On a Local Lyapunov function for the McKean-Vlasov Dynamics

Rajesh Sundaresan

Indian Institute of Science

August 2017

A continuous-time Markov chain

- S = finite set, a state space for a continuous-time Markov chain.

A continuous-time Markov chain

- $S =$ finite set, a state space for a continuous-time Markov chain.
- $L =$ rate matrix of a Markov chain on S.

A continuous-time Markov chain

- $S =$ finite set, a state space for a continuous-time Markov chain.
- $L =$ rate matrix of a Markov chain on S.
- $L(i,j) =$ rate of a transition $i \rightsquigarrow j$. Assume it is irreducible.
- $L(i,i) = -\sum_{j:j\neq i} L(i,j)$.

A continuous-time Markov chain

- $S =$ finite set, a state space for a continuous-time Markov chain.
- $L =$ rate matrix of a Markov chain on S.
- $L(i,j) =$ rate of a transition $i \rightsquigarrow j$. Assume it is irreducible.
- $L(i,i) = -\sum_{j:j\neq i} L(i,j)$.
- $(X(t) : t \geq 0)$: a continuous-time Markov chain on S with $X(0) \sim \nu_0$, some initial distribution.

A continuous-time Markov chain

- S = finite set, a state space for a continuous-time Markov chain.
- L = rate matrix of a Markov chain on S.
- $L(i,j)$ = rate of a transition $i \rightsquigarrow j$. Assume it is irreducible.
- $L(i,i) = -\sum_{j:j \neq i} L(i,j)$.
- $(X(t) : t \geq 0)$: a continuous-time Markov chain on S with $X(0) \sim \nu_0$, some initial distribution.
- $\eta(t) := Law(X(t))$ for $t \geq 0$; column vector.

A continuous-time Markov chain

- $S =$ finite set, a state space for a continuous-time Markov chain.
- $L =$ rate matrix of a Markov chain on S.
- $L(i,j) =$ rate of a transition $i \rightsquigarrow j$. Assume it is irreducible.
- $L(i,i) = -\sum_{j:j\neq i} L(i,j)$.
- $(X(t) : t \geq 0)$: a continuous-time Markov chain on S with $X(0) \sim \nu_0$, some initial distribution.
- $\eta(t) := Law(X(t))$ for $t \geq 0$; column vector.
- Forward equation (FE):

$$\dot{\eta}(t) = L^* \eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

A Lyapunov function for the dynamics

- Let ν be the invariant measure or stationary distribution of the CTMC, a resting point for the FE dynamics, i.e.,

$$L^* \nu = 0.$$

A Lyapunov function for the dynamics

- Let ν be the invariant measure or stationary distribution of the CTMC, a resting point for the FE dynamics, i.e.,

$$L^*\nu = 0.$$

- The dynamics is not necessarily reversible.

A Lyapunov function for the dynamics

- Let ν be the invariant measure or stationary distribution of the CTMC, a resting point for the FE dynamics, i.e.,

$$L^*\nu = 0.$$

- The dynamics is not necessarily reversible.

- Well-known: Relative entropy with respect to the invariant measure $I(\cdot|\nu)$ is a Lyapunov function for the dynamics, i.e.,

$$I(\eta(t)|\nu) \leq I(\eta(s)|\nu) \text{ for } t > s \geq 0.$$

- $I(\mu|\nu) = \sum_{i \in S} \mu_i \log(\mu_i/\nu_i)$.

The nonlinear (McKean-Vlasov) equation

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

- The McKean-Vlasov equation (MVE):

$$\dot{\eta}(t) = L^*_{\eta(t)}\eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

- The McKean-Vlasov equation (MVE):

$$\dot{\eta}(t) = L^*_{\eta(t)}\eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

- Question: What is a Lyapunov function for this dynamics?

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

- The McKean-Vlasov equation (MVE):

$$\dot{\eta}(t) = L^*_{\eta(t)}\eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

- Question: What is a Lyapunov function for this dynamics?

- For simplicity, let us assume that the MVE dynamics has a rest point ν, satisfying $L^*_\nu \nu = 0$, that is globally asympotically stable.

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

- The McKean-Vlasov equation (MVE):

$$\dot{\eta}(t) = L^*_{\eta(t)}\eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

- Question: What is a Lyapunov function for this dynamics?

