M346 (55820), 03/07/2012

Ravi Srinivasan (rav@math.utexas.edu)

Markov chains: an example

Introduction

Consider the following example of a Markov chain. Let the state space be $S = \{1, ..., 100\}$. The transition probabilities are all of the following form: given that we're at state i, we have probability $0 < p < 1$ of moving to state $i + 1$ at the next step and probability $q = 1 - p$ of moving to $i - 1$. At the boundary states 1 and 100, we stay at 1 with probability q and we stay at 100 with probability p. In other words, each state only communicates with its immediate neighbors (a particle sitting at a state can only move to a neighboring state in one time step). In the following, it is useful to think of the chain as a graph of nodes $\{1, ..., 100\}$ with neighboring nodes connected to each other by two directed arrows, except for the boundary nodes 1 and 100 which connect to themselves:

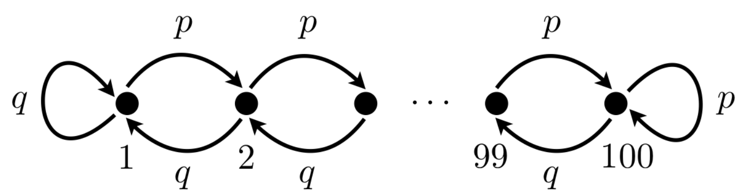

Figure 1. State space $S = \{1, ..., 100\}$ of the Markov chain and corresponding transition probabilities.

This Markov chain, called a (finite-state) random walk, is one of the most well-studied and useful models in all of probability theory. In the limit of larger state spaces and smaller time steps, the model converges to what is known as Brownian motion. Random paths generated from this Markov chain have been used to study an enormous variety of phenomena: molecular diffusion, genetic drift, options pricing in finance, estimating sizes of large networks such as the WWW, to name just a few... look it up on Wikipedia!

Transition matrix

Let's first look at the transition matrix A of the Markov chain. The entry A_{ij} corresponds to the probability of moving to state i in one time step given that we start at state j. Then, A has a string of p's below the diagonal and a string of q's above the diagonal, $A_{1,1} = q$ and $A_{100,100} = p$, and all other entries are 0:

$$A = \begin{pmatrix} q & q & & & & \\ p & 0 & q & & & \\ & p & 0 & \ddots & & \\ & & \ddots & \ddots & q & \\ & & & p & 0 & q \\ & & & & p & p \end{pmatrix}.$$

We can confirm that A is **irreducible**—for any states i, $j \in S$ it is possible to get to state i starting from state j if we wait long enough. A is also **aperiodic**—for any state it is possible to return to it at *any* time larger than some minimum waiting period. In other words, A is **regular** (i.e., it is both irreducible and aperiodic) which can be confirmed directly by observing that for K large enough, A^K has all its entries strictly positive—i.e., $\left(A^K\right)_{ij} > 0$ for all $i, j \in S$.

Analysis of the model

Let's look at three particular versions of this Markov chain, corresponding to different choices of the transition probability p.

- Model 1: $p = 0.525$
- Model 2: $p = 0.5$
- Model 3: $p = 0.475$

For concreteness, Models 1-3 can be thought of as a model of stock prices in a bull market (prices more likely to go up than down), neutral market, or a bear market (prices more likely to go down than up), respectively. Since we work on a finite state space, we have upper and lower boundaries on prices. We will start each of these models with initial probability vector $\boldsymbol{x}(0) = \boldsymbol{e}_{48}$ (that is, every random walker begins at state 48).

