MATH 56A: STOCHASTIC PROCESSES
CHAPTER 1

1. Finite Markov chains

For the sake of completeness of these notes I decided to write a summary of the basic concepts of finite Markov chains. The topics in this chapter are:

(1) definition of a Markov chain
(2) communication classes
(3) classification of states
(4) periodic/aperiodic
(5) invariant probability distribution
(6) return time
(7) substochastic matrix

1.1. **definition.**

A *stochastic process* is a random process which evolves with time. In other words, we have random variables X_t, Y_t, etc. which depend on time t. For example X_t could be your location at time t (where t is in the future).

A *finite Markov chain* is a special kind of stochastic process with the following properties

- There are only a finite number of states. If we are talking about your location, this means we are only interested in which building you are in and not your exact position in the building.
- Time is discrete: For example, things change only at 1pm, 2pm, etc. and never at 1:12pm or any time in between. X_n is the state at time n where $n = 0, 1, 2$, etc.
- No memory. Your (random) movement at time n depends only on X_n and is independent of X_t for $t < n$ (You forget the past and your decision making process is based only on the present state.)
- Time homogeneous: Your rules of movement are the same at all times.

To summarize: You have a finite number of building that you can move around in. You can only move on the hour. Your decision making process is random and depends only on your present location and not

Date: December 15, 2006.

on past locations and does not take into account what time it is. (You move at midnight in the same way that you do at noon.)

Definition 1.1. *A* finite Markov chain *is a sequence of random variables $X_0, X_1, \cdots$ which take values in a finite set S called the* state space *and for any n and any values of $x_0, x_1, \cdots, x_n$ we have:*

$$\mathbb{P}(X_{n+1} = x \mid X_0 = x_0, X_1 = x_1, \cdots, X_n = x_n) = \mathbb{P}(X_1 = x \mid X_0 = x_n)$$

The $S \times S$ matrix P with entries

$$p(x, y) := \mathbb{P}(X_1 = y \mid X_0 = x)$$

is called the transition matrix*.*

For example, suppose that you have 4 points: 0,1,2,3 and at each step you either increase one or decrease one with equal probability except at the end points. Suppose that, when you reach 0, you cannot leave. (0 is called an *absorbing state*.) Suppose that when you reach 3 you always go to 2. (3 is called a *reflecting wall*.) Then the transition matrix is

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Notice that the numbers are all nonnegative and the numbers in each row add up to 1. This characterizes all transition matrices.

In the discussion of Markov chains, there are qualitative non-numerical concepts and quantitative, computational concepts. The qualitative concepts are: communication classes, their classification and periodicity.

1.2. communication classes.

Two states *communicate with each other* if they are equal or if it is possible (with positive probability) to get from either one to the other in a finite amount of time. We write $x \leftrightarrow y$. This is an equivalence relation and the equivalence classes are called *communication classes*.

In the example above, $\{0\}$ and $\{1, 2, 3\}$ are the communication classes.

A Markov chain is *irreducible* if it has only one communication class, i.e., if it is possible to get from any state to any other state.

1.3. classification of states: transient and recurrent.

There are two types of communication classes: recurrent and transient. At this point, I allowed the state space S to be infinite so that you don't get the wrong idea.

A communication class C is *recurrent* if for any state $x \in C$, you will keep returning to x an infinite number of times with probability one. A communication class C is *transient* if, starting at any $x \in C$, you will return to x only a finite number of times with probability one.

The theorem is that these are the only two possibilities. I proved this in class:

Lemma 1.2. *If $p = p(i,j) > 0$ and i is recurrent then so is j. In fact, if you start in state i you will go to state j an infinite number of times with probability one.*

Proof. This is the same as saying that the probability of going to state j only a finite number of times is zero. To prove this suppose that the Markov chain goes to state j only a finite number of times. Then there is a last time, say $X_m = j$. Then you can never return to j.

But i is recurrent. So, with probability one, you will go to i an infinite number of times after time m. Say at times $n_1 < n_2 < n_3 < \cdots$ (all $> m$). But

$$\mathbb{P}(X_{n_1+1} \neq j \,|\, X_{n_1} = i) = 1 - p$$

$$\mathbb{P}(X_{n_2+1} \neq j \,|\, X_{n_2} = i) = 1 - p$$

The product is $(1-p)^2, (1-p)^3$, etc. which converges to 0. So, the probability is zero that in all of these infinite times that you visit i you will never go to j. This is what I was claiming. □

Theorem 1.3. *Once you leave a communication class you can never return.*

Theorem 1.4. *Recurrent communication classes are absorbing.*

The lemma shows that, if you could leave a recurrent communication class, you will with probability one. This would be a contradiction to the definition of recurrent. So, you cannot leave a recurrent class.

