0$ and $\Psi_{\mathcal{L},\mathcal{V}}^* \geq 0$, then $D_s\Psi_{\mathcal{L},\mathcal{V}}(z, 0) = 0$. Differentiating (2.2) then gives

$$D_s\mathcal{L}(z, 0) = D_s\Psi_{\mathcal{L},\mathcal{V}}(z, 0) + \mathcal{V}(z) = \mathcal{V}(z),$$

and hence $\mathcal{V} = \mathcal{V}_{\mathcal{L}}$.

To prove (2.5), observe that $\Psi_{\mathcal{L}}(z, 0) = 0$, and (2.4) gives

$$0 = \Psi_{\mathcal{L}}(z, 0) = \mathcal{L}(z, 0) + \mathcal{H}(z, \mathcal{V}_{\mathcal{L}}(z)).$$

Identity (2.5) follows from resubstituting this expression into (2.4).

(ii) The equivalence of the symmetry of $\Psi_{\mathcal{L}}$ and $\Psi_{\mathcal{L}}^*$ follows directly from the duality relation (1.8). By identity (2.5), the symmetry property $\Psi_{\mathcal{L}}(z, s) = \Psi_{\mathcal{L}}(z, -s)$ is equivalent to (2.6). By (2.3), condition (2.7) is equivalent to symmetry of $\Psi_{\mathcal{L}}^*$. □

2.2 The first main result

By Proposition 2.2, the main conditions on \mathcal{L} to supply the induced evolution $\dot{z}_t = \mathcal{A}(z_t)$ with a gradient structure are that the covector field $\mathcal{V}(z)$ equals $D_s\mathcal{L}(z, 0)$, and that this covector field satisfies the *integrability condition* $D_s\mathcal{L}(z, 0) = DS(z)$ for some S . This leads to the following abstract result:

Theorem 2.3. *For any given L-function \mathcal{L} there holds:*

- (i) *If the covector field $D_s\mathcal{L}(z, 0) = DS_{\mathcal{L}}(z)$ for some $S_{\mathcal{L}} : \mathcal{Z} \rightarrow \mathbb{R}$, and $\Psi_{\mathcal{L}}, \Psi_{\mathcal{L}}^*$ are constructed according to Lemma 2.1 with $\mathcal{V} = DS_{\mathcal{L}}$, then $(\mathcal{Z}, \Psi_{\mathcal{L}}, S_{\mathcal{L}})$ is a generalized gradient system that induces the evolution $\dot{z}_t = \mathcal{A}_{\mathcal{L}}(z_t)$ and satisfies (2.2).*
- (ii) *Moreover, any generalized gradient system (\mathcal{Z}, Ψ, S) inducing the evolution $\dot{z}_t = \mathcal{A}_{\mathcal{L}}(z_t)$ and satisfying (2.2) equals $(\mathcal{Z}, \Psi_{\mathcal{L}}, S_{\mathcal{L}})$, up to adding a constant to S .*
- (iii) *In addition, $\Psi_{\mathcal{L}}$ is symmetric if and only if \mathcal{L} satisfies the time-symmetry condition (2.6).*

Proof. The theorem is a direct consequence of Lemma 2.1 and Proposition 2.2, using the covector field $\mathcal{V}_{\mathcal{L}} = DS_{\mathcal{L}}$. □

We return to the interpretation of this theorem in the discussion in Section 5.

Remark 2.4. As discussed in the introduction of this section, we have assumed that $\mathcal{L}(z, \cdot)$ is differentiable in 0. Many of these results generalize to the case where $\mathcal{L}(z, \cdot)$ is convex, lower semicontinuous, and may attain the value $+\infty$. Then the subdifferential $\partial_s\mathcal{L}(z, 0) := \{\xi \in T_z^*\mathcal{Z} : \forall s \in T_z\mathcal{Z} \quad \mathcal{L}(z, s) \geq \mathcal{L}(z, 0) + \langle \xi, s \rangle\}$ is well-defined, and for instance the condition $D_s\mathcal{L}(z, 0) = DS(z)$ can be generalized to finding an $S_{\mathcal{L}}$ such that $DS_{\mathcal{L}}(z) \in \partial_s\mathcal{L}(z, 0)$. In some cases the subdifferential $\partial_s\mathcal{L}(z, 0)$ may be empty. This happens for example if \mathcal{L} corresponds to the large-deviation rate functional of the empirical process on a finite state space $\{1, \dots, J\}$, where the underlying Markov processes are reducible and the tangent space includes tangents of the form $\mathbb{1}_{\omega_1} - \mathbb{1}_{\omega_2}$ for ω_1 and ω_2 from different irreducible components of $\{1, \dots, J\}$. □

3 Large-deviation L-functions of empirical processes

In this section we decrease the level of abstraction by one notch, and focus on a specific class of L-functions, namely those that arise from a large-deviation principle of an empirical process. Throughout this section, we consider a general Polish (i.e. complete separable metric) space Ω , and a general linear generator $\mathcal{Q} : \text{Dom } \mathcal{Q} \rightarrow C_b(\Omega), \text{Dom } \mathcal{Q} \subset C_b(\Omega)$. As usual, we assume that $\mathcal{Q}1 \equiv 0$, which is equivalent to requiring conservation of probability. As described in Section 1.5, we study the large deviations of the empirical process $\rho^{(n)}$, defined in (1.13), corresponding to a sequence of independent Markov processes with generator \mathcal{Q} .

We define $\mathcal{M}(\Omega)$ as the space of finite, signed Borel measures on Ω . Since $\rho^{(n)}$ is a (random) probability measure, the state space \mathcal{Z} is a subspace of $\mathcal{P}(\Omega) := \{\rho \in \mathcal{M}(\Omega) : \rho \geq 0, \rho(\Omega) = 1\}$. Note that we denote elements of \mathcal{Z} as ρ , to emphasize that we are now working with measures. The tangent bundle $T\rho\mathcal{Z}$ will formally be a subspace of $\mathcal{M}_0(\Omega) := \{s \in \mathcal{M}(\Omega) : \rho(\Omega) = 0\}$. For general purposes it suffices to take the cotangent bundle $T^*\mathcal{Z}$ as a subset of $C_b(\Omega)$, with the usual dual pairing

$$\langle \xi, s \rangle := \int_{\Omega} \xi(x) s(dx).$$

The exact spaces $\mathcal{Z} \subset \mathcal{P}(\Omega)$, $T\mathcal{Z} \subset \mathcal{Z} \times \mathcal{M}_0(\Omega)$ and $T^*\mathcal{Z} \subset \mathcal{Z} \times C_b(\Omega)$ depend on the situation, and must be chosen case-by-case.

3.1 Large-deviation principles and L-functions

Today, there are two common methods to derive large-deviation principles of the form (1.15). Both methods yield the convex dual $\mathcal{H} : T^*\mathcal{Z} \rightarrow \mathbb{R}$ of the L-function \mathcal{L} . As we have seen, this duality plays a crucial role in the connection with generalized gradient flows.

The first method, sometimes called Cramér's trick, is based on the idea that, in order to calculate the large deviation rate for any path ρ , one can perturb the generator in such a way that the curve ρ becomes the most likely one (see for example [KL99, Ch. 10]). As an illustration and for the sake of completeness, we use this method to prove a large-deviation principle for the empirical process of finite-state Markov chains in Section 4.1. Since the arguments are more-or-less standard, we postpone this proof to the appendix.

The other method, due to Feng and Kurtz [FK06], provides a formal way to calculate the large-deviation rate directly, although it requires more advanced techniques to make the arguments rigorous. It can be used to calculate a broad class of large deviations of a family of Markov processes. For empirical processes of independent particles, the formalism gives us the explicit formula (see [FK06, Example 1.14 and Section 13.3]):

$$\mathcal{H}(\rho, \xi) := \int_{\Omega} e^{-\xi} (\mathcal{Q}e^{\xi}) d\rho. \quad (3.1)$$

This function also plays an important role in Cramér's method; for example, it appears in our large-deviation proof for finite-state Markov chains in equation (A.3).

By either method one arrives at the large-deviation principle (1.15), with some function $\mathcal{L} : T\mathcal{Z} \rightarrow \mathbb{R}$, related to \mathcal{H} through

$$\mathcal{H}(\rho, \xi) = \sup_s \langle \xi, s \rangle - \mathcal{L}(\rho, s). \quad (3.2)$$

We point out that for general large deviations, \mathcal{L} may not be convex (this can occur for example when the limit process is not deterministic). In that case, \mathcal{L} will not be an L-function in the sense of Definition 1.1 and we can not connect to the theory from Section 2; in particular, the other duality relation (2.1) may not hold. We show here that these problems do not occur for empirical processes.

Proposition 3.1. *Assume the empirical process (1.13) satisfies the large-deviation principle (1.15) in the Skorohod space $\mathcal{D}([0, T]; \mathcal{P}(\Omega))$, where \mathcal{H} is given by (3.1), and \mathcal{L} is related to \mathcal{H} through (3.2). Then \mathcal{L} is an L-function, i.e.*

(i) $\mathcal{L} \geq 0$,

(ii) $\mathcal{L}(\rho, \cdot)$ is convex for all ρ , hence $\mathcal{L}(\rho, s) = \sup_{\xi} \langle \xi, s \rangle - \mathcal{H}(\rho, \xi)$, and

(iii) $\mathcal{L}(\rho, s) = 0 \iff s = \mathcal{Q}^T \rho$,

where \mathcal{Q}^T is the adjoint of the one-particle generator \mathcal{Q} , acting on probability measures.

Proof. The non-negativity of \mathcal{L} follows immediately from the definition of the large-deviation principle.

The proof of the convexity of \mathcal{L} is inspired by [DG87, Lem. 4.6]: first define the empirical measure on the curve space

$$\theta^{(n)} := \frac{1}{n} \sum_{k=1}^n \delta_{(X_k(t))_{t=0}^T},$$

which lies in $\mathcal{P}(\mathcal{D}([0, T]; \Omega))$, where $\mathcal{D}([0, T]; \Omega)$ is the Skorohod space of curves $x : [0, T] \rightarrow \Omega$. Then by Sanov's Theorem [DZ87, Th. 6.2.10], the random variable $\theta^{(n)}$ satisfies a large-deviation principle, where the rate functional is the relative entropy on the curve space. We only exploit the convexity of this rate functional. Define the transformation $F : \mathcal{P}(\mathcal{D}([0, T]; \Omega)) \rightarrow \mathcal{D}([0, T]; \mathcal{P}(\Omega))$, with

$$F(\theta)(dx) := \left(\theta(\{y \in \mathcal{D}([0, T]; \Omega) : y_t \in dx\}) \right)_{t=0}^T.$$

We can then retrieve the empirical process through $F(\theta^{(n)}) = \rho^{(n)}$, and the corresponding large-deviation rate is retrieved by a contraction principle [DZ87, Th. 4.2.1]. By linearity of the transformation F we find that the rate functional $I_0 + I_T$ corresponding to $\rho^{(n)}$ also is convex [dH00, Th. III.32]. For any $\rho \in \mathcal{P}(\Omega)$, and any two $s_1, s_2 \in T_\rho \mathcal{P}(\Omega)$, take two arbitrary curves $(\rho_t)_{t=0}^T, (\mu_t)_{t=0}^T$ in $\mathcal{D}([0, T]; \mathcal{P}(\Omega))$ such that $\rho_0 = \rho = \mu_0$ and $\dot{\rho}_0 = s_1$, and $\dot{\mu}_0 = s_2$. Applying the convexity of $I_T + I_0$ to these two curves, it then follows from the arbitrariness of T that \mathcal{L} is convex in the second argument. Naturally, this implies that \mathcal{L} is the convex dual of \mathcal{H} .