- For simplicity, let us assume that the MVE dynamics has a rest point ν, satisfying $L^*_\nu \nu = 0$, that is globally asympotically stable. We then know that there is a Lyapunov function.

The nonlinear (McKean-Vlasov) equation

- Suppose L now depends on η, i.e.,

$$L(i,j) \text{ is replaced by } L_\eta(i,j).$$

- The McKean-Vlasov equation (MVE):

$$\dot{\eta}(t) = L^*_{\eta(t)}\eta(t), \quad t \geq 0, \text{ with } \eta(0) = \nu_0.$$

- Question: What is a Lyapunov function for this dynamics?

- For simplicity, let us assume that the MVE dynamics has a rest point ν, satisfying $L^*_\nu \nu = 0$, that is globally asympotically stable.
  We then know that there is a Lyapunov function.
  Identify one.

How does the MVE arise?

How does the MVE arise?

- N interacting particles.

How does the MVE arise?

- N interacting particles.
- $X_n(t) =$ state of the n^{th} particle at time t; $X_n(t) \in S$.

How does the MVE arise?

- N interacting particles.
- $X_n(t)$ = state of the n^{th} particle at time t; $X_n(t) \in S$.
- To describe a weak mean-field interaction, let
- $\eta_N(t)$ = empirical measure of the system (distribution of the N particles across states) at time t.

$$\eta_N(t) = \frac{1}{N} \sum_{n=1}^{N} \delta_{X_n(t)}.$$

How does the MVE arise?

- N interacting particles.
- $X_n(t)$ = state of the n^{th} particle at time t; $X_n(t) \in S$.
- To describe a weak mean-field interaction, let
- $\eta_N(t)$ = empirical measure of the system (distribution of the N particles across states) at time t.

$$\eta_N(t) = \frac{1}{N} \sum_{n=1}^{N} \delta_{X_n(t)}.$$

- Mean-field interactions: The transition rate of a particle depends on its state, and on the states of the other particles, but only through the empirical measure.

How does the MVE arise?

- N interacting particles.
- $X_n(t)$ = state of the n^{th} particle at time t; $X_n(t) \in S$.
- To describe a weak mean-field interaction, let
- $\eta_N(t)$ = empirical measure of the system (distribution of the N particles across states) at time t.

$$\eta_N(t) = \frac{1}{N} \sum_{n=1}^{N} \delta_{X_n(t)}.$$

- Mean-field interactions: The transition rate of a particle depends on its state, and on the states of the other particles, but only through the empirical measure.
- $L_{\eta_N(t)}(i,j)$ = rate of an $i \rightsquigarrow j$ transition at time t for a particle n in state $X_n(t) = i$.

Assumptions and a limit theorem

Assumptions and a limit theorem

- 1. Suppose $L_\eta \geq \rho \underline{L}$, where $\underline{L}(i,j) \in \{0,1\}$ and $\rho > 0$.

Assumptions and a limit theorem

- 1. Suppose $L_\eta \geq \rho \underline{L}$, where $\underline{L}(i,j) \in \{0,1\}$ and $\rho > 0$.
- $\underline{L}(i,j)$ encodes information on which transitions are allowed and which are disallowed. Those that are allowed happen at a uniform rate.

Assumptions and a limit theorem

- 1. Suppose $L_\eta \geq \rho \underline{L}$, where $\underline{L}(i,j) \in \{0,1\}$ and $\rho > 0$.
- $\underline{L}(i,j)$ encodes information on which transitions are allowed and which are disallowed. Those that are allowed happen at a uniform rate.
- 2. The mapping $\eta \mapsto L_\eta(i,j)$ is Lipschitz.

Assumptions and a limit theorem

- 1. Suppose $L_\eta \geq \rho \underline{L}$, where $\underline{L}(i,j) \in \{0,1\}$ and $\rho > 0$.
- $\underline{L}(i,j)$ encodes information on which transitions are allowed and which are disallowed. Those that are allowed happen at a uniform rate.
- 2. The mapping $\eta \mapsto L_\eta(i,j)$ is Lipschitz.
- 3. Let $\eta_N(0) \to \nu_0$, a deterministic measure, in probability.

