Lecture notes on Stopping Times

Aurko Roy, Ben Cousins

November 17, 2014

1 Introduction

Markov chains are awesome and can solve some important computational problems. A classic example is computing the volume of a convex body, where Markov chains and random sampling provide the only known polynomial time algorithm. Arguably the most important question about a Markov chain is how long to run it until it converges to its stationary distribution. Generally, a chain is run for some fixed number of steps, until the current point is provably within some threshold distance from the target distribution.

However, this ignores the information we gain as the chain takes its steps, which could potentially be beneficial for determining convergence. For instance, we might converge much faster than the mixing time suggests, or perhaps the mixing time is unknown, in which case we have no a priori way to know how long to run the Markov chain. So instead, perhaps we observe some quantities of the walk, and then announce convergence once some condition is met. Such algorithms which, at a high level, use the walk to determine convergence are referred to as stopping rules.

Also note that we can use stopping rules to create distributions other than the stationary distribution of the Markov chain. A trivial example of such a stopping rule is to stop the walk when we reach a current state x. We will see other more interesting stopping rules that can generate complex distributions over the state space.

2 Preliminaries

We assume that M is an irreducible Markov chain with transition matrix $M = \{p_{ij}\}$ and initial distribution σ over a state space Ω such that $|\Omega| = n$. We denote by Ω^* the set of finite "walks" on Ω, i.e. the set of strings $w = (w_0, w_1, \ldots, w_t)$ where $w_i \in \Omega$. Given σ and M, the set Ω^* inherits the following probability distribution.

$$\Pr(w) = \sigma_{w_0} \prod_{i=1}^{t} p_{w_{i-1}w_i}$$

Definition 1 (Stopping Rule)**.** *A stopping rule Γ is a partial map from the set Ω^* to $[0,1]$. It is defined precisely when $\Pr(w) > 0$ and $\Gamma(w_0, \ldots, w_i)$ is defined and non-zero for each $0 \leq i \leq t-1$.*

Intuitively the stopping rule gives us for every finite walk w attainable by the Markov chain, the probability of continuing the walk. One can also think of Γ as a random variable taking values in $\mathbb{Z}_+$, such that the chain stops at w_Γ. We will use the two notions interchangeably; their meaning will become clear from the context.

The *mean length* $\mathbb{E}\Gamma$ of a stopping rule Γ is the expected number of steps taken by the chain before it stops. Observe that if $\mathbb{E}\Gamma < \infty$ then with probability 1 the chain eventually stops. This can be seen for instance by defining the sequence of random variables $\{X_t\}_{t\in\mathbb{Z}_+}$ where $X_t = 1$ if the chain has not stopped at time t and 0 otherwise. Then clearly $\Gamma = \sum_{t=0}^{\infty} X_t$ and so $\mathbb{E}\Gamma = \sum_{t=0}^{\infty} \mathbb{E}X_t = \sum_{t=0}^{\infty} \Pr(X_t > 0) < \infty$ and so by the first Borel-Cantelli lemma the event that $X_t > 0$ happens infinitely often is 0. In this case we get a probability distribution σ^Γ over Ω

$$\sigma_j^\Gamma := \sum_{w=(w_0,\ldots,w_t=j)} \sigma_{w_0} \left(\prod_{i=1}^{t} \Gamma(w_0,\ldots,w_{i-1}) \right) (1 - \Gamma(w))$$

where σ_j^Γ is the probability of stopping in state j. Given a distribution τ over Ω, we say that Γ is a stopping rule from σ to τ if $\sigma^\Gamma = \tau$. We first make the observation that for any pair of distributions σ, τ over Ω, there is a stopping rule Γ_τ (called the "naive" rule) such that $\sigma^\Gamma = \tau$. The "naive" rule selects a state $j \in \Omega$ drawn according to τ and stops when the chain reaches j. The mean length of the "naive" rule is

$$\mathbb{E}\Gamma_\tau = \sum_{i,j\in\Omega} \sigma_i \tau_j H(i,j)$$

where $H(i,j)$ is the hitting time from i to j. The *maximum length* $\max(\Gamma)$ of a stopping rule Γ is defined as the length of the longest walk w that has positive probability according to Γ. Note that the probability of a walk $w = (w_0, \ldots, w_t)$ according to Γ is given by

$$\Pr_\Gamma(w) = \sigma_{w_0} \prod_{i=1}^{t} \Gamma(w_0,\ldots,w_i) p_{w_{i-1},w_i}.$$