To show that $\mathcal{L}(\rho, \mathcal{Q}^T \rho) = 0$, note that because of the convexity of \mathcal{L} we can write

$$\mathcal{L}(\rho, \mathcal{Q}^T \rho) = \sup_{\xi} \langle \xi, \mathcal{Q}^T \rho \rangle - \mathcal{H}(\rho, \xi), \quad (3.3)$$

and that from (3.1), we have

$$\langle D_{\xi} \mathcal{H}(\rho, 0), \xi \rangle = - \int_{\Omega} \xi \underbrace{(\mathcal{Q}1)}_{=0} d\rho + \int_{\Omega} (\mathcal{Q}\xi) d\rho = \langle \mathcal{Q}\xi, \rho \rangle = \langle \xi, \mathcal{Q}^T \rho \rangle.$$

Since this identity is exactly the stationarity equation for the supremum in (3.3), we know that $\xi = 0$ is a critical point of $\xi \mapsto \langle \mathcal{Q}\xi, \rho \rangle - \mathcal{H}(\rho, \xi)$. In fact, $\xi = 0$ is a maximizer, because \mathcal{H} is convex by (3.2). Therefore (3.3) becomes

$$\mathcal{L}(\rho, \mathcal{Q}^T \rho) = -\mathcal{H}(\rho, 0) \stackrel{(3.1)}{=} 0.$$

Finally, for the minimizer $\mathcal{Q}^T \rho$ of $\mathcal{L}(\rho, \cdot)$ to be unique, it suffices that \mathcal{L} is strictly convex. This follows from the fact that \mathcal{L} is the dual of \mathcal{H} , and \mathcal{H} is differentiable in ξ by (3.1). \square

3.2 The second main result

From Section 2 we know that two conditions on L-functions play a crucial role in constructing a corresponding generalized gradient flow: the integrability condition (1.11) and the time-symmetry condition (2.6). We will now see that these two conditions are strongly related to the detailed balance condition of each $X_i(t)$ (and hence of the empirical process $\rho^{(n)}$). We recall the definition here:

Definition 3.2. *A Markov process with generator \mathcal{Q} satisfies detailed balance with respect to a locally finite measure π if for all $\psi, \phi \in \text{Dom } \mathcal{Q}$ such that $\mathcal{Q}(\phi\pi), \mathcal{Q}(\psi\pi) \in L^1(\pi)$,*

$$\int_{\Omega} \phi(x)(\mathcal{Q}\psi)(x) \pi(dx) = \int_{\Omega} \psi(x)(\mathcal{Q}\phi)(x) \pi(dx). \quad (3.4)$$

Naturally, this definition immediately implies that π is invariant under \mathcal{Q} , i.e.

$$\int_{\Omega} (\mathcal{Q}\psi) d\pi = 0 \quad \text{for all } \psi \in \text{Dom } \mathcal{Q} \text{ such that } \mathcal{Q}\psi \in L^1(\pi).$$

With an invariant measure π at hand, one can construct the relative entropy

$$\mathcal{E}_{\pi}(\rho) := \int_{\Omega} \log \frac{d\rho}{d\pi}(x) \rho(dx). \quad (3.5)$$

To avoid technicalities, we assume that the state space \mathcal{Z} is chosen such that \mathcal{E}_{π} is always finite. This excludes reducible processes, and we have $\pi > 0$ on \mathcal{Z} .

The relative entropy (3.5) plays an important role in the theory of large deviations of empirical measures. It appears as the large-deviation rate of the empirical measure (i.e. when no time dependence is involved) of independent, identically distributed random variables, by Sanov's Theorem. It turns out that the relative entropy also plays an important role in the connection between large deviations and generalized gradient flows. In fact, we prove below that, if detailed balance holds, then $\mathcal{E}_{\pi}/2$ is exactly the right choice for the functional driving a generalized gradient flow, in the sense that both the integrability condition and the time symmetry condition are satisfied. We also prove the reverse statement. With detailed balance, we can therefore connect to the theory set out in Section 2, and construct a generalized gradient flow with symmetric potentials Ψ, Ψ^* . Note that the factor $1/2$ compensates for the factor 2 in the time-symmetry relation (2.6).

We present the second main result in a compact version where the assumptions are stronger than needed. We see the result as a principle that holds in general contexts, but needs an extra proof for these cases. In the next section we give possible extensions.

Theorem 3.3. *Let an empirical process $\rho^{(n)}$ satisfy a large-deviation principle of the form (1.15) for all $T > 0$ with corresponding L-function \mathcal{L} . Assume that the invariant measure satisfies $\pi \in \mathcal{P}(\Omega)$ and*

that \mathcal{L} is twice Fréchet differentiable near $(\rho, s) = (\pi, 0)$. Then the time symmetry (2.6) holds with respect to $\mathcal{V}_{\mathcal{L}} = \frac{1}{2}D\mathcal{E}_{\pi}$ if and only if the generator \mathcal{Q} satisfies detailed balance with respect to π . In particular, this implies that the integrability condition (1.11) holds with $S = \frac{1}{2}\mathcal{E}_{\pi}$.

By Theorem 2.3, it follows that \mathcal{L} generates a generalized gradient system driven by the functional $\frac{1}{2}\mathcal{E}_{\pi}$.

Remark 3.4. The twice differentiability of \mathcal{L} , which is only used for the direction “ \Rightarrow ”, can be weakened by working on a suitable subspace. This requires a subspace $\mathcal{G} \subset C_b(\Omega)$ with the following properties:

- 1 $\mathcal{G}_1 := \text{span}(\mathcal{G} \cup \{1\})$ is a C_b -dense subset of $\text{Dom } \mathcal{Q}$, and $\mathcal{Q}\mathcal{G}_1 = \{\mathcal{Q}\phi : \phi \in \mathcal{G}_1\}$ is a dense subset of $L^1(\pi)$;
- 2 For every $\phi \in \mathcal{G}$ we have $\int_{\Omega} \phi d\pi = 0$ and $\mathcal{Q}^T(\phi\pi)$ is a finite measure;
- 3 For every $\phi, \psi \in \mathcal{G}$, the function

$$a(\varepsilon_1, \varepsilon_2, \varepsilon_3) := \mathcal{L}\left((1 + \varepsilon_1\phi)\pi, \mathcal{Q}^T[(\varepsilon_2\phi + \varepsilon_3\psi)\pi]\right) \quad (3.6)$$

is well-defined and twice continuously differentiable in a neighborhood of $(0, 0, 0)$.

□

Proof. We first prove the reverse statement. “ \Leftarrow ” Since $\pi \in \mathcal{P}(\Omega)$, detailed balance (3.4) is equivalent to (see for example [Lig85, Prop. II.5.3])

$$\text{Prob}\left((\rho_t^{(n)})_{t=0}^T \in A\right) = \text{Prob}\left((\rho_t^{(n)})_{t=0}^T \in \sigma_T A\right) \quad (3.7)$$

for all measurable $A \subset \mathcal{D}([0, T]; \mathcal{P}(\Omega))$, where $(\sigma_T \rho)_t := \rho_{T-t}$ denotes the time-reversal. In the probabilities above the curves $\rho^{(n)}$ are assumed to start, at $t = 0$, from empirical measures of points drawn from π . Therefore, by definition of the large-deviation principle,

$$I_0(\rho_0) + I_T(\rho) = I_0(\rho_T) + I_T(\sigma_T \rho)$$

for any curve $(\rho_t)_{t=0}^T$. Moreover, by Sanov’s Theorem $I_0 = \mathcal{E}_{\pi}$, so that

$$\mathcal{E}_{\pi}(\rho_0) + \int_0^T \mathcal{L}(\rho_t, \dot{\rho}_t) dt = \mathcal{E}_{\pi}(\rho_T) + \int_0^T \mathcal{L}(\rho_{T-t}, -\dot{\rho}_{T-t}) dt.$$

Differentiating with respect to T then yields (2.6) with $\mathcal{V}_{\mathcal{L}} = \frac{1}{2}D\mathcal{E}_{\pi}$, because $(\rho_t)_{t=0}^T$ was arbitrary.

“ \Rightarrow ”. We now assume (2.6) and prove the detailed-balance property (3.4) relying only on the weakened differentiability of Remark 3.4. By the density assumption on \mathcal{G} it is sufficient to prove (3.4) for all $\phi, \psi \in \mathcal{G}$; fix such $\phi, \psi \in \mathcal{G}$ and the corresponding function a defined in (3.6).

Using the definition of \mathcal{E}_π , we calculate that $D^2\mathcal{E}_\pi(\pi)[s_1, s_2] = \int_\Omega \frac{ds_1}{d\pi} \frac{ds_2}{d\pi} d\pi$, and therefore

$$\int_\Omega \psi(\mathcal{Q}\phi) d\pi = \int_\Omega \phi d\mathcal{Q}^T(\psi\pi) = D^2\mathcal{E}_\pi(\pi)[\phi\pi, \mathcal{Q}^T(\psi\pi)].$$

The property of detailed balance is therefore proved if we show that for all $\phi, \psi \in \mathcal{G}$

$$D^2\mathcal{E}_\pi(\pi)[\phi\pi, \mathcal{Q}^T(\psi\pi)] = D^2\mathcal{E}_\pi(\pi)[\psi\pi, \mathcal{Q}^T(\phi\pi)]. \quad (3.8)$$

To prove this, first note that since $\mathcal{L}((1 + \varepsilon_1\phi)\pi, \mathcal{Q}^T[(1 + \varepsilon_1\phi)\pi] + s)$ is minimal at $s = 0$, we have

$$D_3a(\varepsilon_1, \varepsilon_1, 0) = 0.$$

Differentiating this expression with respect to ε_1 we find that

$$D_{13}a(0, 0, 0) + D_{23}a(0, 0, 0) = 0. \quad (3.9)$$

Secondly, the time symmetry (2.6) with respect to $\frac{1}{2}D\mathcal{E}_\pi$ reads explicitly

$$\mathcal{L}(\rho, s) - \mathcal{L}(\rho, -s) = \langle D\mathcal{E}_\pi(\rho), s \rangle. \quad (3.10)$$

Using appropriate ρ and s we obtain

$$a(\varepsilon_1, 0, \varepsilon_3) - a(\varepsilon_1, 0, -\varepsilon_3) = 2\langle D\mathcal{E}_\pi((1 + \varepsilon_1\phi)\pi), \mathcal{Q}^T(\varepsilon_3\psi\pi) \rangle,$$

and by differentiating with respect to ε_3 we find

$$D_3a(\varepsilon_1, 0, 0) = \langle D\mathcal{E}_\pi((1 + \varepsilon_1\phi)\pi), \mathcal{Q}^T(\psi\pi) \rangle.$$

Differentiating again with respect to ε_1 yields

$$D_{13}a(0, 0, 0) = D^2\mathcal{E}_\pi(\pi)[\phi\pi, \mathcal{Q}^T(\psi\pi)],$$

which combines with (3.9) to give

$$D_{23}a(0, 0, 0) = -D^2\mathcal{E}_\pi(\pi)[\phi\pi, \mathcal{Q}^T(\psi\pi)].$$

This identity implies (3.8), since the left-hand side is unchanged upon exchanging ϕ and ψ , because $D_{23}a(0, 0, 0) = D_{32}a(0, 0, 0)$. This proves (3.4).