Ensemble of paths

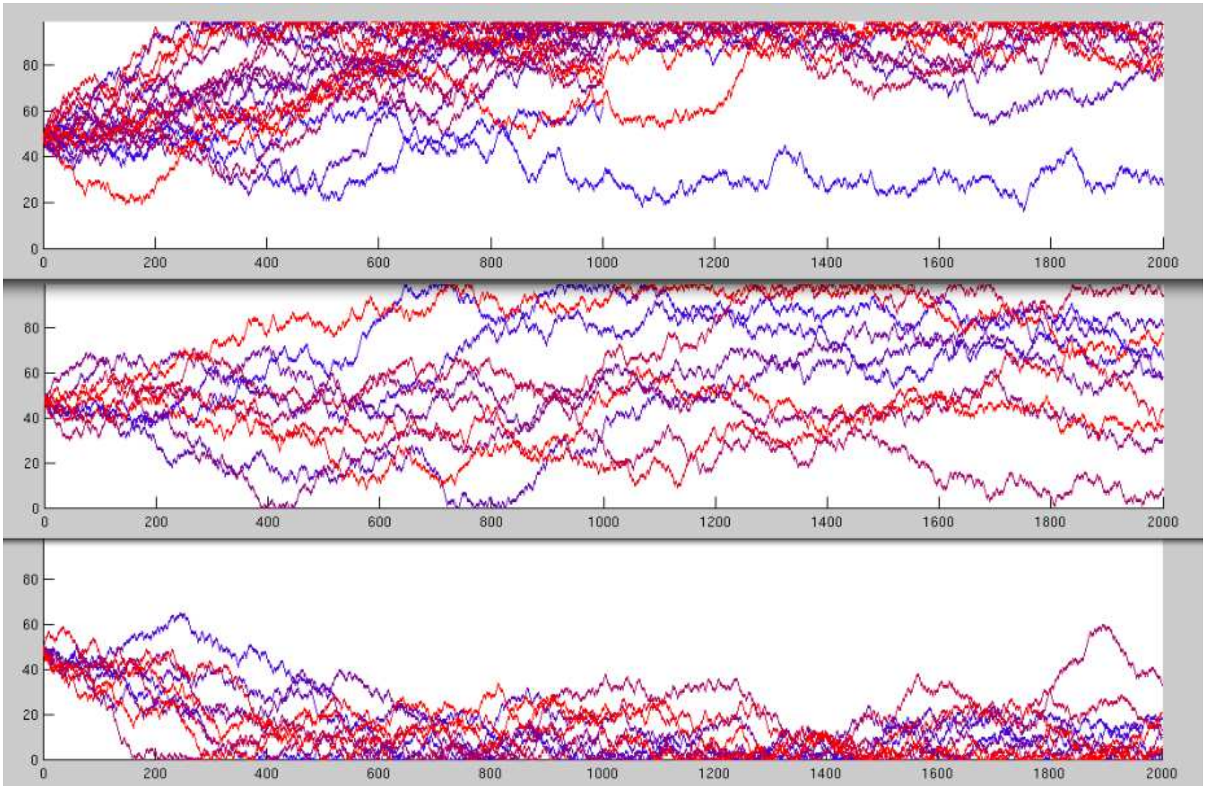

Figure 2. Ensembles of paths for Models 1-3. x-axis is time $k = 0, \ldots, 2000$, y-axis is state space $S = 1, \ldots 100$.

Figure 2 shows 10 sample paths generated from Models 1-3. Each path $\{Z(k): k = 0, \ldots, 2000\}$ is generated by starting a random walker at state 48, then randomly moving up or down with probability p or q, respectively, at every time step. These paths are only a subset of the set of all possible paths (typically referred to as an **ensemble**) that could have generated by the models.