The lemma that I proved can be rephrased as follows:

Lemma 1.5. *If you make an infinite number of attempts and you have a fixed positive probability p of success then, with probability one, you will succeed an infinite number of times.*

The *strong law of large numbers* says that, with probability one, the proportion of trials which are successful will converge to p as the number of trials goes to infinity. I.e., the *experimental value* of the probability p will converge to the *theoretical value* with probability one.

1.4. **periodic chains.** We are interested in the time it takes to return to a state i. The *return time* T_i to state i is the smallest positive integer n so that $X_n = i$ given that $X_0 = i$. In other words, you start at state i and count how many turns it takes to return to the same state i. This number is T_i. It is random. For example, in the random walk example given above, $\mathbb{P}(T_3 = 2) = 1/2, \mathbb{P}(T_2 = 2) = 3/4$.

The *period* of a state i is the greatest common divisor of all possible return times to state i. For the random walk on an infinite straight line (or on a finite line with reflecting walls), the period of every state is 2 because it always takes an even number of steps (the same number right as left) to get back to the same point.

A state i is *aperiodic* if the period is 1.

Theorem 1.6. *States in the same communication class have the same period.*

1.5. **Invariant probability distribution.**

Definition 1.7. *A probability distribution π is called* invariant *if*

$$\pi P = \pi$$

Remember that a *probability distribution* is a vector with nonnegative coordinates adding up to 1:

$$\sum_{i=1}^{n} \pi(i) = 1, \quad \pi(i) \geq 0$$

where n is the number of states. As an application of the Perron-Frobenius theorem we get:

Theorem 1.8. *If the Markov chain is irreducible and aperiodic then P^n converges to a matrix in which every row is equal to the invariant distribution π which is unique:*

$$P^\infty = \begin{pmatrix} \pi \\ \pi \\ \cdots \\ \pi \end{pmatrix}$$

If the Markov chain is periodic then P^n depends on n modulo the period d. However, the average value of P^n will still converge:

Theorem 1.9. *For any finite Markov chain the average value of P^n converges to a matrix in which every row is equal to the unique invariant distribution π:*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} P^k = \begin{pmatrix} \pi \\ \cdots \\ \pi \end{pmatrix}$$

Proof. Let

$$S_n = \frac{1}{n}\sum_{k=1}^{n} P^k$$

Then each row of S_n adds to 1 and

$$S_nP = S_n + \frac{1}{n}(P^{n+1} - P)$$

So, $S_nP \approx S_n$ and $S_\infty P = S_\infty$. Since each row of S_∞ adds to 1, each row is equal to the invariant distribution π (which is unique by Perron-Frobenius). □

1.6. **Return time.** I explained in class the relation between the return time T_i to state i and the value of the invariant distribution $\pi(i)$:

$$\mathbb{E}(T_i) = \frac{1}{\pi(i)}$$

Proof. Begin with the last theorem:

$$\pi(j) = \lim_{n\to\infty} \frac{1}{n}\sum_{m=1}^{n} \mathbb{P}(X_m = j)$$

Now use the fact that probability is the expected value of the indicator function: $\mathbb{P}(A) = \mathbb{E}(I(A))$ and expected value is linear: $\sum \mathbb{E} = \mathbb{E}\sum$

$$\pi(j) = \lim_{n\to\infty} \frac{1}{n}\mathbb{E}\left(\sum_{m=1}^{n} I(X_m = j)\right)$$

This is the average expected number of visits to state j. If, in n steps, you visit a state k times then the average time from one visit to the next is n/k and the average number of visits is k/n. So they are inverse to each other:

$$\pi(j) = \frac{k}{n} = \frac{1}{n/k} = \frac{1}{\mathbb{E}(T_j)}$$

□

1.7. **substochastic matrix.** If there are both absorbing classes and transient classes in the Markov chain then you get a *substochastic matrix* Q which is the transient-to-transient transition matrix: $Q = (p(x, y))$ where we take x, y only from transient states. (Actually, you can take any subset of the set of states.) Since this is only part of the transition matrix P, the rows may not add up to 1. But we know that the entries are all nonnegative and the rows add up to at most 1.