Finally, differentiating (3.10) with respect to s provides the desired integrability condition $D_s\mathcal{L}(\rho, 0) = \frac{1}{2}D\mathcal{E}_\pi(\rho)$. \square

Theorem 3.3 implies that the entropy functional (3.5) is the only functional (up to constants) for which \mathcal{L} is time-symmetric. This entropy functional enters the proof because of Sanov's theorem, and measures the cost to deviate from the equilibrium π .

Remark 3.5. Using the explicit form of \mathcal{H} from (3.1) together with Proposition 2.2(ii), we can also see more directly that detailed balance of the generator \mathcal{Q} is equivalent to the time symmetry (2.7):

$$\begin{aligned}
\mathcal{H}(\rho, \tfrac{1}{2}D\mathcal{E}_\pi(\rho) + \xi) &= \int_{\Omega} e^{-\frac{1}{2}\log\frac{d\rho}{d\pi} - \frac{1}{2} - \xi} (\mathcal{Q}e^{\frac{1}{2}\log\frac{d\rho}{d\pi} + \frac{1}{2} + \xi}) d\rho \\
&= \int_{\Omega} e^{\frac{1}{2}\log\frac{d\rho}{d\pi} - \frac{1}{2} - \xi} (\mathcal{Q}e^{\frac{1}{2}\log\frac{d\rho}{d\pi} + \frac{1}{2} + \xi}) d\pi \\
&\stackrel{(3.4)}{=} \int_{\Omega} e^{\frac{1}{2}\log\frac{d\rho}{d\pi} + \frac{1}{2} + \xi} (\mathcal{Q}e^{\frac{1}{2}\log\frac{d\rho}{d\pi} - \frac{1}{2} - \xi}) d\pi \\
&= \int_{\Omega} e^{-\frac{1}{2}\log\frac{d\rho}{d\pi} - \frac{1}{2} + \xi} (\mathcal{Q}e^{\frac{1}{2}\log\frac{d\rho}{d\pi} + \frac{1}{2} - \xi}) d\rho = \mathcal{H}(\rho, \tfrac{1}{2}D\mathcal{E}_\pi(\rho) - \xi),
\end{aligned}$$

where we have used the linearity of \mathcal{Q} to take the constant e^{-1} out of the parenthesis. \square

3.3 Generalizations

Locally finite invariant measures

Observe that in (3.4) and (3.5), the invariant measure π does not need to be a probability measure. This relaxation is not uncommon in probability theory (see for example [Lig85, Sect. II.5]), and can be quite useful when dealing with generators on unbounded domains for which an invariant probability measure does not exist, such as the diffusion equation on \mathbb{R}^d . It also implies that the value $\pi(\Omega)$ is not relevant. Indeed, if it is finite, this constant will vanish when one considers entropy differences $\mathcal{E}_\pi(\rho_2) - \mathcal{E}_\pi(\rho_1)$ or derivatives $D\mathcal{E}_\pi(\rho)$.

We now show that Theorem 3.3 can be generalized to locally finite invariant measures, under the condition that the rate functional has compact level sets. This assumption is very common in the theory of large deviations, where it is often called ‘goodness’ of the rate functional.

Proposition 3.6. *Theorem 3.3 remains true if the invariant measure π is only locally finite, and the rate functional $\rho \mapsto I_0(\rho_0) + \int_0^T \mathcal{L}(\rho_t, \dot{\rho}_t) dt$ has compact level sets in the topology of the large-deviation principle.*

Proof. The forward proof of Theorem 3.3 remains valid; we only prove the reverse statement here.

By [FK06, Lem. 3.5] the compact level sets of the rate functional imply that the sequence $\rho^{(n)}$ is exponentially tight, which allows the use of a time-discrete approximation. In the time-discrete setting, detailed balance with respect to a locally finite invariant measure implies the following discrete time-symmetry property [Ren13, Prop. 8.2.2]:

$$\hat{I}_t(\rho_t|\rho_0) - \hat{I}_t(\rho_0|\rho_t) = \mathcal{E}_\pi(\rho_t) - \mathcal{E}_\pi(\rho_0), \quad (3.11)$$

where \hat{I}_t is the conditional large-deviation rate (1.18).

First observe that, due to the contraction principle [DZ87, Th. 4.2.1], the sequence $(\rho_{t_k}^{(n)})_{k=0}^K$ satisfies a large-deviation principle in $\mathcal{P}(\Omega)^K$ for any subset $\{t_0, \dots, t_K\}$ of $[0, T]$. Moreover, because of the Markov property, there holds

$$\begin{aligned} \text{Prob}(\rho_{t_K}^{(n)} \in A_K, \rho_{t_{K-1}}^{(n)} \in A_{K-1}, \dots, \rho_{t_1}^{(n)} \in A_1 | \rho_{t_0}^{(n)} = \rho_{t_0}) \\ = \text{Prob}(\rho_{t_K}^{(n)} \in A_K | \rho_{t_{K-1}}^{(n)} \in A_{K-1}) \dots \text{Prob}(\rho_{t_1}^{(n)} \in A_1 | \rho_{t_0}^{(n)} = \rho_{t_0}). \end{aligned}$$

By definition of the large-deviation principle, the rate functional corresponding to the sequence $\rho_{t_k}^{(n)}$ can therefore be written as $\sum_{k=1}^K \hat{I}_{t_k - t_{k-1}}(\rho_{t_k} | \rho_{t_{k-1}})$.

It then follows from (3.11), the exponential tightness, and [FK06, Th. 4.28] that the pathwise large-deviation rate can be written as

$$\begin{aligned} I_T(\rho) &= \sup_{0=t_0 < t_1 < \dots < t_K=T} \sum_{k=1}^K \hat{I}_{t_k - t_{k-1}}(\rho_{t_k} | \rho_{t_{k-1}}) \\ &= \sup_{0=t_0 < t_1 < \dots < t_K=T} \sum_{k=1}^K \hat{I}_{t_k - t_{k-1}}(\rho_{t_{k-1}} | \rho_{t_k}) + \mathcal{E}_\pi(\rho_{t_k}) - \mathcal{E}_\pi(\rho_{t_{k-1}}) \\ &= \mathcal{E}_\pi(\rho_T) - \mathcal{E}_\pi(\rho_0) + \sup_{0=t_0 < t_1 < \dots < t_K=T} \sum_{k=1}^K \hat{I}_{t_k - t_{k-1}}(\rho_{t_{k-1}} | \rho_{t_k}) \\ &= \mathcal{E}_\pi(\rho_T) - \mathcal{E}_\pi(\rho_0) + I_T(\sigma_T \rho), \end{aligned}$$

where σ_T is again the time-reversal operator. Differentiation with respect to T then gives (2.6). \square

General large-deviation principles

A large part of the theory developed above remains true if we consider more general sequences of random variables, that is, not necessarily empirical measures. Let $Z_t^{(n)}$ be a sequence of stochastic processes with generator $\mathcal{Q}^{(n)}$, in a state space \mathcal{Z} , with a formal differentiable structure as in Section 2. (Note that in case of the empirical process as before, the $\mathcal{Q}^{(n)}$ plays the role of the generator of the entire empirical process). Assume that

- 1 $\mathcal{Q}^{(n)}$ satisfies detailed balance with respect to an invariant measure $\pi^{(n)} \in \mathcal{P}(\mathcal{Z})$,
- 2 $Z^{(n)}$ converges narrowly (pointwise in t) to a deterministic curve, satisfying the limiting equation

$$\dot{z}_t = \mathcal{A}(z_t),$$

- 3 For all $T > 0$, $Z^{(n)}$ satisfies a large-deviation principle of the form

$$\text{Prob}\left(\left(Z_t^{(n)}\right)_{t=0}^T \approx z\right) \underset{n \rightarrow \infty}{\sim} e^{-n(I_0(z_0) + I_T(z))}, \quad I_T(z) := \int_0^T \mathcal{L}(z_t, \dot{z}_t) dt, \quad (3.12)$$

where I_0 is the large-deviation rate functional of $Z_0^{(n)}$.

- 4 The sequence of random variables in \mathcal{Z} with law $\pi^{(n)}$ satisfies a large-deviation principle with rate \mathcal{E} .

We then have the following generalization of Theorem 3.3:

Proposition 3.7. *Under the assumptions 1–4 above, the time symmetry (2.6) holds with respect to $\mathcal{V}_{\mathcal{L}} = \frac{1}{2}D\mathcal{E}$, and the integrability condition (1.11) holds with $\mathcal{S} = \frac{1}{2}\mathcal{E}$.*

As before, a consequence is that \mathcal{L} generates a generalized gradient system driven by the functional $\frac{1}{2}\mathcal{E}$ (if \mathcal{L} is an L-function; see Remark 3.8 below).

Proof. The proof is analogous to the proof of “ \Leftarrow ” in Theorem 3.3, where we replace (3.7) by

$$\text{Prob}\left(\left(Z_t^{(n)}\right)_{t=0}^T \in A\right) = \text{Prob}\left(\left(Z_t^{(n)}\right)_{t=0}^T \in \sigma_T A\right),$$

and \mathcal{E}_π by \mathcal{E} . □

Remark 3.8. In most applications, the \mathcal{L} from (3.12) will again be an L-function, as can be seen as follows.

- (i) The non-negativity of \mathcal{L} follows from the definition of the large-deviation principle.
- (ii) Since rate functions can be assumed to be lower semicontinuous, in many cases this forces the dependence of \mathcal{L} on \dot{z} to be convex (see e.g. [Giu03, Ch. 4], [ABM06, Th. 13.1.3] and [Dac00, Th. 3.15]).
- (iii) By the Portmanteau Theorem, the pointwise narrow convergence of $Z_t^{(n)}$ is equivalent to the statement that $\liminf_{n \rightarrow \infty} \text{Prob}\left(Z_t^{(n)} \in U\right) \geq 1$ for all narrowly open $U \subset \mathcal{Z}$ containing the deterministic solution at time t . Then for any $\epsilon > 0$ there is a $N \geq 1$ such that for all $n \geq N$

$$0 \leq \frac{1}{n} \log \text{Prob}\left(Z_t^{(n)} \in U\right) \leq \frac{1}{n} \log(1 - \epsilon) \rightarrow 0.$$

Hence for any fixed t , the conditional rate functional \hat{I}_t satisfies $\hat{I}_t(z_t|z_0) = 0$ whenever z_t solves $\dot{z}_\tau = \mathcal{A}(z_\tau)$ up to time $\tau = t$. Therefore, by a similar time-discretization argument as in the proof of Proposition 3.6, this gives $\mathcal{L}(z_t, \mathcal{A}(z_t)) = 0$.