Evolution of probability distribution

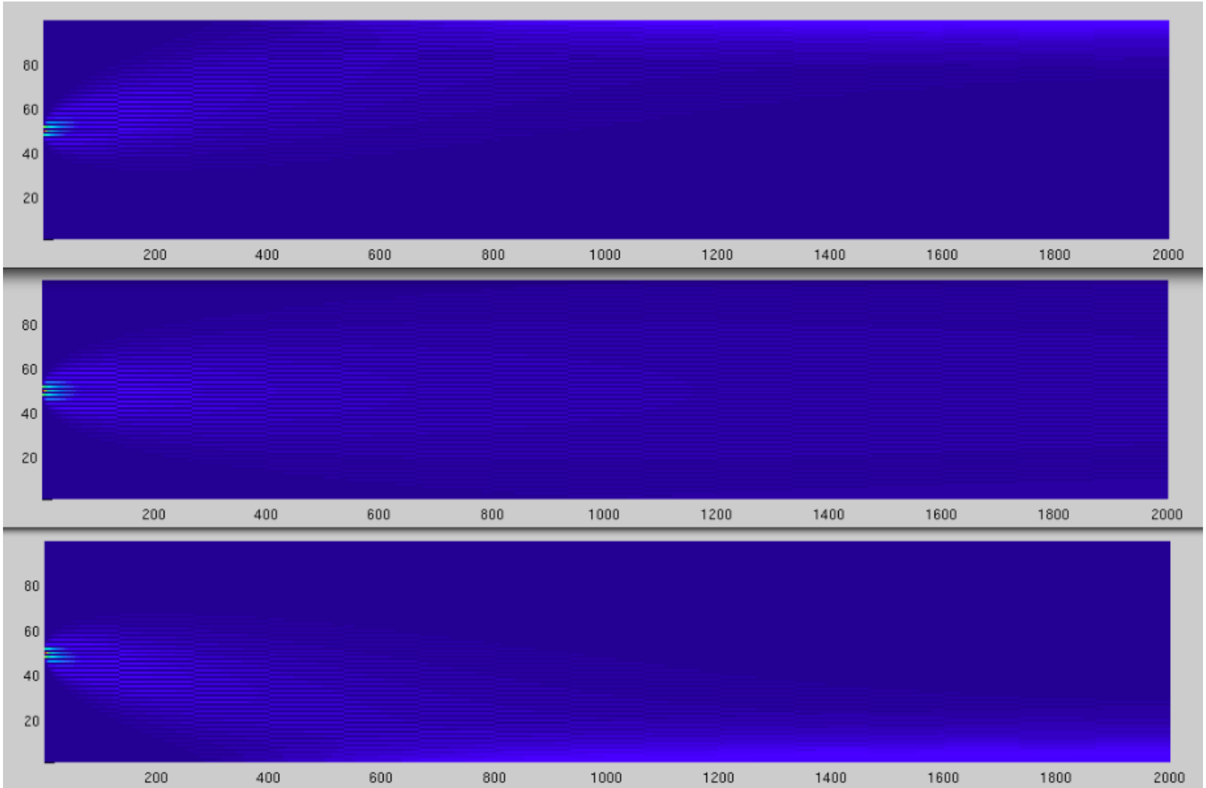

Figure 3. Distribution $\boldsymbol{x}(k) = A^k\boldsymbol{x}(0)$ for Models 1-3. x-axis is time $k = 0, ..., 2000$, y-axis is state space $S = 1, ...100$. Brighter colors correspond to larger values of $\boldsymbol{x}(k)$.

Figure 3 shows a colormap of the probability vector $\boldsymbol{x}(k) = A^k\boldsymbol{x}(0)$ at each time $k = 0, ..., 2000$. Brighter colors correspond to larger values. The distribution $\boldsymbol{x}(k)$ can be interpreted in one of two equivalent ways: (1) as the *probability that any one walker* sits at a state at time k, or (2) as the *proportion of all walkers* which sit at a given state at time k. In your homework problems, $\boldsymbol{x}(k)$ is what you're usually asked to analyze or compute.

Note the upward drift in time for the values of $\boldsymbol{x}(k)$ in Model 1 and the downward drift for $\boldsymbol{x}(k)$ in Model 3, as should be expected. This agrees with the previous figure in which we see that most paths in Model 1 will drift upwards initially while those in Model 3 will tend to drift downwards initially, before settling down to a 'statistical equilibrium' due to the upper and lower boundaries. In Model 2, paths tend to spread out symmetrically (there is no mechanism that favors upward or downward movement). From Figure 2 we see that no individual random walker ever stops moving around—it is the *distribution* $\boldsymbol{x}(k)$ of walkers that converges to some stationary distribution $\boldsymbol{\pi}$. The next figure also confirms this.

At this point, here are some good questions to ask yourself based on the figures above:

1. How long does it look like it takes before $\boldsymbol{x}(k)$ is reasonably close to a stationary distribution $\boldsymbol{\pi}$?

2. If we looked at any one particular path after this given time what would we expect it to look like?

3. If we wanted to use the Markov chain to generate a random number from $\{1, ..., 100\}$ according to the distribution $\boldsymbol{\pi}$, how could we go about doing so? (Hint: Start anywhere and wait long enough. In fact, this simple example illustrates the underlying principle behind what is known as Markov Chain Monte Carlo (MCMC), a very useful algorithm used in computer science and machine learning!)

