

Sampling-based optimization

Richard Combes

October 11, 2013

The topic of this lecture is a family of mathematical techniques called “sampling-based methods”. These methods are called “sampling methods” because they consist in finding a Markov process $\{X_n\}_{n \in \mathbb{N}}$ whose stationary distribution p allows to calculate a quantity of interest which is hard to calculate directly (say by a closed form formula, a summation or numerical integration). Three observations should be made:

- the quantity of interest should be a function of p ,
- $\{X_n\}$ should be easy to simulate, so that the quantity of interest can be estimated by simulating $\{X_n\}$ for long enough and observing its empirical distribution,
- p *need not be known* in general. The only requirement is that $\{X_n\}$ can be simulated. For instance knowing p up to a multiplicative constant (the normalization factor) is enough.

1 Markov Chain Monte Carlo

First we give a short introduction to Markov Chain Monte Carlo (MCMC) methods, which forms that basis for most sampling approaches. The interested reader might refer to [1] for a good introduction to the topic, and its applications to Bayesian statistics, machine learning and related fields.

1.1 The original problem: Maxwell-Boltzmann statistics

We give an introduction to MCMC by the first problem which motivated its development. The original problem was the calculation of Maxwell-Boltzmann statistics. This is an important problem in physics since Maxwell-Boltzmann statistics describe systems non-interacting particles (such as a perfect gas) in thermodynamical equilibrium.

We have a thermodynamical system whose state is denoted by s and $s \in \mathcal{S}$ a finite state space. It is noted that although \mathcal{S} is finite, it is generally very large. Namely, its dimension is the number of particles, so that $|\mathcal{S}|$ grows exponentially with the number of particles. To state s we associate its potential energy $E(s)$. We further define the temperature $T > 0$, and $k_B \approx 1.310^{-23} J.K^{-1}$ the Boltzmann constant. Both $E(\cdot)$ and T are assumed to be known. If the system stays at temperature T for a sufficient amount of time, it reaches thermodynamical

equilibrium, so that the probability $p(s)$ for the system to be at state s at a given time is given by the Boltzmann distribution:

$$p(s) = \frac{\exp\left(-\frac{E(s)}{T k_B}\right)}{Z(T)},$$

$$Z(T) = \sum_{s \in \mathcal{S}} \exp\left(-\frac{E(s)}{T k_B}\right).$$

with $Z(T)$ the normalizing constant also known as the partition function. If $|\mathcal{S}|$ is small, calculating $Z(T)$ can be done by summation over \mathcal{S} . However in most cases of interests, $|\mathcal{S}|$ is prohibitively large (exponential in the number of particles) so that direct summation is out of the question. An intuition why it might be possible to do better is that at low temperature, the only states with a significant contribution to $Z(T)$ are the states which have high probability i.e the states of *low energy*. Hence our attention should be restricted to those states.

The solution found by Metropolis ([6]) is to define a (reversible) Markov chain $\{X_n\}$ which admits p as a stationary distribution, and which is easy to simulate given the knowledge of the energy function $E(\cdot)$ and T, k_B . Before defining the update mechanism, it is noted that if $\{X_n\}$ is an ergodic Markov chain with stationary distribution, given any positive function $f : \mathcal{S} \rightarrow \mathbb{R}$, we have that (by the ergodic theorem for Markov chains):

$$\hat{f}_n = \frac{1}{n} \sum_{t=1}^n f(X_t) \xrightarrow{n \rightarrow \infty} \mathbb{E}_p[f] \quad \text{a.s.}$$

$$\mathbb{E}_p[f] = \sum_{s \in \mathcal{S}} p(s) f(s)$$

so that any function of p can be estimated as long as we are able to sample $\{X_n\}$ for sufficiently long.