In general, the uniqueness of the minimizer $\mathcal{A}(z)$ of $\mathcal{L}(z, \cdot)$ does not follow from the large-deviation principle (3.12). In practice, this is often related to finding the right order of n for which the large-deviation \mathcal{L} does have a unique minimizer. □

4 Applications

One of the questions that led to this study concerns the connection between large deviations of finite-state-space Markov processes and the entropy-driven Riemannian gradient structure uncovered in [Maa11, CHLZ12, Mie13b]. As we described in the introduction, this structure is a generalization of the Wasserstein gradient-flow structure in continuous spaces, and therefore the question naturally arises whether this structure connects to the large deviations of the empirical process of i.i.d. copies of the underlying Markov process, following the line of Section 2. Below we calculate these large deviations, and interestingly these do *not* generate the above-mentioned Riemannian gradient structure, although the driving functional is a multiple of the relative entropy in both cases. However, by the arguments in Sections 2 and 3 they do generate a well-defined and symmetric generalized gradient structure, as we show in the first application.

In the second application we consider the diffusion-drift equation. The connection between the large deviations of the empirical measure and the Wasserstein gradient flow of the free energy was already made in [ADPZ11, DLR13] for the discrete-time setting, and in [ADPZ12, Sect. 4.2] in the (continuous-time) L-function setting. Nevertheless, we discuss below how the theory developed in Sections 2 and 3 can be used to reproduce this result.

4.1 Finite-state Markov chains

We now apply the theory of Sections 2 and 3 to continuous-time Markov chains on a finite state space $\Omega = \{1, \dots, J\}$. In this case we can identify

$$\mathcal{Z} := \mathcal{P}(\{1, \dots, J\}) = \left\{ \rho \in \mathbb{R}_+^J : \sum_{j=1}^J \rho_j = 1 \right\},$$

and the generator $\mathcal{Q} : \mathbb{R}^J \rightarrow \mathbb{R}^J$ can be identified with a Markov intensity matrix $Q \in \mathbb{R}^{J \times J}$, i.e.

$$Q_{ij} \geq 0 \quad \text{for } i \neq j, \quad \text{and} \quad \sum_{j=1}^J Q_{ij} = 0. \quad (4.1)$$

The natural tangent and cotangent spaces to $\mathcal{P}(\{1, \dots, J\})$ are

$$T_\rho \mathcal{Z} = T_\rho^* \mathcal{Z} := \left\{ s \in \mathbb{R}^J : \sum_{i=1}^J s_i = 0 \right\},$$

with the usual Euclidean inner product $\langle s, \xi \rangle := s \cdot \xi$.

Particle system and large deviations. We now take a sequence of independent Markov processes $X_1(t), X_2(t), \dots$ in $\{1, \dots, J\}$ with generator matrix Q . By identifying probability measures on the

finite state space $\{1, \dots, J\}$ with vectors in \mathbb{R}^J , the empirical process (1.13) is replaced by

$$t \mapsto \rho_t^{(n)} := \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{X_k(t)},$$

which is a random variable in $\mathcal{P}(\{1, \dots, J\})$ that counts the weighted number of particles at each site. Under suitable initial conditions, this process converges almost surely to the solution of (1.16). We prove in Theorem A.1 in the Appendix that a corresponding large-deviation principle (1.15) holds in $BV(0, T; \mathcal{P}(\{1, \dots, J\}))$ with the weak-* topology, with

$$\mathcal{L}(\rho, s) = \sup_{\xi \in T_\rho^* \mathcal{Z}} \xi \cdot s - \mathcal{H}(\rho, \xi), \quad \text{and} \quad \mathcal{H}(\rho, \xi) = \sum_{i,j=1}^J \rho_i Q_{ij} (e^{\xi_j - \xi_i} - 1). \quad (4.2)$$

Towards generalized gradient flows. Let $\pi \in \mathcal{P}(\Omega)$ be invariant, i.e. $Q^T \pi = 0$. We implicitly assume that π has only non-zero coordinates if $\pi_i = 0$ for some i , then the space \mathcal{Z} needs to be restricted to those ρ for which $\rho_i = 0$. The relative entropy functional (3.5) then becomes

$$\mathcal{E}_\pi(\rho) = \sum_{i=1}^J \rho_i \log \frac{\rho_i}{\pi_i}.$$

For this application, the role of the detailed-balance condition will appear explicitly in the calculations; let us not assume this condition for now. Following the construction from Lemma 2.1, we define (taking $\mathcal{S} = \frac{1}{2} \mathcal{E}_\pi$)

$$\begin{aligned} \Psi_{\mathcal{L}}^*(\rho, \xi) &:= \mathcal{H}\left(\rho, \frac{1}{2} D\mathcal{E}_\pi(\rho) + \xi\right) - \mathcal{H}\left(\rho, \frac{1}{2} D\mathcal{E}_\pi(\rho)\right) \\ &= \sum_{i,j=1}^J \rho_i Q_{ij} \left[\exp\left(\frac{1}{2} \log \frac{\rho_j}{\pi_j} + \xi_j - \frac{1}{2} \log \frac{\rho_i}{\pi_i} - \xi_i\right) - \exp\left(\frac{1}{2} \log \frac{\rho_j}{\pi_j} - \frac{1}{2} \log \frac{\rho_i}{\pi_i}\right) \right] \\ &= \sum_{i,j=1}^J \sqrt{\rho_i \rho_j \frac{\pi_i}{\pi_j}} Q_{ij} (e^{\xi_j - \xi_i} - 1). \end{aligned} \quad (4.3)$$

Under a condition of *weak reversibility* (cf. [Mie11b]),

$$Q_{ij} > 0 \iff Q_{ji} > 0 \quad \text{for all } i, j = 1, \dots, J, \quad (4.4)$$

expression (4.3) can be written as:

$$\begin{aligned} \Psi_{\mathcal{L}}^*(\rho, \xi) &= \frac{1}{2} \sum_{i,j=1}^J \left[\sqrt{\rho_i \rho_j \frac{\pi_i}{\pi_j}} Q_{ij} (e^{\xi_j - \xi_i} - 1) + \sqrt{\rho_i \rho_j \frac{\pi_j}{\pi_i}} Q_{ji} (e^{\xi_i - \xi_j} - 1) \right] \\ &= \frac{1}{2} \sum_{i,j=1}^J \sqrt{\rho_i \rho_j Q_{ij} Q_{ji}} \left(\cosh\left(\xi_j - \xi_i + \frac{1}{2} \log \underbrace{\frac{Q_{ij} \pi_i}{Q_{ji} \pi_j}}_{=0 \text{ if d.b. holds}}\right) - 1 \right). \end{aligned} \quad (4.5)$$

From this form it is immediate that $\inf_{\xi} \Psi_{\mathcal{L}}^*(z, \xi) = 0 = \Psi_{\mathcal{L}}^*(z, 0)$ whenever the detailed balance condition (3.4) holds. By (1.9a) and (1.9b), this implies that $\Psi, \Psi^* \geq 0$. Therefore, under the condition of detailed balance, the linear system (1.16) is described by the symmetric generalized gradient system $(\mathcal{P}(\{1, \dots, J\}), \Psi_{\mathcal{L}}, \frac{1}{2}\mathcal{E}_{\pi})$, as prescribed by Theorem 3.3.

Remark 4.1. In general, the explicit formula for $\Psi_{\mathcal{L}}$ can become much more involved than the formula for $\Psi_{\mathcal{L}}^*$. For example, $\Psi_{\mathcal{L}}^*$ is linear in Q , while $\Psi_{\mathcal{L}}$ is not. \square

To illustrate the principle of generating a generalized gradient structure out of the large deviations, we describe two examples; one for which detailed balance holds, and one for which it does not.

Example 4.2 (finite-state Markov chain with detailed balance). For this example we take $J = 2$; then detailed balance (3.4) always holds whenever (4.4) holds. We now write $s = (s_1, -s_1)$ and $\xi = (\xi_1, -\xi_1)$. Then (4.5) can be written as

$$\Psi_{\mathcal{L}}^*(\rho, \xi) = \sqrt{\rho_1 \rho_2 Q_{12} Q_{21}} (\cosh(2\xi_1) - 1),$$

and its dual becomes

$$\Psi_{\mathcal{L}}(\rho, s) = \sqrt{\rho_1 \rho_2 Q_{12} Q_{21}} \left(\cosh^* \left(\frac{s_1}{2\sqrt{\rho_1 \rho_2 Q_{12} Q_{21}}} \right) + 1 \right),$$

where \cosh^* is the Legendre transform of the hyperbolic cosine. By Theorems 2.3 and 3.3, we know that the entropy functional $\mathcal{S}_{\mathcal{L}}$ has to be the relative entropy (3.5), and we find

$$\mathcal{L}(\rho, s) = \Psi_{\mathcal{L}}(\rho, s) + \Psi_{\mathcal{L}}^*(\rho, -\frac{1}{2}D\mathcal{E}_{\pi}(\rho)) + \frac{1}{2}D\mathcal{E}_{\pi}(\rho) \cdot s,$$

where the left-hand side is the large-deviation function, and the right-hand side describes a generalized gradient system. Note especially that $\Psi_{\mathcal{L}}$ and $\Psi_{\mathcal{L}}^*$ are indeed symmetric, but not quadratic in s and ξ .

Example 4.3 (finite-state Markov chain without detailed balance). We now take $J = 3$ and

$$Q := \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ 1 & 0 & -1 \end{pmatrix}.$$

The invariant measure is $\pi = (1/3, 1/3, 1/3)$, but detailed balance does not hold. Therefore, we can no longer apply Theorem 3.3, so that we can not tell a priori whether the integrability condition (1.11) holds.

Therefore, we need to go one step back to Proposition 2.2. We have the explicit formula $\mathcal{H}(\rho, \xi) \stackrel{(4.2)}{=} \rho_1 e^{\xi_2 - \xi_1} + \rho_2 e^{\xi_3 - \xi_2} + \rho_3 e^{\xi_1 - \xi_3} - 1$ and find

$$\mathcal{V}_{\mathcal{L}}(\rho) := D_s \mathcal{L}(\rho, 0) = \operatorname{argmin} \mathcal{H}(\rho, \cdot) = \frac{1}{3} \begin{pmatrix} \log(\rho_1/\rho_3) \\ \log(\rho_2/\rho_1) \\ \log(\rho_3/\rho_2) \end{pmatrix}.$$

By Proposition 2.2(i), we can still use this covector field $\mathcal{V}_{\mathcal{L}}$ to construct $\Psi_{\mathcal{L}}, \Psi_{\mathcal{L}}^*$ such that:

$$\mathcal{L}(\rho, s) = \Psi_{\mathcal{L}}(\rho, s) + \Psi_{\mathcal{L}}^*(\rho, -\mathcal{V}_{\mathcal{L}}(\rho)) + \mathcal{V}_{\mathcal{L}}(\rho) \cdot s.$$

The right-hand side of this equation however, does not represent a generalized gradient system, since the vector field $\mathcal{V}_{\mathcal{L}}(\rho)$ is non-conservative, i.e. the integrability condition (1.11) breaks down.