Assumptions and a limit theorem

- 1. Suppose $L_\eta \geq \rho \underline{L}$, where $\underline{L}(i,j) \in \{0,1\}$ and $\rho > 0$.
- $\underline{L}(i,j)$ encodes information on which transitions are allowed and which are disallowed. Those that are allowed happen at a uniform rate.
- 2. The mapping $\eta \mapsto L_\eta(i,j)$ is Lipschitz.
- 3. Let $\eta_N(0) \to \nu_0$, a deterministic measure, in probability.

Theorem

Under the above assumptions, we have the following convergence (uniform over compacts) in probability:

$$\eta_N(\cdot) \to \nu(\cdot),$$

where $\nu(\cdot)$ is the solution to the MVE with initial condition ν_0.

Interesting asides

Interesting asides

- Arises in the analysis of Wi-Fi and closed queueing systems.

Interesting asides

- Arises in the analysis of Wi-Fi and closed queueing systems.
- If the initial measure is exchangeable ...
  - Asymptotically, each individual particle's state is distributed according to $\nu(t)$,

Interesting asides

- Arises in the analysis of Wi-Fi and closed queueing systems.
- If the initial measure is exchangeable ...
  - Asymptotically, each individual particle's state is distributed according to $\nu(t)$, and it evolves according to a time-inhomogeneous Markov process with transition rates given by the rate matrix $L_{\nu(t)}$.

Interesting asides

- Arises in the analysis of Wi-Fi and closed queueing systems.
- If the initial measure is exchangeable ...
  - Asymptotically, each individual particle's state is distributed according to $\nu(t)$, and it evolves according to a time-inhomogeneous Markov process with transition rates given by the rate matrix $L_{\nu(t)}$.
  - Asymptotically, the particle evolutions decouple; "propagation of chaos".

Interesting asides

- Arises in the analysis of Wi-Fi and closed queueing systems.
- If the initial measure is exchangeable ...
  - Asymptotically, each individual particle's state is distributed according to $\nu(t)$, and it evolves according to a time-inhomogeneous Markov process with transition rates given by the rate matrix $L_{\nu(t)}$.
  - Asymptotically, the particle evolutions decouple; "propagation of chaos".
  - If ν, the rest point for the MVE, is globally asymptotically stable, then the equilibrated distribution of k tagged particles is $\nu^{\otimes k}$.

Transition rates for the Markov process $\eta_N(\cdot)$

Transition rates for the Markov process $\eta_N(\cdot)$

- $(\eta_N(t), t \geq 0)$ is a Markov process with state-space $\Delta_N(S)$, the probability simplex of empirical measures of N particles.

Transition rates for the Markov process $\eta_N(\cdot)$

- $(\eta_N(t), t \geq 0)$ is a Markov process with state-space $\Delta_N(S)$, the probability simplex of empirical measures of N particles.

- The transition
$$\xi \rightsquigarrow \xi^{i,j} := \xi + \frac{1}{N}\delta_j - \frac{1}{N}\delta_i$$
occurs at rate $N\xi_i L_\xi(i,j)$.

Transition rates for the Markov process $\eta_N(\cdot)$

- $(\eta_N(t), t \geq 0)$ is a Markov process with state-space $\Delta_N(S)$, the probability simplex of empirical measures of N particles.
- The transition
$$\xi \rightsquigarrow \xi^{i,j} := \xi + \frac{1}{N}\delta_j - \frac{1}{N}\delta_i$$
occurs at rate $N\xi_i L_\xi(i,j)$.
- The rate matrix of this Markov process is
$$\mathscr{L}_N(\xi, \xi^{i,j}) = N\xi_i L_\xi(i,j).$$
Also it is irreducible if $\underline{L}$ is so.

A candidate prelimit Lyapunov function

A candidate prelimit Lyapunov function

- $p_N(t) = Law(\eta_N(t))$ for $t \geq 0$.