Consider a set $N(s) \subset \mathcal{S}$ called the neighbors of s , typically states whose configuration is similar to s . For instance $N(s)$ could denote all the states obtained by taking s and changing the configuration of at most one particle. We assume symmetry: $s' \in N(s)$ iff $s \in N(s')$. The update equation considered by Metropolis is:

$$X_0 \in \mathcal{S}$$

$$Y_n \sim \text{Uniform}(N(X_n))$$

$$X_{n+1} = Y_n \text{ with probability } \min\left(e^{-\frac{E(Y_n) - E(X_n)}{k_B T}}, 1\right)$$

$$X_{n+1} = X_n \text{ otherwise.}$$

Namely, at time n , a neighbor Y_n of the current state X_n is chosen at random, and the chain moves to Y_n with probability $e^{-\frac{E(Y_n) - E(X_n)}{k_B T}}$. It is noted that if Y_n has higher probability (lower energy) than X_n , the chain always moves. To the contrary, if Y_n has a much lower probability (higher energy) than X_n , then the chain stays static with overwhelming probability. This is the reason why this approach is good compared with rejection sampling: most of the time $\{X_n\}$ dwells in regions of low energy.

We have that $\{X_n\}$ is a Markov chain, with transition probability ($s' \in N(s)$):

$$P(s, s') = \frac{\min(\exp(-\frac{E(s')-E(s)}{T k_B}), 1)}{|N(s)|},$$

and we can readily check that the detailed balance equations hold for all $s, s' \in N(s)$:

$$p(s)P(s, s') = p(s')P(s', s),$$

so that $\{X_n\}$ is a reversible Markov chain with stationary distribution p .

1.2 MCMC: sampling a distribution known up to a constant

While initially developed with the Boltzmann distribution in mind, the Metropolis-Hastings algorithm is naturally extended to sample from any distribution $p(\cdot)$ *known up to a constant*. In this lecture we restrict our attention to the case where \mathcal{S} is finite, but it is noted that the methods presented here can be extended in a natural way to uncountable and unbounded state spaces. The interested reader can refer to [1] for this case.

Consider $Q(\cdot, \cdot)$ a symmetric transition matrix on \mathcal{S} , and $R(\cdot, \cdot)$ an $|\mathcal{S}|$ by $|\mathcal{S}|$ matrix with positive entries. R must satisfy some constraints derived below. Matrix Q is often referred to as the “proposal distribution”, since it defines the candidate for the next value of the chain. We define a reversible Markov chain $\{X_n\}$ with stationary distribution p by the following update equation:

$$\begin{aligned} X_0 &\in \mathcal{S} \\ Y_n &\sim Q(X_n, \cdot) \\ X_{n+1} &= Y_n \text{ with probability } R(X_n, Y_n) \\ X_{n+1} &= X_n \text{ with probability } 1 - R(X_n, Y_n). \end{aligned}$$

Since we want $\{X_n\}$ to be reversible with distribution p , detailed balance must hold:

$$\begin{aligned} p(s)P(s, s') &= p(s')P(s', s), \\ p(s)Q(s, s')R(s, s') &= p(s')Q(s', s)R(s', s), \\ p(s)R(s, s') &= p(s')R(s', s), \end{aligned}$$

Therefore the acceptance probability must be chosen as:

$$R(s, s') = \begin{cases} 1 & \text{if } p(s') \geq p(s) \\ \frac{p(s')}{p(s)} & \text{otherwise.} \end{cases}$$

1.3 MCMC and mixing time

As seen above, for any proposal distribution Q , we can define a corresponding $\{X_n\}$ chain used for estimating the quantity of interest $\mathbb{E}_p[f]$ by the empirical average \hat{f}_n . The distribution