Convergence to stationary distribution

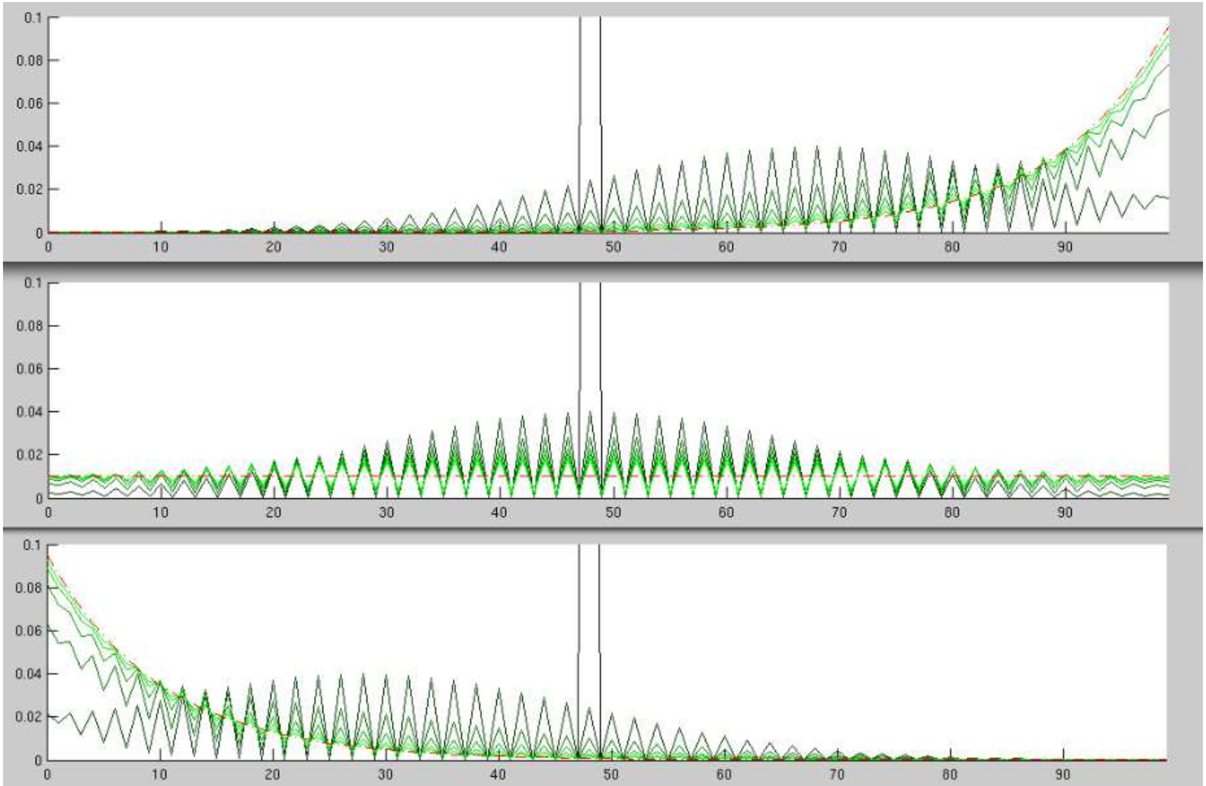

Figure 4. Distribution $\boldsymbol{x}(k)$ for Models 1-3 at times $k = 0, 400, 800, ..., 2000$. x-axis is state space S, y-axis is value of probability vector at states $i \in S$. Brighter green curves correspond to larger times k. Stationary distribution $\boldsymbol{\pi}$ given by dashed red curve.

In Figure 4, we plot the distribution $\boldsymbol{x}(k)$ at particular values of time k (think of these graphs as vertical slices of the surface graphed in Figure 3). The brighter green the color the larger the time. The stationary distribution $\boldsymbol{\pi}$ is calculated by finding the eigenvector of A corresponding to eigenvalue 1 and is plotted as a dashed red curve. In particular, we can see that $\boldsymbol{x}(k)$ converges to $\boldsymbol{\pi}$ for large times k.

There are some interesting features we see here that we might have missed in Figure 3. First, notice the oscillations in $\boldsymbol{x}(k)$ which tend be become less pronounced for larger values of k. This is not some sort of numerical error. It is a reflection of the fact that for the first 50 time steps or so the distribution $\boldsymbol{x}(k)$ alternates between having zero values on odd-numbered states and zeros on even-numbered states, before interactions with the boundary begin to destroy this effect. If we had plotted the graphs at times $k = 1, 401, 801, ..., 1601$ we would have seen $\boldsymbol{x}(k)$ having values close to 0 on even-numbered states instead of odd-numbered ones. These effects are even more pronounced in Model 2 since it takes longer for the typical walker to find its way to the boundary, so oscillations persist and convergence to the stationary distribution takes significantly longer.

To summarize, this Markov chain initially behaves like a periodic system—a walker can only return to its present state at times divisible by 2. If our boundary conditions were changed so that walkers at 1 or 100 were not allowed to remain there for more than one time step, this periodicity would be preserved and we would never converge to any stationary state (averaging over two time steps, however, we would).