4.2 Diffusion with drift

In this application, we study the drift-diffusion equation (1.17), and derive the Wasserstein gradient flow from [JKO98, Ott01]. The arguments that follow below motivate the choice of the spaces:

$$\mathcal{Z} := \left\{ \rho \in W^{1,1}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \rho = 1 \text{ and } \int_{\mathbb{R}^d} \left(x^2 \rho + \frac{|\nabla \rho|^2}{\rho} + |F|^2 \rho \right) < \infty \right\},$$

with formal cotangent and tangent bundles (see [AGS08, Sec. 12.4] for the rigorous notions):

$$T_\rho^* \mathcal{Z} := \left\{ \xi \in H^1(\rho) : \int_{\mathbb{R}^d} \xi(x) dx = 0 \right\}, \text{ and}$$

$$T_\rho \mathcal{Z} := H^{-1}(\rho) = \left\{ s \in \mathcal{D}'(\mathbb{R}^d) : \exists v \in \overline{\{\nabla \phi : \phi \in C_c^\infty(\mathbb{R}^d)\}}^{L^2(\rho)} \text{ such that } s + \operatorname{div}(\rho v) = 0 \right\},$$

and with dual pairing

$$\langle \xi, s \rangle := \int_{\mathbb{R}^d} \rho v \cdot \nabla \xi \quad \text{if } s + \operatorname{div}(\rho v) = 0,$$

where each occurrence of $s + \operatorname{div}(\rho v) = 0$ is to be interpreted in distributional sense. Naturally, for the generator we take $\mathcal{Q} : \operatorname{Dom} \mathcal{Q} \rightarrow C_b(\mathbb{R}^d)$, with $\operatorname{Dom} \mathcal{Q} \supset C_b^2(\mathbb{R}^d)$ and

$$(\mathcal{Q}\phi)(x) := \Delta \phi(x) - F(x) \cdot \nabla \phi(x). \quad (4.6)$$

Wherever needed, we tacitly assume sufficient regularity on covector functions $\xi \in H^1(\rho)$ such that all derivatives exist in a weak sense.

Particle system and large deviations. We now take independent processes $X_1(t), X_2(t), \dots$ in \mathbb{R}^d , each with generator \mathcal{Q} as described above. The convergence of the empirical process (1.13) and its corresponding large deviations are well-known results (see for example [DG87]). By formula (3.1),

$$\mathcal{H}(\rho, \xi) := \int_{\mathbb{R}^d} e^{-\xi} (\mathcal{Q} e^\xi) \rho = \int_{\mathbb{R}^d} (\Delta \xi + |\nabla \xi|^2 - F \cdot \nabla \xi) \rho. \quad (4.7)$$

The corresponding large-deviation L-function is found through

$$\begin{aligned} \mathcal{L}(\rho, s) &= \sup_{\xi \in T_\rho^* \mathcal{Z}} \langle \xi, s \rangle - \mathcal{H}(\rho, \xi) \\ &= \sup_{\xi \in T_\rho^* \mathcal{Z}} \langle \xi, s - \Delta \rho - \operatorname{div}(\rho F) \rangle - \int_{\mathbb{R}^d} |\nabla \xi|^2 \rho \\ &=: \frac{1}{4} \|s - \Delta \rho - \operatorname{div}(\rho F)\|_{H^{-1}(\rho)}^2. \end{aligned} \quad (4.8)$$

Indeed by [DG87, Th 4.5], the empirical process $\rho^{(n)}$ satisfies the large-deviation principle (1.15) in the space of continuous functions with values in $\mathcal{P}(\mathbb{R}^d)$, equipped with the narrow topology, with L-function (4.8).

Towards gradient flows. Again, following the construction of Proposition 2.2, we need to find a covector $\mathcal{V}_{\mathcal{L}}(\rho) := D_s \mathcal{L}(\rho, 0) = \operatorname{argmin} \mathcal{H}(\rho, \cdot)$. To this aim, we fix a $\rho \in \mathcal{Z}$ and set, for all $\xi \in T_{\rho}^* \mathcal{Z}$, for now neglecting the dependence of $\mathcal{V}_{\mathcal{L}}$ on ρ ,

$$0 = \langle D_{\xi} \mathcal{H}(\rho, \mathcal{V}_{\mathcal{L}}), \xi \rangle = \int_{\mathbb{R}^d} (\Delta \xi + 2 \nabla \xi \cdot \nabla \mathcal{V}_{\mathcal{L}} - F \nabla \xi) \rho,$$

which implies

$$\int_{\mathbb{R}^d} \nabla \xi \cdot \nabla \mathcal{V}_{\mathcal{L}} \rho = \frac{1}{2} \int_{\mathbb{R}^d} \nabla \xi \cdot \left(\frac{\nabla \rho}{\rho} + F \right) \rho. \quad (4.9)$$

The left-hand side is an inner product on $H^1(\rho)$; the right-hand side is a linear functional in $\xi \in H^1(\rho)$, bounded by definition of the space \mathcal{Z} . Therefore, Riesz's Theorem gives us the unique existence of $\mathcal{V}_{\mathcal{L}}(\rho) \in T_{\rho}^* \mathcal{Z}$.

With this covector field at hand, we use (2.3) to construct

$$\begin{aligned} \Psi_{\mathcal{L}}^*(\rho, \xi) &:= \mathcal{H}(\rho, \mathcal{V}_{\mathcal{L}}(\rho) + \xi) - \mathcal{H}(\rho, \mathcal{V}_{\mathcal{L}}(\rho)) \\ &\stackrel{(4.7)}{=} \int_{\mathbb{R}^d} (\Delta \xi + |\nabla \xi|^2 + 2 \nabla \xi \cdot \nabla \mathcal{V}_{\mathcal{L}}(\rho) - F \cdot \nabla \xi) \rho \\ &\stackrel{(4.9)}{=} \int_{\mathbb{R}^d} |\nabla \xi|^2 \rho, \end{aligned}$$

and

$$\Psi_{\mathcal{L}}(\rho, s) := \sup_{\xi \in T_{\rho}^* \mathcal{Z}} \langle \xi, s \rangle - \int_{\mathbb{R}^d} |\nabla \xi|^2 \rho = \frac{1}{4} \|s\|_{H^{-1}(\rho)}^2.$$

Up to now, we have not assumed anything about the force field F . We see that even in the most general case, the construction yields dissipations potentials $\Psi_{\mathcal{L}}, \Psi_{\mathcal{L}}^*$ that are symmetric (and even quadratic), and independent of F . In general however, the integrability condition $\mathcal{V}_{\mathcal{L}}(\rho) = D\mathcal{S}(\rho)$ is likely to fail, and the best we can do is to split \mathcal{L} into

$$\mathcal{L}(\rho, s) = \frac{1}{4} \|s\|_{H^{-1}(\rho)}^2 + \|\mathcal{V}_{\mathcal{L}}(\rho)\|_{H^1(\rho)}^2 + \langle \mathcal{V}_{\mathcal{L}}(\rho), s \rangle,$$

where the left-hand side is the large-deviation L-function (4.8), but the right-hand side can not be interpreted as a generalized gradient flow.

By contrast, if we assume that $F = \nabla P$ for some potential P , then $\pi(dx) = e^{-P(x)} dx$ is an (locally finite) invariant measure for (4.6), and detailed balance (3.4) is always satisfied for this π . Therefore, by Theorem 3.3, the time-symmetry condition (2.6) and the integrability condition $\mathcal{V}_{\mathcal{L}}(\rho) = D\mathcal{S}_{\mathcal{L}}(\rho)$ both hold, with entropy functional

$$\mathcal{S}_{\mathcal{L}}(\rho) := \frac{1}{2} \mathcal{E}_{\pi}(\rho) = \frac{1}{2} \int_{\mathbb{R}^d} \rho(x) \log \frac{\rho(x)}{\pi(x)} dx = \frac{1}{2} \int_{\mathbb{R}^d} \rho(x) \log \rho(x) dx + \frac{1}{2} \int_{\mathbb{R}^d} P(x) \rho(x) dx.$$

Note that the HWI inequality [Vil09, Cor. 20.13] assures that $\mathcal{S}_{\mathcal{L}}$ is indeed finite on \mathcal{Z} . Putting the parts together, we find:

$$\mathcal{L}(\rho, s) = \frac{1}{4} \|s\|_{H^{-1}(\rho)}^2 + \|D\mathcal{S}_{\mathcal{L}}(\rho)\|_{H^1(\rho)}^2 + \langle D\mathcal{S}_{\mathcal{L}}(\rho), s \rangle,$$

where the right-hand side is the well-known quadratic Wasserstein-entropy gradient system.

5 Conclusion and discussion

In the introduction we asked the question “Does there exist a general connection between large-deviation rate functionals and generalized gradient flows?” The results of this paper give an affirmative answer to this question, that we split into steps.

Step 1: Careful definition. What the ‘connection’ above means can be interpreted in multiple ways. The most restricted interpretation would be *find Ψ , Ψ^* , and \mathcal{S} such that*

$$\mathcal{L}(z, \dot{z}) = 0 \iff \Psi(z, \dot{z}) + \Psi^*(z, -D\mathcal{S}(z)) + \langle D\mathcal{S}(z), \dot{z} \rangle = 0. \quad (5.1)$$

This interpretation leaves much room: given *any* function \mathcal{S} , a function Ψ can be chosen such that $D\Psi^*(z, -D\mathcal{S}(z)) = \mathcal{A}_{\mathcal{L}}(z)$, implying (5.1). In a Hilbert space an example would be $\Psi^*(z, \xi) := \frac{1}{2}\|\xi + D\mathcal{S}(z) - \mathcal{A}(z)\|^2$.

Instead we require more, namely

$$\mathcal{L}(z, \dot{z}) = \Psi(z, \dot{z}) + \Psi^*(z, -D\mathcal{S}(z)) + \langle D\mathcal{S}(z), \dot{z} \rangle. \quad (5.2)$$

This property formulates a relationship between not only the zero values of the left- and right-hand sides, but all values. We comment on the interpretation of the other values below.

However, this requirement is still too weak to be meaningful. Lemma 2.1 shows the slightly surprising fact that it is still possible, for *any* choice of \mathcal{S} , to construct a pair Ψ, Ψ^* such that (5.2) is satisfied. Therefore even (5.2) does not carry enough information.

To improve the situation we require a further property, which is that Ψ and Ψ^* are minimized at zero, which is equivalent to Ψ and Ψ^* being non-negative. This additional condition is inspired by the fact that it makes \mathcal{S} into a Lyapunov function and that stationary points of \mathcal{S} are also stationary points of the evolution. It is with this additional condition that we call $(\mathcal{Z}, \mathcal{S}, \Psi)$ a generalized gradient system, and this is the connection that we seek.

Step 2: Abstract characterization. Now that we have defined what we are looking for, Theorem 2.3 gives a clear characterization of the possibility of finding it. Given an L-function \mathcal{L} , a corresponding generalized gradient system exists *if and only if*

$$D_s \mathcal{L}(z, 0) \text{ is a derivative (of a functional } \bar{\mathcal{S}}, \text{ say)}. \quad (5.3)$$

In that case there is *exactly one* system $(\mathcal{Z}, \mathcal{S}, \Psi)$ satisfying (5.2) (up to addition of constants in \mathcal{S}), and $\mathcal{S} = \bar{\mathcal{S}} + \text{constant}$.