of \hat{f}_n depends on Q , and we would like choose Q such that \hat{f}_n is as close as possible to $\mathbb{E}_p[f]$. For instance we would like to make sure that the variance of \hat{f}_n is small. If $\{X_n\}$ was i.i.d, $\mathbf{var}(\hat{f}_n) = (1/n)(\mathbb{E}_p[f^2] - (\mathbb{E}_p[f])^2)$. The samples $\{X_n\}$ are in fact correlated so that $\mathbf{var}(\hat{f}_n)$ depends on the correlations $\mathbb{E}[f(X_n)f(X_{n'})]$. Roughly, we would like the samples $\{X_n\}$ to be *as de-correlated as possible*, so that the Markov chain $\{X_n\}$ goes through the state space quickly and has a small mixing time. The intuition for the choice of Q is the following: the possible moves (the neighborhood) of a state should be large enough to allow quick movement, but not include moves $s \rightarrow s'$ which have a small likelihood $p(s')/p(s)$, since these moves are rejected with high probability. Another way to put it is that if there are too few neighbors, it takes time to go through the state space, and if there are too many neighbours, then most proposed moves will be rejected and the chain will stay static most of the time.

It is noted that choosing a good proposal distribution Q is generally a difficult problem.

1.4 Gibbs Sampling

Gibbs sampling (proposed by [2]) is a version of the Metropolis-Hastings algorithm where the chain $\{X_n\}$ is updated one component at a time (or one block of components at a time). Gibbs sampling is also sometimes referred to as Glauber dynamics or the “heat bath”.

Going back to the example of the Boltmann distribution, consider the case where the system is made of K particles with 2 possible states: $\mathcal{S} = \{0, 1\}^K$, $k \leq K$. The state is $s = (s_1, \dots, s_K)$. We denote by s_{-k} the state without the k -th component, so that $s = (s_k, s_{-k})$. As said previously, the joint distribution p too complex since we cannot calculate the partition function. However the conditional distribution of s_k knowing s_{-k} is very simple, in fact it is a Bernoulli distribution, and:

$$p(s_k = 0 | s_{-k}) = \frac{e^{-\frac{E(0, s_{-k})}{k_B T}}}{e^{-\frac{E(0, s_{-k})}{k_B T}} + e^{-\frac{E(1, s_{-k})}{k_B T}}}.$$

The idea of Gibbs sampling is at time n to choose a component $k(n)$ at random, and update only this component. The Gibbs sampler is defined as:

$$\begin{aligned} X_0 &\in \mathcal{S} \\ k(n) &\sim \text{Uniform}(\{1, \dots, K\}) \\ Y_n &\sim p(\cdot | X_{n, -k(n)}) \\ X_{n+1, k(n)} &= Y_n \\ X_{n+1, k} &= X_{n, k} \text{ if } k \neq k(n) \end{aligned}$$

Once again, one can readily verify that the detailed balance conditions hold so that $\{X_n\}$ is a reversible Markov chain with stationary distribution p . Two remarks can be made:

- there is no rejection in Gibbs sampling, the proposed state is always accepted
- because of its component-by-component update, the Gibbs sampler lends itself to distributed updates (see for instance the corresponding literature on CSMA(Carrier Sense Medium Access) in wireless networks).

2 Simulated annealing

We now turn to simulated annealing which is a well-known sampling method for minimizing a function with possibly many local minima. Simulated annealing comes from an analogy with annealing in metallurgy. Annealing is a procedure which improves the crystalline structure of a metal by heating it up and then cooling it *slowly*. At low temperature, the Boltzmann distribution is concentrated on the states of minimal energy (also called “ground states”). The ground state here corresponds to a perfectly regular crystal. The metal might have an imperfect structure because it is trapped in a “glass state”: its configuration has high energy but the lack of thermal agitation prevents it from leaving that state and going to the ground state. The slow decay of temperature guarantees that at all time, the state of the system is close to thermodynamical equilibrium.