Definition 2 (Mean Optimal Stopping Rule)**.** *A stopping rule Γ from σ to τ is called* mean optimal *(for σ and τ) if for any stopping rule Γ' from σ to τ, $\mathbb{E}\Gamma \le \mathbb{E}\Gamma'$. If Γ is mean optimal for σ, τ then $H(\sigma,\tau) := \mathbb{E}\Gamma$ is called the* access time *from σ to τ.*

Definition 3 (Max Optimal Stopping Rule)**.** *A stopping rule Γ from σ to τ is called* max optimal *(for σ and τ) if for any stopping rule Γ' from σ to τ, $\max(\Gamma) \le \max(\Gamma')$.*

Definition 4 (Exit frequencies)**.** *The* exit frequencies *$x = \{x_j\}_{j\in\Omega}$ is defined by setting x_j equal to the expected number of times the walk leaves state j before stopping. The exit frequencies of the naive stopping rule is denoted by $\{\tilde{x}_j\}_{j\in\Omega}$.*

We have the following lemma about exit frequencies of stopping rules from σ to τ.

Lemma 1. *The exit frequencies of any stopping rule from σ to τ satisfy*

$$\sum_i p_{ij} x_i - x_j = \tau_j - \sigma_j$$

A simple corollary of the above lemma is the following theorem.

Theorem 1. *Let Γ and Γ' be two stopping rules with exit frequencies x and x' respectively. Let $D = \mathbb{E}\Gamma - \mathbb{E}\Gamma'$ be the difference between their mean lengths. Then $\sigma^{\Gamma} = \sigma^{\Gamma'}$ if and only if $x' - x = D\pi$, where π is the stationary distribution of M.*

In the next section we talk about a few mean optimal stopping rules that hold for any Markov chain. The first of these called the *filling rule* also gives us a useful characterization of mean optimal rules in terms of exit frequencies.

3 Examples

3.1 Random walk on a hypercube

Consider the following random walk on the hypercube $\{0,1\}^n$ where we select a direction uniformly at random and flip the bit in that direction. Then a stopping rule that generates the uniform distribution on $\{0,1\}^n$ is as follows: stop when you have selected every direction. By the coupon collector argument the mean length of this rule is $O(n \log n)$.

3.2 Top card shuffling

Consider the following shuffling algorithm: select the top card from the deck and insert it with equal probability into any of the n slots among the remaining $n-1$ cards. Then a stopping rule that generates the uniform distribution is to perform one more shuffle when the card that was originally at the bottom of the deck reaches the top and stop. Again by the coupon collector argument, the mean length of this stopping rule is $O(n \log n)$.

3.3 Random walk on a cycle

Consider the following stopping rule for the random walk on a cycle starting from a vertex u: stop when every node has been visited once. Somewhat counterintuitively, the probability that v is the last node visited is the same for every $v \neq u$. Thus this also gives us a stopping rule for generating the uniform distribution on the cycle.

4 Mean Optimal Stopping Rules

4.1 Filling Rule

Definition 5. *The filling rule $\Phi_{\sigma,\tau}$ is defined recursively as follows. Suppose we have defined $\Phi_{\sigma,\tau}(w)$ for every w such that $|w| \leq k$. Let p_i^k be the probability* *according to $\Phi_{\sigma,\tau}$ of being in state i after k steps and let q_i^k be the probability of stopping at state i in fewer than k steps. Then if we are in state i in the $k+1^{th}$ step, then we stop with probability* $\min(1, (\tau_i - q_i^k)/p_i^k)$.