At this level it is also possible to characterize when Ψ and Ψ^* are symmetric: if \mathcal{L} satisfies the relation $\mathcal{L}(z, s) - \mathcal{L}(z, -s) = 2\langle D\bar{\mathcal{S}}(z), s \rangle$ for some functional $\bar{\mathcal{S}}$, then (5.3) is automatically fulfilled, and in addition Ψ and Ψ^* are symmetric.

Step 3: Application to empirical processes. We next specialize to the case when \mathcal{L} is the large-deviation rate functional of an empirical measure of i.i.d. Markov processes. Theorem 3.3 states that the existence of a generalized gradient system *is equivalent to the property of detailed balance* of the underlying

processes. In this case the unique functional \mathcal{S} is one-half of the relative entropy with respect to the invariant measure. This property is well known in the probability community, in an informal way, but we are not aware of a proof at this level of abstraction. In Section 4 we illustrate these abstract concepts on two examples.

Summarizing, with these results we identify a class of L-functions that have corresponding generalized gradient systems, which coincides with the class of detailed-balance systems in the case of empirical measures of i.i.d. Markov processes. Let us continue with a few remarks and questions.

Interpretation of non-zero values of \mathcal{L} . Although the deterministic evolution described by $\mathcal{L}(z, \dot{z}) = 0$ depends only on the level set at value zero of \mathcal{L} , other values of \mathcal{L} are relevant when \mathcal{L} is the rate function of a large-deviation principle: $\mathcal{L}(z, \dot{z})$ can be interpreted as a ‘cost’ to deviate from the optimal value $\dot{z} = \mathcal{A}(z)$. Relation (2.2) shows that this cost is simultaneously characterized in terms of the components of the generalized gradient flow.

If \mathcal{L} is interpreted as a cost function, then (5.2) shows that this cost consists of two parts:

- the cost not to move at all rather than in the optimal direction is given by $\mathcal{L}(z, 0) - \mathcal{L}(z, \mathcal{A}(z)) = \mathcal{L}(z, 0) = \Psi^*(z, -D\mathcal{S}_{\mathcal{L}}(z))$;
- the net cost to move in direction s rather than not to move at all is given by $\mathcal{L}(z, s) - \mathcal{L}(z, 0) = \Psi(z, s) + \langle D\mathcal{S}_{\mathcal{L}}(z), s \rangle$.

It is because of this splitting that the integrability condition needs to be checked for the tangent $s = 0$, although it yields a gradient system for the tangent $s = \mathcal{A}_{\mathcal{L}}(\rho)$.

Deviations from detailed balance. In the case of empirical measures *without* detailed balance, certain parts of the properties above remain. The existence of a ‘generalized covector system’ $(\mathcal{Z}, \Psi, \mathcal{V})$ with non-negative Ψ and Ψ^* follows under similar conditions, as we show in Lemma 2.1 and Proposition 2.2. Of course such systems do not identify a functional, and therefore the question whether the system is ‘driven’ by some functional \mathcal{S} remains open.

Interestingly, the relative entropy with respect to any invariant measure π is always a Lyapunov function, as is even every convex functional of the form $\rho \mapsto \int f(d\rho/d\pi) d\pi$. Therefore there is no shortage of Lyapunov functionals; but at this stage we do not have a ‘canonical’ way of coupling any of these to a given function \mathcal{L} .

In specific cases one can say more. In related work [DPZ13b] we showed that the large-deviation rate function \mathcal{L} for the empirical measure of inertial particles (even with interaction) can be written as

$$\mathcal{L}(z, \dot{z}) = \Psi(z, \dot{z} - \mathcal{H}z) + \Psi^*(z, -D\mathcal{S}(z)) + \langle D\mathcal{S}(z), \dot{z} \rangle,$$

where $\mathcal{H}z$ is the non-detailed-balance part of the limiting evolution resulting from inertia, and \mathcal{S} is a relative entropy. Indeed this functional \mathcal{L} does not satisfy the integrability condition (5.3). What forms of \mathcal{L} this example suggests us to expect in more general situations is still unclear.

Multiple gradient-flow structures for finite-state Markov chains. One of the original motivations was to understand the origin of the gradient structure for time-continuous Markov chains with detailed balance

found independently by Maas [Maa11], Chow et al [CHLZ12], and [Mie11b, Mie13b]. This gradient structure is a true gradient structure for the relative entropy in the sense that the dissipation potential is quadratic. However, it turned out that the LDP produces a generalized gradient structure. In fact, these two gradient structures are special cases of a general family where Ψ^* is given in terms of general scalar dissipation functions ψ_{ij} in the form

$$\Psi^*(\rho, \xi) = \sum_{i,j=1}^J L_{ij}(\rho) \psi_{ij}(\xi_j - \xi_i) \quad \text{with } L_{ij}(\rho) = \pi_i Q_{ij} \frac{\frac{\rho_j}{\pi_j} - \frac{\rho_i}{\pi_i}}{\psi'_{ij}\left(\log \frac{\rho_j}{\pi_j} - \log \frac{\rho_i}{\pi_i}\right)} \text{ for } i \neq j,$$

and $L_{jj} \equiv 0$. Here $L_{ij}(\rho)$ is always non-negative since ψ_{ij} is strictly convex. The case $\psi_{ij}(\zeta) = \frac{1}{2}\zeta^2$ leads to the above-mentioned gradient structure, while $\psi_{ij}(\zeta) = \cosh \zeta - 1$ generates the generalized gradient structure obtained from the LDP. It remains an open question to understand the role of these other gradient structures and their relation with stochastic processes.

Another interesting challenge is the stochastic origin of the gradient structure for general nonlinear reaction systems or even reaction-diffusion systems as discussed in [Mie11b]. There again a nonlinear version of the detailed balance condition (also called Wegscheider condition for reaction systems) is crucial for the existence of the gradient structure with the relative entropy as the driving functional.

A Rigorous proof of the large-deviation principle

In this appendix we prove the large-deviation principle (1.15) for finite-state continuous-time Markov chains (the formal calculation can also be found in [Wan12, Sec. 4.3]). For this result we will need the space $BV(0, T; \mathcal{P}(\{1, \dots, J\}))$, defined as a subspace of the functions of bounded variation with values in \mathbb{R}^J , equipped with the total-variation norm,

$$\|\rho\|_{TV} := \sup \left\{ \int_0^T \dot{\phi}_t \cdot \rho_t dt : \phi \in C^1([0, T]; \mathbb{R}^J), \|\phi\|_\infty \leq 1 \right\}.$$

If the dual pairing $\int_0^T \dot{\phi}_t \rho_t dt$ is used, then the space $BV(0, T; \mathbb{R}^J)$ can be interpreted as the dual of $C^1([0, T]; \mathbb{R}^J)$. We then define Θ to be the space $BV(0, T; \mathcal{P}(\{1, \dots, J\}))$, endowed with the weak-* topology.

Theorem A.1. *Consider the Markov processes $X_1(t), X_2(t), \dots$ and the corresponding counting process $\rho^{(n)} : t \mapsto \sum_{k=1}^n \mathbb{1}_{X_k(t)}$ as defined in Section 4.1. Assume that the initial states $\rho_0^{(n)}$ are deterministic and converge in \mathbb{R}^J to the invariant measure μ . Then $\rho^{(n)}$ satisfies a large-deviation principle in the space Θ with good rate function*

$$I(\rho) := \begin{cases} \int_0^T \mathcal{L}(\rho_t, \dot{\rho}_t) dt & \text{if } \rho \text{ is absolutely continuous,} \\ +\infty & \text{otherwise,} \end{cases}$$

where \mathcal{L} is defined in (4.2).

The proof consists of three lemmas. Lemmas A.3 and A.4 together prove the ‘weak’ large-deviation principle. The exponential tightness, proven in Lemma A.2, then guarantees that the (‘strong’) large-deviation principle is also satisfied [DZ87, Lem. 1.2.18].

Lemma A.2. (Exponential tightness) For all $\alpha > 0$ there exists a compact set $K \subset \Theta$ such that

$$\text{Prob}(\rho^{(n)} \notin K) \leq e^{-\alpha n}.$$

Lemma A.3. (Lower bound) For all open sets $O \subset \Theta$,

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in O) \geq -\inf_O I.$$

Lemma A.4. (Upper bound) For all compact sets $C \subset \Theta$,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C) \leq -\inf_C I.$$

We now prove these lemmas.

Proof of Lemma A.2, Exponential tightness. Take an arbitrary $\alpha > 0$, and define for some $M > 0$, to be chosen later, the set

$$K := \overline{\{\rho \in \Theta : \|\rho\|_{TV} \leq M\}}.$$

Here the closure is in the weak-* topology, such that K is automatically compact by the Banach-Alaoglu Theorem.

To estimate $\text{Prob}(\rho^{(n)} \in K^c)$, we first estimate

$$\|\rho^{(n)}\|_{TV} \leq \frac{1}{n} \sum_{k=1}^n \|\mathbb{1}_{X_k}\|_{TV} = \frac{2}{n} \sum_{k=1}^n \Lambda_k, \quad (\text{A.1})$$

where Λ_k is the number of jumps of X_k in the time interval $[0, T]$. Therefore we estimate the probability of the set K^c by estimating the number of jumps of all particles.

We apply a Chernoff estimate to find

$$\begin{aligned} \text{Prob}\left(\frac{2}{n} \sum_{k=1}^n \Lambda_k \geq M\right) &= \text{Prob}\left(\exp \sum_{k=1}^n \Lambda_k \geq \exp \frac{nM}{2}\right) \\ &\leq e^{-\frac{nM}{2}} \mathbb{E}\left(\exp \sum_{k=1}^n \Lambda_k\right) = e^{-\frac{nM}{2}} \prod_{k=1}^n \mathbb{E}(\exp \Lambda_k), \end{aligned}$$

by the independence of the Λ_k . Let $\gamma := \max_{i=1, \dots, J} \sum_{j=1, j \neq i}^J Q_{ij}$ be the maximal rate of jumping away from any site. Comparing the random variable Λ_k with a Poisson process L with rate γ , we find that

$$\mathbb{E}(\exp \Lambda_k) \leq \mathbb{E}(\exp L_T) = e^{\gamma T(e-1)}.$$

Then, continuing the calculation above,

$$\text{Prob} \left(\frac{2}{n} \sum_{k=1}^n \Lambda_k \geq M \right) \leq e^{-\frac{nM}{2}} \prod_{k=1}^n e^{\gamma T(e-1)} = \exp -n \left(\frac{M}{2} - \gamma T(e-1) \right).$$

By choosing $M > 2(\alpha + \gamma T(e-1))$ we have

$$\text{Prob}(\rho^{(n)} \notin K) \leq \text{Prob}(\|\rho^{(n)}\|_{TV} \geq M) \stackrel{(A.1)}{\leq} \text{Prob} \left(\sum_{k=1}^n \Lambda_k \leq \frac{nM}{\sqrt{2}} \right) \leq e^{-\alpha n} \quad \text{for all } n.$$

This proves the assertion. \square

For the upper and lower bound it will be useful to have a dual characterization of the proposed rate function $\int_0^T \mathcal{L}$ similar to the dual characterization of \mathcal{L} itself. For $\rho \in \Theta$ and $\xi \in C_b([0, T]; \mathbb{R}^J)$, define

$$G(\rho, \xi) := \int_0^T \left[\langle \xi_t, \dot{\rho}_t \rangle - \mathcal{H}(\rho_t, \xi_t) \right] dt.$$

Here and below we use the notation ξ_t for the value of ξ at time t , and $\xi_t(i)$ for the i -coordinate of the vector $\xi_t \in \mathbb{R}^J$, for $i \in \{1, \dots, J\}$. Moreover, to simplify notation we write $\langle \xi, \rho \rangle := \sum_{i=1}^J \xi(i) \rho(i)$.