2.1 Intuition: a Boltzmann distribution with slowly decreasing temperature

We consider a cost function $V : \mathcal{S} \rightarrow \mathbb{R}^+$, with \mathcal{S} finite. We define $H = \{\arg \min_s V(s)\}$ the set of local minima. The goal is to minimize V . We define the distribution p (which is a Boltzmann distribution):

$$p(s, T) = \frac{\exp(-\frac{V(s)}{T})}{\sum_{s' \in \mathcal{S}} \exp(-\frac{V(s')}{T})}.$$

We see that p is concentrated on H for small temperatures:

$$p(H, T) \xrightarrow{T \rightarrow 0} 1.$$

Therefore it makes sense (at least intuitively) to sample from $p(\cdot, T)$ using the methods exposed above, while slowly decreasing the temperature.

2.2 Constant by parts schedules

The main question is to determine which “cooling schedules” ensure a.s convergence to the set of minima. In this lecture we consider a simple setting where the temperature is constant by parts, and we show the well known fact that the cooling schedule should be logarithmic. We restrict ourselves to this setting to be able to use elementary arguments in the proof. The interested reader should consult [3] for a more general setting and a stronger convergence result.

Temperature is decreased in a series of steps, with step $m \in \mathbb{N}$ taking place at time instants $\{t_m, \dots, t_{m+1} - 1\}$. We define α_m the length of the m -th step and T_m the corresponding temperature. It is noted that $t_m = \sum_{m' < m} \alpha_{m'}$. Finally we denote by $X^m = X_{t_m-1}$ the value of the chain at the end of step m .

Theorem 1 shows that providing that we choose a logarithmic cooling schedule $T_m = O(1/\log(m))$ and steps of polynomial length $\alpha_m = O(m^a)$ with a large enough, then the simulated annealing converges a.s. It is noted that Theorem 1 is weaker than the results of [3], since [3] proves the convergence of $\{X_n\}$ while we prove convergence of $\{X^m\}$ which is a *subsequence* of $\{X_n\}$.

Theorem 1. *There exists $a_0 > 0$ such that by choosing $T_m = \frac{\delta}{2 \log(m)}$, $\alpha_m = \lceil m^{a_0} \rceil$, $a \geq a_0$, the simulated annealing converges a.s. to the set of global minima:*

$$X^m \xrightarrow{m \rightarrow \infty} H, \text{ a.s.}$$

Without loss of generality, we assume that $\min_{s \in \mathcal{S}} V(s) = 0$, and we define $\delta = \min_{s \notin H} V(s)$, $V_\infty = \max_{s \in \mathcal{S}} V(s)$. We first prove that if α_m is large enough, then a.s. convergence occurs (Lemma 1). We then finish the proof by upper bounding the mixing time of the chain for a fixed value of the temperature T_m .

Lemma 1. *There exists a positive sequence $\{\beta_m\}$ such that if $\alpha_m \geq \beta_m$, and $T_m = \frac{\delta}{2 \log(m)}$, then:*

$$X^m \xrightarrow{m \rightarrow \infty} H, \text{ a.s.}$$

Proof of Lemma 1. Define the distribution of X^m : $p^m(s) = \mathbb{P}(X^m = s)$. We denote by $P(\cdot, \cdot, T)$ the transition matrix of the Markov chain $\{X_n\}$ when the temperature is constant equal to T .

Probability of not finding a minimum at the end of the m -th period:

For $\epsilon > 0$, we define the mixing time of $\{X_n\}$ when the temperature is constant and equal to T :

$$\tau(\epsilon, T) = \inf\{n : \sup_{s \in \mathcal{S}, p_0 \in \Delta(\mathcal{S})} |p(s, T) - p_0 P^n(\cdot, \cdot, T)| \leq \epsilon\},$$

with $\Delta(\mathcal{S})$ the set of distributions on \mathcal{S} . From discrete Markov chain theory, we know that for all $\epsilon > 0$ and $T > 0$, we have that the mixing time is finite $\tau(\epsilon, T) < \infty$ (this is a consequence of the Perron-Frobenius theorem). We will upper bound τ later.