A candidate prelimit Lyapunov function

- $p_N(t) = Law(\eta_N(t))$ for $t \geq 0$.
- The corresponding forward equation is:

$$\dot{p}_N(t) = \mathscr{L}_N^* p_N(t), \quad t \geq 0, \text{ with } p_N(0) = Law(\eta_N(0)).$$

A candidate prelimit Lyapunov function

- $p_N(t) = Law(\eta_N(t))$ for $t \geq 0$.
- The corresponding forward equation is:

$$\dot{p}_N(t) = \mathscr{L}_N^* p_N(t), \quad t \geq 0, \text{ with } p_N(0) = Law(\eta_N(0)).$$

- Let p_N be the (unique) invariant measure, i.e., $\mathscr{L}_N^* p_N = 0$. Then by the Lyapunov function property of relative entropy, for $t > s \geq 0$, we get

$$I(p_N(t)|p_N) \leq I(p_N(s)|p_N)$$

A candidate prelimit Lyapunov function

- $p_N(t) = Law(\eta_N(t))$ for $t \geq 0$.
- The corresponding forward equation is:

$$\dot{p}_N(t) = \mathscr{L}_N^* p_N(t), \quad t \geq 0, \text{ with } p_N(0) = Law(\eta_N(0)).$$

- Let p_N be the (unique) invariant measure, i.e., $\mathscr{L}_N^* p_N = 0$. Then by the Lyapunov function property of relative entropy, for $t > s \geq 0$, we get

$$I(p_N(t)|p_N) \leq I(p_N(s)|p_N)$$

or equivalently, after division by N,

$$\frac{1}{N} \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \log \left(\frac{p_N(t)(\xi)}{p_N(\xi)} \right) \leq \frac{1}{N} \sum_{\xi \in \Delta_N(S)} p_N(s)(\xi) \log \left(\frac{p_N(s)(\xi)}{p_N(\xi)} \right)$$

Manipulating the prelimit Lyapunov function

- Consider the LHS which is nonincreasing in t:

$$\frac{1}{N} \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \log \left(\frac{p_N(t)(\xi)}{p_N(\xi)} \right)$$

Manipulating the prelimit Lyapunov function

- Consider the LHS which is nonincreasing in t:

$$\frac{1}{N} \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \log \left(\frac{p_N(t)(\xi)}{p_N(\xi)} \right)$$

- Rewriting

$$\sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \left(-\frac{1}{N} \log p_N(\xi) \right) - \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \left(-\frac{1}{N} \log p_N(t)(\xi) \right)$$

Manipulating the prelimit Lyapunov function

- Consider the LHS which is nonincreasing in t:

$$\frac{1}{N}\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\log\left(\frac{p_N(t)(\xi)}{p_N(\xi)}\right)$$

- Rewriting

$$\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\left(-\frac{1}{N}\log p_N(\xi)\right)-\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\left(-\frac{1}{N}\log p_N(t)(\xi)\right)$$

- By the limit theorem, we know $\eta_N(t)\to\nu(t)$ in probability, where $\nu(t)$ is solution to the MVE evaluated at t, or $p_N(t)\to\delta_{\nu(t)}$ weakly.

Manipulating the prelimit Lyapunov function

- Consider the LHS which is nonincreasing in t:

$$\frac{1}{N} \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \log \left(\frac{p_N(t)(\xi)}{p_N(\xi)} \right)$$

- Rewriting

$$\sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \left(-\frac{1}{N} \log p_N(\xi) \right) - \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) \left(-\frac{1}{N} \log p_N(t)(\xi) \right)$$

- By the limit theorem, we know $\eta_N(t) \to \nu(t)$ in probability, where $\nu(t)$ is solution to the MVE evaluated at t, or $p_N(t) \to \delta_{\nu(t)}$ weakly.

- Interpolate $-\frac{1}{N} \log p_N(t)(\cdot)$ to get a (cts) function $J_{N,t}(\cdot)$ on $\Delta(S)$.

Manipulating the prelimit Lyapunov function

- Consider the LHS which is nonincreasing in t:

$$\frac{1}{N}\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\log\left(\frac{p_N(t)(\xi)}{p_N(\xi)}\right)$$

- Rewriting

$$\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\left(-\frac{1}{N}\log p_N(\xi)\right)-\sum_{\xi\in\Delta_N(S)} p_N(t)(\xi)\left(-\frac{1}{N}\log p_N(t)(\xi)\right)$$

- By the limit theorem, we know $\eta_N(t)\to\nu(t)$ in probability, where $\nu(t)$ is solution to the MVE evaluated at t, or $p_N(t)\to\delta_{\nu(t)}$ weakly.