Ergodicity and sample path properties

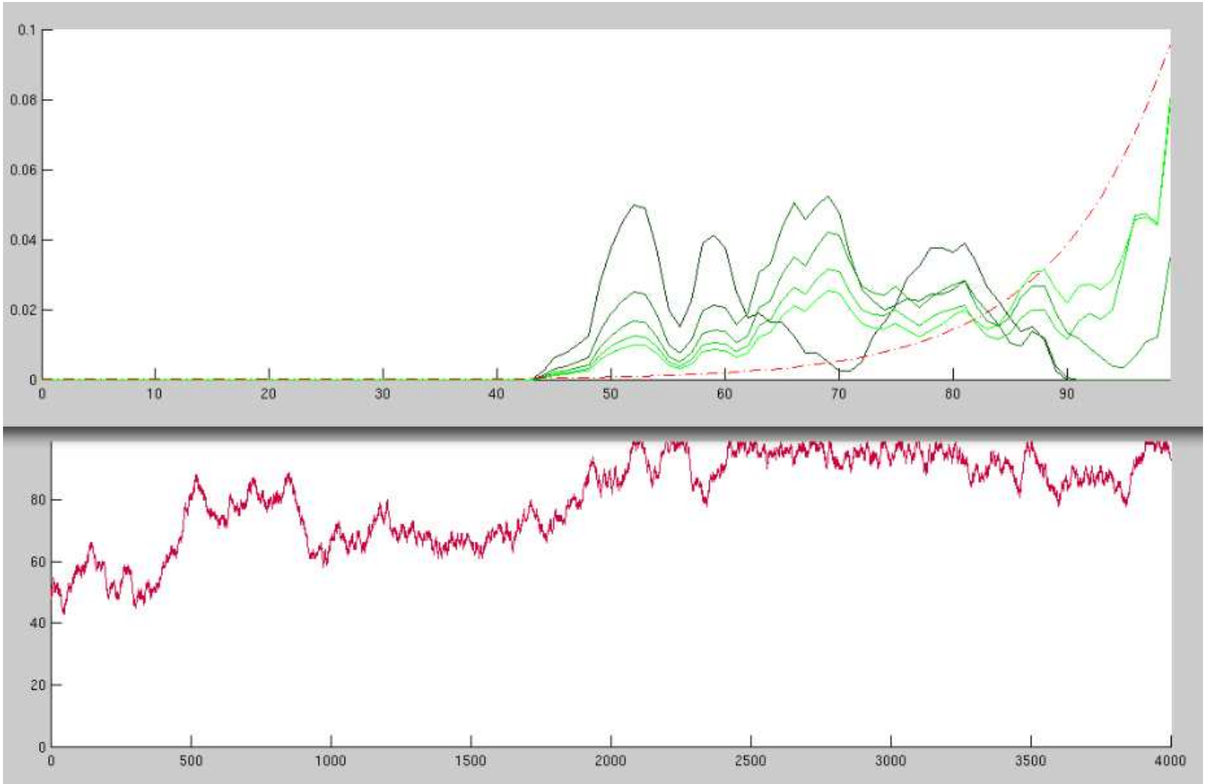

Figure 5. Time averages $\hat{\boldsymbol{x}}_T = \frac{1}{T}\sum_{k=0}^{T-1} \mathbf{1}_{Z(k)=i}$ of a sample path $Z(k)$ generated from Model 1, for $T = 800$, 1600, 2400, 3200, 4000. x-axis is state space, y-axis is proportion of time spent at state $i \in S$. Brighter green curves correspond to larger length T of time interval.

Finally, we end with a short discussion about the path properties of any given walker. The lower part of Figure 5 shows a randomly generated path $Z(k)$ taken by one walker in Model 1 over 4000 time steps. The upper part of the figure shows the proportion of time between $k = 0$ and $k = T - 1$ spent by the walker at each state $i \in S$, with $T = 800, 1600, 2400, 3200, 4000$. That is, we graph

$$\hat{\boldsymbol{x}}_T = \frac{1}{T}\sum_{k=0}^{T-1} \mathbf{1}_{Z(k)=i}, \qquad i \in S$$

where $\mathbf{1}_{Z(k)=i} = 1$ if $Z(k) = i$ and 0 otherwise. This generates a series of histograms, for which the brighter green the color the larger the length T of the time interval on which to average. Note that as T gets larger this histogram approaches the stationary distribution $\boldsymbol{\pi}$. In other words, if we look at just one random path over long enough time intervals, it will mimic the behavior of the entire ensemble of paths! Therefore, a time-average $\hat{\boldsymbol{x}}_T$ of any one random path $Z(k)$ converges to the ensemble average $\boldsymbol{x}(k)$ at a *fixed* (large) time k:

$$\lim_{T\to\infty} \hat{\boldsymbol{x}}_T = \lim_{k\to\infty} \boldsymbol{x}(k) = \boldsymbol{\pi}.$$

This is what is known as **ergodicity** of the random process $Z(k)$.