We set $\beta_m = \tau(\exp(-\frac{\delta}{T_m}), T_m)$. Consider $s \notin H$, the probability to be in state s is upper bounded by:

$$\begin{aligned} p^m(s) &\leq p(s, T_m) + \exp(-\delta/T_m) \\ &= \frac{\exp(-V(s)/T_m)}{\sum_{s' \in \mathcal{S}} \exp(-V(s')/T_m)} + \exp(-\delta/T_m) \\ &\leq \exp(-\delta/T_m) + \exp(-\delta/T_m) = 2 \exp(-\delta/T_m). \end{aligned}$$

where we have used the fact that:

- $V(s) \geq \delta$ since $s \notin H$ and
- $\sum_{s' \in \mathcal{S}} \exp(-V(s')/T_m) \geq 1$, since $\min_s V(s) = 0$.

Summing over $s \notin H$:

$$p^m(\mathcal{S} \setminus H) \leq 2|\mathcal{S}| \exp(-\delta/T_m).$$

Almost sure convergence by the Borel-Cantelli theorem:

By definition of a.s convergence, denoting by ω a sample path:

$$\mathbb{P}[\{\omega : X^m(\omega) \not\rightarrow_{m \rightarrow \infty} H\}] = \mathbb{P}[\limsup_m \{\omega : X^m(\omega) \notin H\}].$$

Setting $T_m = \delta/(2\log(m))$:

$$\sum_{m \geq 1} p_m(\mathcal{S} \setminus H) \leq |\mathcal{S}| \sum_{m \geq 1} m^{-2} < \infty.$$

Hence applying the Borel-Cantelli theorem gives the announced result:

$$\mathbb{P}[X^m \not\rightarrow_{m \rightarrow \infty} H] = 0.$$

□

Proof of Theorem 1. From Lemma 1, and the definition of mixing time, it suffices to show that for large m :

$$\tau(\exp(-\delta/T_m), T_m) \leq O(m^a). \quad (1)$$

for $a \geq 0$ large enough. To bound the mixing time we use theorem 3 (presented in the appendix) and lower bound the conductance of the chain with temperature T_m .

Consider $\mathcal{S}' \subset \mathcal{S}$, $p(\mathcal{S}', T_m) \leq 1/2$. By assumption on the neighborhood function N , there exists a pair (s, s') , such that $s \in \mathcal{S}'$, $s' \in \mathcal{S} \setminus \mathcal{S}'$ and $s' \in N(s)$. Therefore the probability of observing a transition from s to s' is lower bounded by:

$$\begin{aligned} p(s)P(s, s') &\geq \frac{1}{|N(s)|} \exp\left(-\frac{V(s') - V(s)}{T_m}\right) \frac{\exp(-V(s)/T_m)}{\sum_{s'' \in \mathcal{S}} \exp(-V(s'')/T_m)} \\ &\geq \frac{\exp(-V(s')/T_m)}{|N(s)| \sum_{s'' \in \mathcal{S}} \exp(-V(s'')/T_m)} \\ &\geq \frac{\exp(-V(s')/T_m)}{|N(s)||\mathcal{S}|} \\ &\geq \frac{\exp(-V_\infty/T_m)}{|\mathcal{S}|^2}. \end{aligned}$$

where we have used the fact that $\sum_{s'' \in \mathcal{S}} \exp(-V(s'')/T_m) \leq |\mathcal{S}|$ since $\min_s V(s) = 0$, and the fact that $|N(s)| \leq |\mathcal{S}|$.

Therefore the conductance is lower bounded by:

$$\Phi \geq \frac{2 \exp(-V_\infty/T_m)}{|\mathcal{S}|^2}.$$

Also, we have that $\min_s p(s, T_m) \geq \exp(-V_\infty/T_m)/|\mathcal{S}|$.