- Interpolate $-\frac{1}{N}\log p_N(t)(\cdot)$ to get a (cts) function $J_{N,t}(\cdot)$ on $\Delta(S)$. Interpolate $-\frac{1}{N}\log p_N(\cdot)$ to get a (cts) function $J_N(\cdot)$ on $\Delta(S)$.

A Lyapunov function

A Lyapunov function

Theorem

$J_{N,t}(\cdot) \rightarrow J_t(\cdot)$, and $J_N(\cdot) \rightarrow J(\cdot)$, and the convergences are uniform. Further, $J_t(\nu(t)) = 0$.

A Lyapunov function

Theorem

$J_{N,t}(\cdot) \to J_t(\cdot)$, and $J_N(\cdot) \to J(\cdot)$, and the convergences are uniform. Further, $J_t(\nu(t)) = 0$.

Consequences:

A Lyapunov function

Theorem

$J_{N,t}(\cdot) \to J_t(\cdot)$, and $J_N(\cdot) \to J(\cdot)$, and the convergences are uniform. Further, $J_t(\nu(t)) = 0$.

Consequences:

- Thanks to uniform convergence, the prelimit Lyapunov function is as close to

$$\sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) J(\xi) - \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) J_t(\xi)$$

as one wishes, for sufficiently large N.

A Lyapunov function

Theorem

$J_{N,t}(\cdot) \to J_t(\cdot)$, and $J_N(\cdot) \to J(\cdot)$, and the convergences are uniform. Further, $J_t(\nu(t)) = 0$.

Consequences:

- Thanks to uniform convergence, the prelimit Lyapunov function is as close to

$$\sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) J(\xi) - \sum_{\xi \in \Delta_N(S)} p_N(t)(\xi) J_t(\xi)$$

  as one wishes, for sufficiently large N.

- Thanks to continuity of J_t, J, and $p_N(t) \xrightarrow{w} \delta_{\nu(t)}$, we get

$$J(\nu(t)) - J_t(\nu(t)) = J(\nu(t))$$

  is nonincreasing with t and thus a Lyapunov function for the MVE.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N\geq 1}$, is a Lyapunov function for the MVE.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N\geq 1}$, is a Lyapunov function for the MVE.

- Result already known (Budhiraja, Dupuis, Fischer, Ramanan 2015). But our proof (with L. Miclo) is simple and direct.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N\geq 1}$, is a Lyapunov function for the MVE.

- Result already known (Budhiraja, Dupuis, Fischer, Ramanan 2015). But our proof (with L. Miclo) is simple and direct.
- Earlier, we had established (with V. S. Borkar)
  - Large deviation principle (LDP) of the empirical measure in path space.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N \geq 1}$, is a Lyapunov function for the MVE.

- Result already known (Budhiraja, Dupuis, Fischer, Ramanan 2015). But our proof (with L. Miclo) is simple and direct.
- Earlier, we had established (with V. S. Borkar)
  - Large deviation principle (LDP) of the empirical measure in path space.
  - By contraction, LDP of the flow of empirical measure $(\eta_N(t), t \in [0, T])$. Rate function $S_{[0,T]}(\eta(\cdot))$. Rate function for state at terminal time J_T.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N\geq 1}$, is a Lyapunov function for the MVE.

- Result already known (Budhiraja, Dupuis, Fischer, Ramanan 2015). But our proof (with L. Miclo) is simple and direct.
- Earlier, we had established (with V. S. Borkar)
  - Large deviation principle (LDP) of the empirical measure in path space.
  - By contraction, LDP of the flow of empirical measure $(\eta_N(t), t \in [0, T])$. Rate function $S_{[0,T]}(\eta(\cdot))$. Rate function for state at terminal time J_T.
  - By techniques similar to the previous talk, LDP for the invariant measure p_N. Rate function J.

The main result

Theorem

The function $J(\cdot)$, the large deviation rate function of the sequence of invariant measures $(p_N)_{N\geq 1}$, is a Lyapunov function for the MVE.

- Result already known (Budhiraja, Dupuis, Fischer, Ramanan 2015). But our proof (with L. Miclo) is simple and direct.
- Earlier, we had established (with V. S. Borkar)
  - Large deviation principle (LDP) of the empirical measure in path space.
  - By contraction, LDP of the flow of empirical measure $(\eta_N(t), t \in [0, T])$. Rate function $S_{[0,T]}(\eta(\cdot))$. Rate function for state at terminal time J_T.
  - By techniques similar to the previous talk, LDP for the invariant measure p_N. Rate function J.
- While we could naturally establish that J_t and J are l.s.c., we had not established continuity of these functions.