Lemma A.5. For all $\rho \in \Theta$,

$$I(\rho) = \sup \{ G(\rho, \xi) : \xi \in C_b([0, T]; \mathbb{R}^J) \}. \quad (\text{A.2})$$

Proof. The proof is standard, and proceeds in three steps:

- 1 The inequality “ \geq ” follows directly from the duality relation (4.2).
- 2 The implication “ ρ not absolutely continuous \implies the supremum is $+\infty$ ” follows from an explicit choice of subintervals in which $|\dot{\rho}|$ is large; by concentrating ξ on those subintervals the supremum can be made as large as necessary.
- 3 Finally, the inequality “ \leq ” follows by a standard approximation argument, in which ξ is constructed to approximate the pointwise optimum in the duality (4.2).

\square

We also need a Girsanov-type theorem that characterizes the Radon-Nikodym derivative of the original process with respect to a modified process. It is no coincidence that the functional G defined above arises as in this characterization. To formulate this result we first introduce some notation. Let \mathbb{P} be the law of a single process $X_k(t)$, generated by \mathcal{Q} , and for any $\xi \in C^1([0, T]; \mathbb{R}^J)$, let \mathbb{P}_ξ be the law of a single process generated by the modified, time-dependent generator

$$(\mathcal{Q}_\xi(t)g)_i := \sum_{j=1}^J Q_{ij} e^{\xi_t(j) - \xi_t(i)} (g(j) - g(i)), \quad \text{for } g \in \mathbb{R}^J, i \in \{1, \dots, J\}.$$

We write $\mathbb{P}^{(n)}$ and $\mathbb{E}^{(n)}$ for the law and expectation of the empirical process $\rho^{(n)}$, and $\mathbb{P}_\xi^{(n)}$ and $\mathbb{E}_\xi^{(n)}$ for the law and expectation of the modified empirical process.

Lemma A.6. (A Girsanov-type characterization) For $\mathbb{P}^{(n)}$ -a.e. ρ of the form $t \mapsto \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{x_k(t)}$,

$$\frac{1}{n} \log \frac{d\mathbb{P}_\xi^{(n)}}{d\mathbb{P}^{(n)}}(\rho) = G(\rho, \xi).$$

Proof. Proposition 7.3 of the Appendix 1 of [KL99] gives that for any càdlàg function $x : [0, T] \rightarrow \{1, \dots, J\}$,

$$\log \frac{d\mathbb{P}_\xi}{d\mathbb{P}}(x) = \xi_T(x_{T-}) - \xi_0(x_0) - \int_0^T e^{-\xi_t(x_t)} \left[(\partial_t + \mathcal{Q}) e^{\xi_t} \right](x_t) dt. \quad (\text{A.3})$$

We calculate for $i \in \{1, \dots, J\}$ that (using convention (4.1))

$$e^{-\xi_t(i)} (\mathcal{Q} e^{\xi_t})(i) = \sum_{j=1}^J Q_{ij} e^{\xi_t(j) - \xi_t(i)} = \sum_{j=1}^J Q_{ij} (e^{\xi_t(j) - \xi_t(i)} - 1) = \mathcal{H}(\mathbb{1}_x, \xi_t),$$

and $e^{-\xi_t} \partial_t (e^{\xi_t}) = \dot{\xi}_t$. Therefore, for a càdlàg function $x : [0, T] \rightarrow \{1, \dots, J\}$,

$$\begin{aligned} \log \frac{d\mathbb{P}_\xi}{d\mathbb{P}}(x) &= \xi_T(x_{T-}) - \xi_0(x_0) - \int_0^T e^{-\xi_t(x_t)} (\partial_t e^{\xi_t})(x_t) dt - \int_0^T \mathcal{H}(\mathbb{1}_{x_t}, \xi_t) dt \\ &= \langle \xi_T, \mathbb{1}_{x_{T-}} \rangle - \langle \xi_0, \mathbb{1}_{x_0} \rangle - \int_0^T \langle \dot{\xi}_t, \mathbb{1}_{x_t} \rangle dt - \int_0^T \mathcal{H}(\mathbb{1}_{x_t}, \xi_t) dt. \end{aligned}$$

Since the time evolution of the n processes is independent, and the Radon-Nikodym derivatives below are invariant under the push-forward $(x_1, \dots, x_n) \mapsto \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{x_k} =: \rho$, we have

$$\begin{aligned} \frac{1}{n} \log \frac{d\mathbb{P}_\xi^{(n)}}{d\mathbb{P}^{(n)}}(\rho) &= \frac{1}{n} \log \prod_{k=1}^n \frac{d\mathbb{P}_\xi}{d\mathbb{P}}(x_k) = \frac{1}{n} \sum_{k=1}^n \log \frac{d\mathbb{P}_\xi}{d\mathbb{P}}(x_k) \\ &= \langle \xi_T, \rho_{T-} \rangle - \langle \xi_0, \rho_0 \rangle - \int_0^T \langle \dot{\xi}_t, \rho_t \rangle dt - \int_0^T \mathcal{H}(\rho_t, \xi_t) dt. \end{aligned}$$

With probability one, each process only jumps a finite number of times, and therefore $\mathbb{P}^{(n)}$ -almost every ρ is piecewise constant, with a finite number of jumps. Therefore the time derivative $\dot{\rho}$ is a finite measure, so that we can apply partial integration to obtain

$$\frac{1}{n} \log \frac{d\mathbb{P}_\xi^{(n)}}{d\mathbb{P}^{(n)}}(\rho) = \int_0^T \left[\langle \xi_t, \dot{\rho}_t \rangle - \mathcal{H}(\rho_t, \xi_t) \right] dt = G(\rho, \xi).$$

□

We now continue with the proofs of the upper and lower bounds.

Proof of Lemma A.3, the lower bound. Fix an open set $O \subset \Theta$ and $\delta > 0$. Choose $\rho \in \Theta$, smooth as a function from $[0, T]$ into \mathbb{R}^J , such that

$$I(\rho) \leq \delta + \inf_O I.$$

Define $\xi \in C^1([0, T]; \mathbb{R}^J)$ to achieve the supremum in (A.2).

Choose $\varepsilon > 0$ such that the \mathbb{R}^J -ball $B(\rho, \varepsilon)$ is contained in O , so that

$$\begin{aligned} \text{Prob}(\rho^{(n)} \in O) &\geq \text{Prob}(\rho^{(n)} \in B(\rho, \varepsilon)) = \mathbb{E}^{(n)}(\mathbf{1}_{B(\rho, \varepsilon)}) \\ &= \mathbb{E}_\xi^{(n)} \left(\mathbf{1}_{B(\rho, \varepsilon)} \frac{d\mathbb{P}^{(n)}}{d\mathbb{P}_\xi^{(n)}} \right) \geq \left(\mathbb{P}_\xi^{(n)}\text{-ess inf}_{B(\rho, \varepsilon)} \frac{d\mathbb{P}^{(n)}}{d\mathbb{P}_\xi^{(n)}} \right) \mathbb{E}_\xi^{(n)}(\mathbf{1}_{B(\rho, \varepsilon)}). \end{aligned} \quad (\text{A.4})$$

Since ξ achieves the optimum in (A.2), we have for all g that

$$\begin{aligned} 0 &= \langle D_\xi G(\rho, \xi), g \rangle \\ &= \int_0^T \{ \langle g_t, \dot{\rho}_t \rangle - \langle D_\xi \mathcal{H}(\rho_t, \xi_t), g_t \rangle \} dt \\ &= \int_0^T \left\{ \langle g_t, \dot{\rho}_t \rangle - \sum_{i,j=1}^J \rho_t(i) Q_{ij} e^{\xi_t(j) - \xi_t(i)} (g_t(j) - g_t(i)) \right\} dt \\ &= \int_0^T \{ \langle g_t, \dot{\rho}_t \rangle - \langle \mathcal{Q}_\xi(t) g_t, \rho_t \rangle \} dt. \end{aligned}$$

From this it follows that ξ is chosen such that \mathcal{Q}_ξ generates the process described by ρ . Therefore, by the law of large numbers,

$$\mathbb{E}_\xi^{(n)}(\mathbf{1}_{B(\rho, \varepsilon)}) \longrightarrow 1 \quad \text{as } n \rightarrow \infty. \quad (\text{A.5})$$

Since $\mathbb{P}_\xi^{(n)}$ -almost every $\hat{\rho} \in B(\rho, \varepsilon)$ is of the form $\hat{\rho} = \frac{1}{n} \sum_{k=1}^n \mathbf{1}_{x_k}$, Lemma A.6 gives

$$\begin{aligned} \frac{1}{n} \log \frac{d\mathbb{P}^{(n)}}{d\mathbb{P}_\xi^{(n)}}(\hat{\rho}) &= -G(\hat{\rho}, \xi) = -G(\rho, \xi) + G(\rho - \hat{\rho}, \xi) \\ &\geq -I(\rho) - \sup_{\hat{\rho} \in B(\rho, \varepsilon)} |G(\rho - \hat{\rho}, \xi)| \\ &\geq -\inf_O I - \delta - \sup_{\hat{\rho} \in B(\rho, \varepsilon)} |G(\rho - \hat{\rho}, \xi)|. \end{aligned} \quad (\text{A.6})$$

Here, the last term converges to zero as $\varepsilon \rightarrow 0$, since for fixed ξ the mapping $\rho \mapsto G(\rho, \xi)$ is continuous. Taking the limit first as $\varepsilon \rightarrow 0$ and then $\delta \rightarrow 0$ the claim follows from (A.4), (A.5) and (A.6). \square

Proof of Lemma A.4, the upper bound. Fix a compact set $C \subset \Theta$. First we show that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C) \leq - \sup_{\cup_k C_k \supset C} \min_k \sup_{\xi} \inf_{\rho \in C_k} G(\rho, \xi), \quad (\text{A.7})$$

where the first supremum is taken over all finite coverings $\{C_k\}_{k=1}^K$ of C .