Using theorem 3 with $\epsilon = \exp(-\delta/T_m)$ and the fact that $1/T_m = O(\log m)$, for $a \geq 0$ large enough

- $\Phi^{-2} = O(m^a)$
- $\log(1/p^*) = O(\log m)$
- $\log(1/\epsilon) = \delta/T_m = O(\log m)$

we have that $\tau(\exp(-\delta/T_m), T_m) = O(m^a \log(m))$ which concludes the proof. □

3 Payoff-based learning

In this section we describe a mechanism by which a group of agents can optimize a function without any information exchange [5]. The existence of such a mechanism and its simplicity are in fact rather surprising. This mechanism is a sampling method, where the state of agents follows a Markov chain. The rationale behind this mechanism comes from the theory of resistance trees for perturbed Markov chains introduced in [7].

3.1 The model

The model is the following. We consider N independent agents. Agent i chooses an action $a_i \in \mathcal{A}_i$ in a finite set and observes payoff $U_i(a_1, \dots, a_N) \in [0, 1)$. We can assume that payoffs are strictly smaller than 1 without loss of generality. The agents goal is to maximize the sum of payoffs $U(a) = \sum_{i=1}^N U_i(a)$. We define the set of maxima of U : $H = \arg \max_a U(a)$. This model is called “payoff-based learning” because the agents do not observe the payoffs or actions of the other players, and must select a decision based solely on their past actions and observed utilities. We require the assumption of “interdependence”: agents cannot be separated into two disjoint subsets that do not interact.

3.2 A sampling method

The sampling approach proposed by Peyton-Young is to define the agent’s “state” $(\bar{a}_i, \bar{u}_i, m_i)$, where $\bar{a}_i \in \mathcal{A}_i$ is a benchmark action (an action tried in the past), $\bar{u}_i \in [0, 1)$ is a benchmark payoff (a payoff previously yielded by the benchmark action), and a “mood” $m_i \in \{C, D\}$ (“Content”, “Discontent”). Roughly speaking, the “mood” represents whether or not the agent believes that he has found a good benchmark action. Agents will tend to become content when they observe high payoffs. They become discontent when they detect that their current action is not good or that other agents have changed. We consider a small experimentation rate $\epsilon > 0$ and a constant $c > N$.

The state at time n is $X_n = (\bar{a}_i(n), \bar{u}_i(n), m_i(n))_{1 \leq i \leq N}$. $\{X_n\}$ is a Markov chain and it is updated by the following mechanism:

If i is content ($m_i(n-1) = C$):

- Choose action $a_i(n)$ as:

$$\mathbb{P}[a_i(n) = a] = \begin{cases} \epsilon^c / (|\mathcal{A}_i| - 1) & a \neq \bar{a}_i(n-1) \\ 1 - \epsilon^c & a = \bar{a}_i(n-1) \end{cases}$$

- Observe resulting payoff $u_i(n)$:

- If $(a_i(n), u_i(n)) = (\bar{a}_i(n-1), \bar{u}_i(n-1))$, i stays content : $m_i(n) = C$
- If $(a_i(n), u_i(n)) \neq (\bar{a}_i(n-1), \bar{u}_i(n-1))$: i becomes discontent with probability $1 - \epsilon^{1-u_i(n)}$.

- Update benchmark actions $(\bar{a}_i(n), \bar{u}_i(n)) = (a_i(n), u_i(n))$

If i is discontent ($m_i(n-1) = D$):

- Choose action $a_i(n)$ as:

$$\mathbb{P}[a_i(n) = a] = 1/|\mathcal{A}_i|, \quad a \in \mathcal{A}_i$$

- Observe resulting $u_i(n)$, and become content with probability $\epsilon^{1-u_i(n)}$
- Update benchmark actions $(\bar{a}_i(n), \bar{u}_i(n)) = (a_i(n), u_i(n))$

The rationale behind this update is the following:

- Experiment (a lot) until content: When an agent is discontent, he plays an action at random, and becomes content only if this action yields high reward
- Do not change if content: A content agent remembers the (action, reward) pair that caused him to become content, and he keeps playing that same action with overwhelming probability
- Become discontent when others change: Whenever a content agent detects a change in reward he becomes discontent, because it indicates that another agent has deviated. This rule acts as a change detection mechanism.
- Experiment (a little) if content: Occasionally a content agent experiments, which is mandatory to avoid getting stuck in local minima

The result proven in [5] shows that for a small experimentation rate ϵ , the distribution of $\{X_n\}$ is concentrated on the set of maxima.