Ingredients to show uniform convergence $J_N \to J$

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.

- $\mu_N(y) =$ probability that the configuration is y under the invariant measure.

- Let y and x be neighbours with x arising out of a single particle's legal transition. Then

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.

- $\mu_N(y) =$ probability that the configuration is y under the invariant measure.

- Let y and x be neighbours with x arising out of a single particle's legal transition. Then

  1. $\left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| \leq c \log N$.

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.

- $\mu_N(y) =$ probability that the configuration is y under the invariant measure.

- Let y and x be neighbours with x arising out of a single particle's legal transition. Then

  1. $\left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| \leq c \log N$.
  2. $\text{osc}(\log \mu_N(\cdot)) \leq cN$.

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.
- $\mu_N(y) =$ probability that the configuration is y under the invariant measure.
- Let y and x be neighbours with x arising out of a single particle's legal transition. Then
  1. $\left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| \leq c \log N$.
  2. $\mathrm{osc}(\log \mu_N(\cdot)) \leq cN$.
  3. $\left|\log \frac{p_N(\xi)}{p_N(\xi')}\right| \leq \left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| + c \log N$,
     if ξ and ξ' are the empirical measures of y and x, resp.

Ingredients to show uniform convergence $J_N \to J$

- Let $x = (x_n, 1 \leq n \leq N)$, a configuration of N particles.
  Let $y = (y_n, 1 \leq n \leq N)$, another neighbouring configuration.

- $\mu_N(y) =$ probability that the configuration is y under the invariant measure.

- Let y and x be neighbours with x arising out of a single particle's legal transition. Then

  1. $\left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| \leq c \log N$.
  2. $\text{osc}(\log \mu_N(\cdot)) \leq cN$.
  3. $\left|\log \frac{p_N(\xi)}{p_N(\xi')}\right| \leq \left|\log \frac{\mu_N(y)}{\mu_N(x)}\right| + c \log N$,
     if ξ and ξ' are the empirical measures of y and x, resp.
  4. $\text{osc}(\log p_N(\cdot)) \leq \text{osc}(\log \mu_N(\cdot)) + c \log N$.

- By Arzela-Ascoli, limit of $-\frac{1}{N} \log p_N(\cdot)$ (which we know exists from before) must be continuous.

J_t and J by inspection

- In the rest of the talk, a heuristic approach to write down the rate function by inspection.
- Formal proof takes a different route.

J_t and J by inspection

- In the rest of the talk, a heuristic approach to write down the rate function by inspection.
- Formal proof takes a different route.
  - First LDP for the empirical measure over finite horizons in path space, via Sanov's theorem for noninteracting particles, and then via an application of the Girsanov transformation.

J_t and J by inspection

- In the rest of the talk, a heuristic approach to write down the rate function by inspection.
- Formal proof takes a different route.
  - First LDP for the empirical measure over finite horizons in path space, via Sanov's theorem for noninteracting particles, and then via an application of the Girsanov transformation.
  - Then use the contraction principle to get to flows. Some additional work is needed to identify the form of the rate function.

J_t and J by inspection

- In the rest of the talk, a heuristic approach to write down the rate function by inspection.
- Formal proof takes a different route.
  - First LDP for the empirical measure over finite horizons in path space, via Sanov's theorem for noninteracting particles, and then via an application of the Girsanov transformation.
  - Then use the contraction principle to get to flows. Some additional work is needed to identify the form of the rate function.
  - Finally follow Freidlin-Wentzell to analyse the invariant measure.

The running cost of following a trajectory $\eta(\cdot)$

The running cost of following a trajectory $\eta(\cdot)$

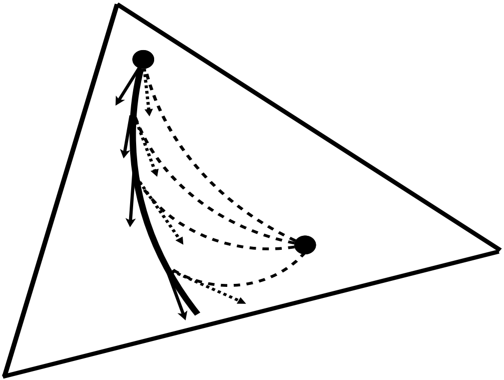

- At each time t, if the current state is $\eta(t)$, the natural tendency is to go along the tangent $L^*_{\eta(t)}\eta(t)$.