The filling rule gives us the following useful characterization of mean optimal stopping rules.

Theorem 2. *A stopping rule Γ is mean optimal for σ, τ if and only if there is a $j \in \Omega$ such that $x_j = 0$.*

Proof Sketch. From Theorem 1 it follows that if $x_j = 0$ for a stopping rule Γ then Γ must be mean optimal. This is because for any Γ' from σ to τ, $x_j - x'_j = -x'_j = (\mathbb{E}\Gamma - \mathbb{E}\Gamma')\pi_j$ and since $x'_j, \pi_j \geq 0$, it follows that $\mathbb{E}\Gamma \leq \mathbb{E}\Gamma'$. For the converse direction we need to exhibit a mean optimal stopping rule that has a state j such that $x_j = 0$. The filling rule is exactly the rule with the desired property. □

4.1.1 Understanding the Filling Rule

Let's consider a small toy example to try and understand the filling rule. Consider the graph in Figure 1 and the Markov chain which randomly selects an outgoing edge with equal probability. Note that this chain will have stationary distribution $(1/2, 1/2)$ and will converge after a single step. Let σ_u and τ_u denote the starting and target distributions, respectively, for the vertex u.

We now run the filling rule with an example starting and target distribution. Suppose $\sigma_u = 1, \sigma_v = 0$ and $\tau_u = 2/3, \tau_v = 1/3$. Note that whenever the Markov chain has nonzero probability at a state, the filling rule will assign as much probability as it can to that state without "overfilling", i.e. exceeding the target probability for that state. So, in the first step, the filling rule will stop at vertex u with probabilty $2/3$ and will take another step of the Markov chain with probability $1/3$, i.e. will be at u or v each with probability $1/6$. Note that vertex u has "filled" up its probability to $2/3$, so the filling rule will no longer stop at state u. Note that v is a halting state, as every time the we visit v for this σ, τ we always stop (and so by the previous theorem, we see that the filling rule is mean optimal here).

So how many steps do we need to take, in expectation, to generate τ starting from σ? With probability $2/3$, we stop after 0 steps (at u); with probability $1/3$, we will stop once we reach v. This will be a geometric distribution with probability $1/2$. So, in expectation, we stop after $(1/3) \cdot 2 = 2/3$ steps.

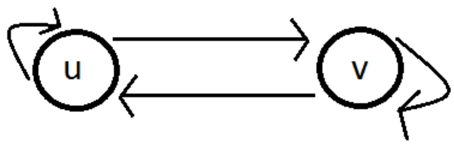

Figure 1: A simple graph.

4.2 Local Rule

Definition 6. *Let τ be arbitrary and let x_i be the exit frequencies for a mean optimal stopping rule from σ to τ. Then if we are at state i then the chain stops with probability $\tau_i/(\tau_i + x_i)$. Note that the stopping rule depends only on the state i and not on the time.*

Theorem 3. *The local rule generates τ, i.e $\sigma^{\Lambda_\tau} = \tau$ and Λ_τ is mean optimal.*

4.3 Chain Rule

Let ρ be a distribution on the subsets U of Ω. Then ρ gives us the following natural stopping rule: choose a set U from Ω according to ρ and keep walking till the chain hits a state in U. Note that the naive rule is a special example of such a rule where ρ is concentrated only on singletons. The chain rule is also a special case of this rule.

Definition 7. *For every pair of distributions σ, τ on Ω, there exists a unique distribution ρ which is concentrated on a chain of subsets of Ω and gives a mean optimal stopping rule from σ to τ. Then the stopping rule obtained as above from ρ is called a* chain rule *from σ to τ.*

4.4 Threshold Rule

Definition 8. *A stopping rule Γ is said to be a threshold rule if there is a threshold vector $h = (h_1, \ldots, h_n), h_i \geq 0$ such that*

$$\Gamma(w_0, \ldots, w_k) = \begin{cases} 0 & \text{if } k \geq h_{w_k} \\ 1 & \text{if } k \leq h_{w_k} - 1 \\ k - h_{w_k} & \text{otherwise} \end{cases}$$

Essentially, a threshold rule is a set of critical times for each state, and if we reach a state after its critical time, we stop. For instance, we could set each state's threshold to be the mixing time of the lazy chain to get a point approximately from the stationary distribution of the chain.