Theorem 2. Consider the Markov chain $\{X_n\}$, denote by $p(\cdot, \epsilon)$ its stationary distribution. Define

$$\tilde{H} = \{(\bar{u}, \bar{a}, m) : \bar{u}_i = U_i(\bar{a}), \bar{a} \in H, m_i = C, \forall i\}.$$

Then \tilde{H} is the only stochastically stable set so that:

$$p(\tilde{H}, \epsilon) \rightarrow 1, \epsilon \rightarrow 0^+.$$

The main difficulty of the proof is that the Markov chain $\{X_n\}$ is *not reversible*. In particular, unlike in simulated annealing for instance, the stationary distribution $p(\cdot, \epsilon)$ seems difficult to calculate in closed form. We do not go through the proof here, but a brief explanation of conductance trees for perturbed Markov chains is given in appendix, since it forms the basis for the proof.

A Mixing time of a reversible Markov chain

The following result gives an upper bound the mixing time of a *reversible* Markov chain based on a characteristic of the chain called “conductance”. The reason why conductance is useful is because in many cases conductance can be found by inspection of the transition matrix. This is in contrast with calculating the second eigenvalue of the transition matrix which is

often much harder. For a complete survey on mixing times in Markov chains, the interested reader should refer to [4].

We consider a reversible Markov chain $\{X_n\}_n$ on a finite set \mathcal{S} with transition matrix $P(.,.)$ and stationary measure $p(.,.)$. The ergodic flow between subsets $\mathcal{S}_1, \mathcal{S}_2$ is defined as:

$$\mathcal{Q}(\mathcal{S}_1, \mathcal{S}_2) = \sum_{s_1 \in \mathcal{S}_1} \sum_{s_2 \in \mathcal{S}_2} p(s)P(s, s'),$$

The conductance of the chain is:

$$\Phi = \min_{\mathcal{S}' \subset \mathcal{S}, p(\mathcal{S}') \leq 1/2} \frac{\mathcal{Q}(\mathcal{S}', \mathcal{S} \setminus \mathcal{S}')}{p(\mathcal{S}')}.$$

Simply said, the ergodic flow $\mathcal{Q}(\mathcal{S}_1, \mathcal{S}_2)$ is the probability of observing a transition from \mathcal{S}_1 to \mathcal{S}_2 , and the conductance is the minimum over all \mathcal{S}' of the probability of observing a transition leaving \mathcal{S}' , knowing that the chain is currently in \mathcal{S}' . The conductance is small whenever there is a particular subset of the state space which is difficult to escape, so that intuitively the mixing time should be a decreasing function of the conductance.

We recall that the mixing time is defined as:

$$\tau(\epsilon) = \inf\{n : \sup_{s \in \mathcal{S}, p_0 \in \Delta(\mathcal{S})} |p(s) - p_0 P^n(.,.)| \leq \epsilon\}.$$

Theorem 3. *With the above definitions, and $p^* = \min_s p(s)$, the following bound on the mixing time holds:*

$$\tau(\epsilon) \leq \frac{2}{\Phi^2} (\log(1/p^*) + \log(1/\epsilon)).$$

B Stochastic Potential

B.1 The main theorem

We consider a family of Markov chains $\{X_n^\epsilon\}_{n,\epsilon}$ on a finite set \mathcal{S} with transition matrices $P^\epsilon(.,.)$ and stationary measures p^ϵ . Once again we do not assume reversibility. Assume that there are constants (which do not depend on ϵ), $r(s, s') \geq 0$ such that $P^\epsilon(s, s') \sim_{\epsilon \rightarrow 0^+} \epsilon^{r(s, s')}$. By analogy with electric circuits, $r(s, s')$ is called the resistance of link (s, s') . We define E_1, \dots, E_M the recurrence classes of $P(.,., 0)$. Consider γ_{ij} a path between class E_i and E_j , namely $\xi_{ij} = \{(s_1, s_2), \dots, (s_{b-1}, s_b)\}$, $s_1 \in E_i$ and $s_b \in E_j$. We define the resistance of a path as the sum of resistances of its links