Definition 1.10. *A* substochastic matrix *is a square matrix Q whose entries are all nonnegative with rows adding up to at most 1.*

I used a model from economics to explain what was the point of doing this.

1.7.1. *Leontief model.* In this model we have r factories which produce goods. For every dollar worth of goods that factor i produces, it needs q_{ij} dollars worth of the output of factory j. In order to be profitable or to at least break even we need the production cost to be less than or equal to one dollar:

$$q_{i1} + q_{i2} + \cdots + q_{ir} \leq 1$$

In other words, each row of the matrix $Q = (q_{ij})$ must add up to at most 1. So, Q is a substochastic matrix.

To analyze how this works we follow the dollar. q_{ij} represents goods going from j to i and it represents money going from i to j.

Now look at the total amount of money and what happens to it. Out of each dollar that factory i gets, it must give $\sum q_{ij}$ to other factories. What remains:

$$1 - \sum q_{ij}$$

is profit. Let's say it puts this money in the bank. When we add the bank to the system we get a Markov chain with $r + 1$ states. The transition matrix is $P = (p(i, j))$ where

$$p(i,j) = \begin{cases} q_{ij} & \text{if } i, j \neq 0 \\ 1 - \sum_k q_{ik} & \text{if } i \neq 0, j = 0 \\ 0 & \text{if } i = 0, j \neq 0 \\ 1 & \text{if } i = j = 0 \end{cases}$$

Note that this formula can be used to convert any substochastic matrix into a Markov chain by adding one absorbing state.

Problem The problem is to figure out how much each factory needs to produce in order for the net production (not counting inter-industry consumption) to be equal to a fixed vector γ. This is the *consumption vector*. The consumer wants γ_i dollars worth of stuff from factory i. To find the answer we just follow the money after the consumer buys the goods.

I explained it in class like this: The consumer orders γ_i worth of goods from factory i. On day zero, each factory i produces the requested goods using its inventory of supplies. Then it orders supplies from the other factories to replenish its inventory. On day one, each factory produces goods to fill the orders from the other factories using its inventory. And so on. Eventually, (in the limit as $n \to \infty$), the inventories are all back to normal and all the money is in the bank *assuming that all the factories are transient states.* The total production is given by adding up the production on each day. Factory i produces γ_i on day 0,

$$\sum_j \gamma_j q_{ji} = (\gamma Q)_i$$

on day 1,

$$\sum_j \gamma_j q_{jk} \sum_k q_{ki} = (\gamma Q^2)_i$$

on day 2, $(\gamma Q^3)_i$ on day 3, etc. So, the total production of factory i is

$$(\gamma(I + Q + Q^2 + Q^3 + \cdots))_i = (\gamma(I - Q)^{-1})_i$$

Lemma 1.11. *The factories are all transient states if and only if $I - Q$ is invertible.*

Proof. If the factories are all transient then the money will all eventually end up in the bank. Equivalently, the matrices Q^n converge to zero. So $I - Q$ is invertible. Conversely, if there is a recurrent class, it must consist of nonprofit organizations which require only the output from other nonprofits. Then these nonprofit factories give a Markov process with equilibrium distribution π. This will be a vector with $\pi Q = \pi$. So, it is a null vector of $I - Q$ showing that $I - Q$ is not invertible. □

1.7.2. *avoiding states.* Substochastic matrices are used to calculate the probability of reaching one set of states A before reaching another set B (assuming that they are disjoint subsets of S). To do this you first combine them into two recurrent states. You also need to assume there are no other recurrent states (otherwise the Markov chain could get stuck in a third state and never reach A or B).

Suppose that there are n transient states and two absorbing states A and B. Let Q be the $n \times n$ transient-to-transient transition matrix and let S be the $n \times 2$ transient-to-recurrent matrix. Then the (i, A) entry of $Q^k S$ is the probability of getting from state i to state A in exactly $k + 1$ steps. So, the total probability of ending up in state A is

$$\mathbb{P}_i(A) = ((I + Q + Q^2 + Q^3 + \cdots)S)_{iA} = ((I - Q)^{-1} S)_{iA}$$