Theorem 4. *For any target distribution τ, there exists a threshold rule which is both mean-optimal and max-optimal.*

In fact, there is an algorithm which, given a target distribution τ, will construct a threshold vector as the walk runs, and runs in time polynomial in the size of the state space.

4.4.1 Understanding the Threshold Rule

Consider the same example graph of Figure 1. The filling rule assigns as much probability to a state as it can in the current step; the threshold rule, however, will wait until a critical time for the state is reached, and then always assign probability when that state is visited. Suppose $\sigma_u = 1$, $\sigma_v = 0$, and $\tau_u = \tau_v = 1/2$. The threshold rule for this example will simply take a single step and then stop. Note that the filling rule and the threshold rule for this example will have the sample number of expected steps, but the threshold rule will always run in a single step. This behavior of the threshold rule is what makes it max optimal. Note that not every threshold rule will be max optimal–the previous theorem simply guarantees the existence of one.

5 Exact Mixing in an Unknown Chain

Suppose that we didn't know the mixing time of a Markov chain, but wanted to generate a point from the stationary distribution π of the chain. It turns out, perhaps suprisingly, that there is a stopping rule that will generate points *exactly* from π. Let M be an irreducible Markov chain with state space Ω, and let $h = \max_{i,j \in \Omega} H(i,j)$ denote the maximum hitting time of the chain.

We first give some definitions. Given a state i and positive integer t, a *t-exit* of i waits X steps after being in state i, where X is chosen uniformly from $\{0, 1, \ldots, t-1\}$, then setting i' equal to the resulting state. A *t-pass* selects a state j uniformly at random from Ω, then runs independent t-exits from every state i except j. A t-pass is deemed *successful* if the directed graph consisting of the $n-1$ arcs (i, i') has no loop or cycle (and is thus an in-directed spanning tree rooted at j).

The stopping rule Γ sets $t = 2$, performs $3|\Omega|$ t-passes, and sees if any of the t-passes are successful. If so, Γ takes one step from the state j, which was the root of the successful pass, and stops. Otherwise, we double t and repeat.

Theorem 5. *Γ runs in expected number of steps bounded by $O(h^3 \log h)$ and stops at state j with probability exactly π_j.*

Also, note that any stopping rule for an unknown Markov chain that generates the stationary distribution never stops before it visits all states. Therefore, we cannot hope to do better than the hitting time h.

6 Sampling from a Convex Body

All of the previous discussion has been on stopping rules for a Markov chain on a discrete state space. But what about a continuous state space? One application that we've come across for a stopping rule is the problem of generating random samples from a convex body. There are a number of random walks which can provable converge in polynomial time to the uniform distribution over a convex body. However, the bounds given on the mixing time are generally too high to be practical, and therefore if one were to implement one of these random walks, it would be far too slow to run the Markov chain for the proven mixing time.

An implementation of these walks seems to suggest that convergence happens faster than the proven mixing time [1]. So how should we determine convergence of the walk? The implementation uses some heuristic tests, such as looking at the proportion of points that lie on one side of a random halfspace. But none of them provide provable guarantees of accuracy. An interesting open question is to develop an efficient stopping rule that would determine when the current point is approximately random. Or in the case of when these points are used for estimating the volume of a convex body, could we develop some set of tests that decides when the current stream of points will provide a good estimate for the volume?

References

[1] Cousins, B. and Vempala, S., “Volume Computation of Convex Bodies”, MATLAB File Exchange, 2013, http://www.mathworks.com/matlabcentral/fileexchange/43596-volume-computation-of-convex-bodies

[2] Lovasz, L. and Winkler, P., “Fast Mixing in a Markov Chain”, 2007

[3] Lovasz, L. and Winkler, P., “Efficient stopping rules for Markov Chains”, 1995