$$r(\xi) = r(s_1, s_2) + \dots + r(s_{b-1}, s_b).$$

We define the minimal resistance $\rho_{ij} = \min r(\xi_{ij})$ where the minimum is taken on all possible paths from E_i to E_j . Consider a weighted graph \mathcal{G} whose vertices represent the recurrence classes $(E_i)_{1 \leq i \leq M}$ with weights $(\rho_{ij})_{1 \leq i, j \leq M}$. Fix i , and consider a directed tree \mathcal{T} on \mathcal{G} which contains exactly one path from j to i (for all $j \neq i$). The stochastic potential ϕ_i of class E_i is the minimum of $\sum_{(i,j) \in \mathcal{T}} \rho_{i,j}$, where the minimum is taken over all possible trees \mathcal{T} .

The following result [7] states that the distribution $p(., \epsilon)$ is concentrated on the sets with minimal stochastic potential.

Theorem 4. *The only stochastically stable recurrence classes E_1, \dots, E_M are the ones with minimum stochastic potential. Namely, if $\phi_i \neq \min_{i'} \phi_{i'}$, then $p(E_i, \epsilon) \rightarrow_{\epsilon \rightarrow 0} 0$.*

B.2 An example

As an illustrative example, consider the case where there are $I + 1$ possible states $\{0, \dots, I\}$, with $P(i, j, \epsilon) = \epsilon^{r(i,j)}$, and $r(i, j) = 0$ if $i \neq 0$ and $j \neq 0$. Namely we are considering a star network with $I + 1$ nodes where node 0 is the central node. It is noted that in this case we do not need the theory of resistance trees since the chain is reversible and its stationary distribution can be calculated in closed form, but we provide this example for the sake of illustrating the definitions. Then:

- The recurrence classes are singletons: $E_i = \{i\}$, $i \in \{0, \dots, I\}$
- The minimal resistances are (by inspection): $\rho_{0i} = r(0, i)$, $\rho_{i0} = r(i, 0)$, $\rho_{ij} = r(i, 0) + r(0, j)$ for $i \neq j$, $i \neq 0$, $j \neq 0$.
- The tree of minimal resistance rooted at 0 is $\{(1, 0), (2, 0), \dots, (I, 0)\}$, and the tree of minimal resistance rooted at $i \neq 0$ is $\{(0, i), (1, 0), (2, 0), \dots, (I, 0)\}$ (the link $(i, 0)$ is not included).

Therefore, the stochastic potential of class j is (using the convention $r(0, 0) = 0$):

$$\phi_j = r(0, j) - r(j, 0) + \sum_{1 \leq i \leq I} r(i, 0).$$

Therefore, the stochastically stable set of nodes is given by:

$$\{\arg \min_j (r(0, j) - r(j, 0))\},$$

which is the net flow through the link $j \leftrightarrow i$ in the electrical network analogy. One can readily derive the same result by calculating the stationary distribution explicitly, since detailed balance holds.

References

- [1] Christophe Andrieu. An introduction to mcmc for machine learning. *Kluwer Academic Publishers*, 2003.
- [2] S. Geman and D Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1984.
- [3] B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 1988.
- [4] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. AMS, 2009.
- [5] J.R. Marden, H. Peyton Young, and L.Y. Pao. Achieving pareto optimality through distributed learning. In *Proc. of CDC*, 2012.

- [6] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 1953.
- [7] H. Peyton Young. The evolution of conventions. *Econometrica*, 1993.