To prove this inequality, first note that for any such covering, $\text{Prob}(\rho^{(n)} \in C) \leq \sum_k \text{Prob}_n(\rho^{(n)} \in C_k)$, and therefore

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C) \leq \max_k \limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C_k). \quad (\text{A.8})$$

We then estimate, using Lemma A.6,

$$\begin{aligned} \text{Prob}(\rho^{(n)} \in C_k) &= \mathbb{E}^{(n)}(\mathbf{1}_{C_k}) = \mathbb{E}_\xi^{(n)} \left(\mathbf{1}_{C_k} \frac{d\mathbb{P}^n}{d\mathbb{P}_\xi^{(n)}} \right) \leq \mathbb{P}_\xi^{(n)}\text{-ess sup}_{C_k} \frac{d\mathbb{P}^n}{d\mathbb{P}_\xi^{(n)}} \\ &\leq -\mathbb{P}_\xi^{(n)}\text{-ess inf}_{C_k} G(\cdot, \xi) \leq -\inf_{C_k} G(\cdot, \xi). \end{aligned} \quad (\text{A.9})$$

Taking the limit $n \rightarrow \infty$ and the infimum over ξ , we find

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C_k) \leq -\sup_\xi \inf_{C_k} G(\cdot, \xi).$$

Therefore, (A.8) can be further estimated by

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \text{Prob}(\rho^{(n)} \in C) \leq -\max_k \sup_\xi \inf_{C_k} G(\cdot, \xi),$$

and since this holds for all coverings $\{C_k\}$ the inequality (A.7) follows.

We now show that

$$\sup_{\cup_k C_k \supset C} \min_k \sup_\xi \inf_{\rho \in C_k} G(\rho, \xi) \geq \inf_{\rho \in C} \sup_\xi G(\rho, \xi) = \inf_C I, \quad (\text{A.10})$$

which, together with (A.7) will conclude the proof. Naturally, the last equality follows from (A.2). To show the inequality in (A.10), fix a $\delta > 0$ and let $\eta := \inf_{\rho \in C} \sup_\xi G(\rho, \xi)$. Observe that for each ρ there exists an ξ_ρ such that $G(\rho, \xi_\rho) \geq \eta - \delta$. Since $\rho \mapsto G(\rho, \xi)$ is continuous for any ξ , it follows that

$$\forall \rho \in C \exists \xi_\rho \exists O_\rho \ni \rho \text{ open such that } \inf_{O_\rho} G(\cdot, \xi_\rho) \geq \eta - 2\delta,$$

from which we deduce that

$$\forall \rho : \sup_\xi \inf_{O_\rho} G(\cdot, \xi) \geq \eta - 2\delta. \quad (\text{A.11})$$

Since C is compact and $C \subset \bigcup_{\rho \in C} O_\rho$, there is a finite covering $C \subset \bigcup_{k=1}^K O_{\rho_k}$; by (A.11) we have

$$\min_k \sup_\xi \inf_{O_{\rho_k}} G(\cdot, \xi) \geq \eta - 2\delta.$$

Trivially, we have

$$\sup_{\cup_k C_k \supset C} \min_k \sup_\xi \inf_{O_{\rho_k}} G(\cdot, \xi) \geq \eta - 2\delta.$$

Since $\delta > 0$ is arbitrary, the inequality (A.10) follows. \square

References

- [ABM06] H. Attouch, G. Buttazzo, and G. Michaille. *Variational analysis in Sobolev and BV spaces : applications to PDEs and optimization*. SIAM, Philadelphia, PA, USA, 2006.
- [ADPZ11] S. Adams, N. Dirr, M. A. Peletier, and J. Zimmer. From a large-deviations principle to the Wasserstein gradient flow: a new micro-macro passage. *Communications in Mathematical Physics*, 307(3):791–815, 2011.
- [ADPZ12] S. Adams, N. Dirr, M. A. Peletier, and J. Zimmer. Large deviations and gradient flows. <http://arxiv.org/abs/1201.4601> (to appear in *Philosophical Transactions A*), 2012.
- [AGS08] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Lectures in Mathematics. ETH Zürich. Birkhauser, Basel, Switzerland, 2nd edition, 2008.
- [AMP⁺12] S. Arnrich, A. Mielke, M.A. Peletier, G. Savaré, and M. Veneroni. Passing to the limit in a wasserstein gradient flow: from diffusion to reaction. *Calculus of Variations and Partial Differential Equations*, 44(3–4):419–454, 2012.
- [CHLZ12] S.N. Chow, W. Huang, Y. Li, and H. Zhou. Fokker–Planck equations for a free energy functional or Markov process on a graph. *Arch. Rational Mech. Anal.*, 203(3):969–1008, 2012.
- [CV90] P. Colli and A. Visintin. On a class of doubly nonlinear evolution equations. *Comm. Partial Differential Equations*, 15(5):737–756, 1990.
- [Dac00] B. Dacorogna. *Direct Methods in the Calculus of Variations*. Springer, New York, NY, USA, second edition, 2000.
- [DG87] D.A. Dawson and J. Gärtner. Large deviations from the McKean-Vlasov limit for weakly interacting diffusions. *Stochastics*, 20(4):247–308, 1987.
- [DG93] Ennio De Giorgi. New problems on minimizing movements. In *Boundary value problems for partial differential equations and applications*, volume 29 of *RMA Res. Notes Appl. Math.*, pages 81–98. Masson, Paris, 1993.
- [dH00] F. den Hollander. *Large deviations*. Fields Institute Monographs. American Mathematical Society, 2000.
- [DLR13] M. H. Duong, V. Laschos, and M. Renger. Wasserstein gradient flows from large deviations of many-particle limits. *ESAIM: Control, Optimisation and Calculus of Variations*, 19(4):1166–1188, 2013.
- [DLZ12] N. Dirr, V. Laschos, and J. Zimmer. Upscaling from particle models to entropic gradient flows. *Journal of Mathematical Physics*, 53(6), 2012.

- [DMT80] Ennio De Giorgi, Antonio Marino, and Mario Tosques. Problems of evolution in metric spaces and maximal decreasing curve. *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Natur.* (8), 68(3):180–187, 1980.
- [DPZ13a] M.H. Duong, M.A. Peletier, and J. Zimmer. Conservative-dissipative approximation schemes for a generalized Kramers equation. <http://arxiv.org/abs/1302.1024>, 2013.
- [DPZ13b] M.H. Duong, M.A. Peletier, and J. Zimmer. GENERIC formalism of a Vlasov-Fokker-Planck equation and connection to large-deviation principles. <http://arxiv.org/abs/1302.1024>, 2013.
- [Dud89] R.M. Dudley. *Real analysis and probability*. Wadsworth & Brooks/Cole, Pacific Grove, CA, USA, 1989.
- [DZ87] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Stochastic modelling and applied probability*. Springer, New York, NY, USA, 2nd edition, 1987.
- [FK06] J. Feng and T.G. Kurtz. *Large deviations for stochastic processes*, volume 131 of *Mathematical surveys and monographs*. American Mathematical Society, Providence, RI, USA, 2006.
- [Giu03] E. Giusti. *Direct Methods in the Calculus of Variations*. World Scientific, Singapore, 2003.
- [JKO98] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [KL99] C. Kipnis and C. Landim. *Scaling limits of interacting particle systems*. Springer, Berlin-Heidelberg, Germany, 1999.
- [Léo07] C. Léonard. A large deviation approach to optimal transport. arxiv.org/abs/0710.1461v1, 2007.
- [Lig85] T.M. Liggett. *Interacting particle systems*. Springer, Berlin-Heidelberg, Germany, 1985.
- [LM12] M. Liero and A. Mielke. Gradient structures and geodesic convexity for reaction-diffusion systems, 2012. To appear. WIAS preprint 1701 (April 2012).
- [LS95] S. Luckhaus and T. Sturzenhecker. Implicit time discretization for the mean curvature flow equation. *Calculus of Variations and Partial Differential Equations*, 3(2):253–271, 1995.
- [Maa11] J. Maas. Gradient flows of the entropy for finite Markov chains. *Functional analysis*, 261(8):2250–2292, 2011.
- [Mae99] C. Maes. The fluctuation theorem as a Gibbs property. *Journal of statistical physics*, 95(1/2):367–392, 1999.

- [Mie03] A. Mielke. Energetic formulation of multiplicative elasto–plasticity using dissipation distances. *Continuum Mechanics and Thermodynamics*, 15(4):351–382, 2003.
- [Mie11a] A. Mielke. Formulation of thermoelastic dissipative material behavior using GENERIC. *Continuum Mechanics and Thermodynamics*, 23(3):233–256, 2011.
- [Mie11b] A. Mielke. A gradient structure for reaction-diffusion systems and for energy-drift-diffusion systems. *Nonlinearity*, 24(4):1329–1346, 2011.
- [Mie13a] A. Mielke. Deriving amplitude equations via evolutionary Γ -convergence. *In preparation*, 2013. In preparation.
- [Mie13b] A. Mielke. Geodesic convexity of the relative entropy in reversible Markov chains. *Calculus of Variations and Partial Differential Equations*, 48(1):1–31, 2013.
- [MM13] J. Maas and A. Mielke. On gradient structures for reaction rate and chemical master equations with detailed balance. *In preparation*, 2013.
- [MMS09] D. Matthes, R. J. McCann, and G. Savaré. A family of nonlinear fourth order equations of gradient flow type. *Communications in Partial Differential Equations*, 34(11):1352–1397, 2009.
- [MRS08] A. Mielke, R. Rossi, and G. Savaré. A metric approach to a class of doubly nonlinear evolution equations and applications. *Ann. Scuola Norm. Sup. Pisa Cl. Sci.*, VII(5):97–169, 2008.
- [MRV00] C. Maes, F. Redig, and A. Van Moffaert. On the definition of entropy production, via examples. *Journal of Mathematical Physics*, 41(3):1528–1554, 2000.
- [Ott01] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. *Communications in partial differential equations*, 26(1&2):101–174, 2001.
- [PP10] J.W. Portegies and M.A. Peletier. Well-posedness of a parabolic moving-boundary problem in the setting of Wasserstein gradient flows. *Interfaces and Free Boundaries*, 12(2):121–150, 2010.
- [PRV13] M.A. Peletier, D.R.M. Renger, and M. Veneroni. Variational formulation of the Fokker-Planck equation with decay: a particle approach. *Communications in Contemporary Mathematics*, 15(5):1350017, 2013.
- [Ren13] D.R.M. Renger. *Microscopic interpretation of Wasserstein gradient flows*. PhD thesis, Technische Universiteit Eindhoven, Eindhoven, the Netherlands, 2013. <http://alexandria.tue.nl/extra2/749143.pdf>.
- [RMS08] R. Rossi, A. Mielke, and G. Savaré. A metric approach to a class of doubly nonlinear evolution equations and applications. *Ann. Scuola Norm. Sup. Pisa Cl. Sci.*, VII(5):97–169, 2008.

- [RSV13] T. Roubíček, O. Souček, and R. Vodička. A model of rupturing lithospheric faults with reoccurring earthquakes. *SIAM J. Appl. Math.*, 73(4):1460–1488, 2013.
- [SS04] E. Sandier and S. Serfaty. Gamma-convergence of gradient flows with applications to Ginzburg-Landau. *Communications on Pure and Applied Mathematics*, 57(12):1627–1672, 2004.
- [Vil09] C. Villani. *Optimal transport: old and new*. Springer, Berlin-Heidelberg, Germany, 2009.
- [Wan12] F. Wang. *Dynamical Gibbs-non-Gibbs transitions: a study via coupling and large deviations*. PhD thesis, Leiden Universiteit, Leiden, the Netherlands, 2012. <http://hdl.handle.net/1887/20110>.