The running cost of following a trajectory $\eta(\cdot)$

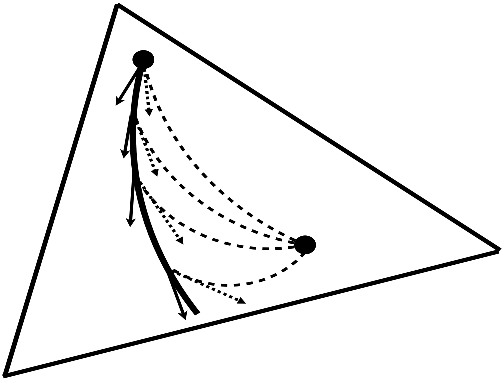

- At each time t, if the current state is $\eta(t)$, the natural tendency is to go along the tangent $L^*_{\eta(t)}\eta(t)$.
- To follow $\eta(t)$ however, the system needs to work against the McKean-Vlasov gradient and move along a different tangent $\dot{\eta}(t)$.

The running cost of following a trajectory $\eta(\cdot)$

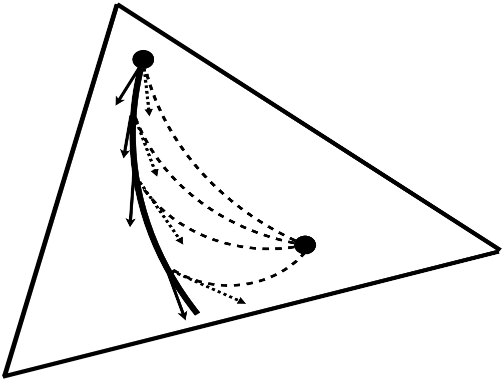

- At each time t, if the current state is $\eta(t)$, the natural tendency is to go along the tangent $L^*_{\eta(t)}\eta(t)$.
- To follow $\eta(t)$ however, the system needs to work against the McKean-Vlasov gradient and move along a different tangent $\dot{\eta}(t)$.
- This incurs a running cost $\mathcal{L}(\eta(t), \dot{\eta}(t))$.

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.
- Natural tendency for the $N\eta(t)(i)$ particles in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $L_{\eta(t)}(i, j)$.

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.
- Natural tendency for the $N\eta(t)(i)$ particles in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $L_{\eta(t)}(i, j)$.
- But to move along $\eta(t)$ they must have an instantaneous rate of $G_t(i, j)$.

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.
- Natural tendency for the $N\eta(t)(i)$ particles in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $L_{\eta(t)}(i,j)$.
- But to move along $\eta(t)$ they must have an instantaneous rate of $G_t(i,j)$.
- The $N\eta(t)(i)$ Bernoulli ($p = L_{\eta(t)}(i,j)\ dt$) random variables must have undergone a large deviation to exhibit an empirical measure close to ($q = G_t(i,j)\ dt$).

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.
- Natural tendency for the $N\eta(t)(i)$ particles in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $L_{\eta(t)}(i,j)$.
- But to move along $\eta(t)$ they must have an instantaneous rate of $G_t(i,j)$.
- The $N\eta(t)(i)$ Bernoulli $(p = L_{\eta(t)}(i,j)\ dt)$ random variables must have undergone a large deviation to exhibit an empirical measure close to $(q = G_t(i,j)\ dt)$. By Sanov's theorem, the negative exponent of the probability of this event is:

$$N\eta(t)(i)I(q||p) \cong N\eta(t)(i)(q \log \frac{q}{p} - q + p)$$

Guessing the running cost

- Write $\dot{\eta}(t) = G_t^* \eta(t)$.
- By decoupling, each particle's state is iid $\eta(t)$.
- Natural tendency for the $N\eta(t)(i)$ particles in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $L_{\eta(t)}(i,j)$.
- But to move along $\eta(t)$ they must have an instantaneous rate of $G_t(i,j)$.
- The $N\eta(t)(i)$ Bernoulli ($p = L_{\eta(t)}(i,j)\ dt$) random variables must have undergone a large deviation to exhibit an empirical measure close to ($q = G_t(i,j)\ dt$). By Sanov's theorem, the negative exponent of the probability of this event is:

$$N\eta(t)(i) I(q||p) \cong N\eta(t)(i)(q \log \frac{q}{p} - q + p)$$

- Sum over i and j and integrate over $[0, T]$ to get the action functional:

$$\int_0^T \mathcal{L}(\eta(t), \dot{\eta}(t))\ dt.$$

The case of a globally asymptotically stable equilibrium ν

Theorem

$J(\xi)$ is given by

$$J(\xi) = \inf \left\{ \int_0^T \mathcal{L}(\eta(t), \dot{\eta}(t))\, dt \mid \eta(0) = \nu, \eta(T) = \xi, T \in (0, \infty) \right\}.$$

- Any deviation that puts the system at ξ must have started its effort from ν.
- $J(\nu) = 0$.

The path to ξ

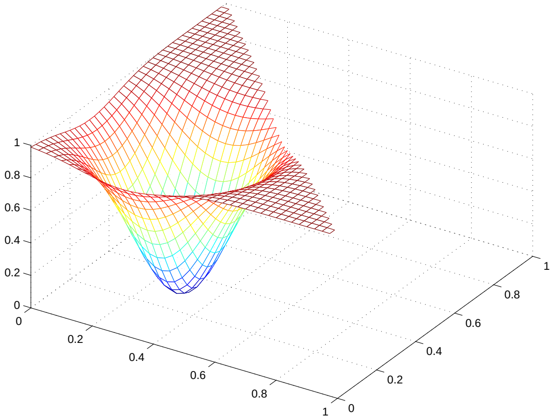

Can specify not only exponent $J(\xi)$ of the probability, but also the path.

Any deviation that puts the system near ξ must have started from ν, and must have taken the least cost path.

When there are multiple stable limit sets

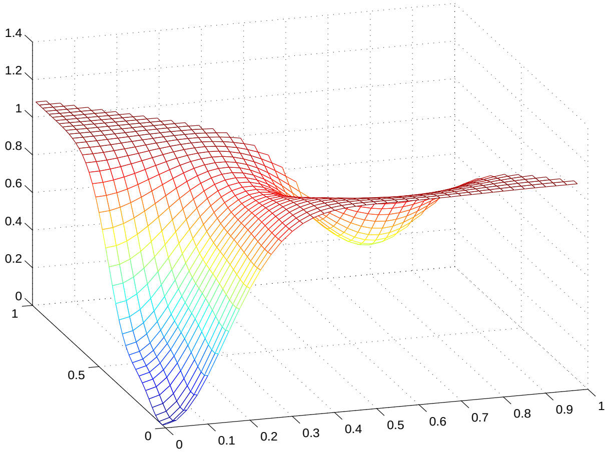

The case of two stable equilibria is easy to describe.

- V_{12} = cost of moving from $\nu^{(1)}$ to $\nu^{(2)}$.
- V_{21} = cost of the reverse move.
- If $V_{12} > V_{21}$, then $v_1 = 0$ and $v_2 = V_{12} - V_{21}$.

When there are multiple stable limit sets

Theorem

$J(\xi)$ is given by

$$J(\xi) = \inf \left\{ v_i + \int_0^T \mathcal{L}(\eta(t), \dot{\eta}(t))\, dt \mid \eta(0) = \nu^{(i)}, i = 1, 2, \eta(T) = \xi, T > 0 \right\}.$$

- Start from the global minimum $\nu^{(1)}$ and move to the attractor in the basin in which ξ lies along the least cost path.
- Then move to ξ from there along the least cost path.
- A similar result when there are more than two, but finite, number of stable limit sets.

Summary

Summary

- The large deviation rate function $J(\cdot)$ for the sequence of invariant measure $(p_N)_{N\geq 1}$ serves as a local Lyapunov function for the McKean-Vlasov dynamics.

Summary

- The large deviation rate function $J(\cdot)$ for the sequence of invariant measure $(p_N)_{N \geq 1}$ serves as a local Lyapunov function for the McKean-Vlasov dynamics.

- It is local Lyapunov if the dynamics has multiple stable rest points.