

Markov Chains

Arnoldo Frigessi* Bernd Heidegott†

November 4, 2015

1 Introduction

Markov chains are stochastic models which play an important role in many applications in areas as diverse as biology, finance, and industrial production. Roughly speaking, Markov chains are used for modeling how a system moves from one state to another in time. Transitions between state are random and governed by a conditional probability distribution which assigns a probability to the move into a new state, given the current state of the system. This dependence represents the memory of the system. A basic example of a Markov chain is the so-called random walk on the non-negative integers, defined as follows. Let $X_t \in \mathbb{N}$, for $t \in \mathbb{N}$, be a sequence of random variables with initial value $X_0 = 0$. Assume that $\mathbb{P}(X_{t+1} = X_t + 1 | X_t \geq 1) = p = 1 - \mathbb{P}(X_{t+1} = X_t - 1 | X_t \geq 1)$ and $\mathbb{P}(X_{t+1} = 1 | X_t = 0) = 1$ for some given value of $p \in [0, 1]$. The sequence $X = \{X_t : t \in \mathbb{N}\}$ is an example of a Markov chain (for the definition see below). There are many aspects of the chain X one can be interested in, including (i) whether X returns to 0 in a finite number of steps (this holds for $0 \leq p \leq 1/2$), (ii) the expected number of steps until the chain returns to 0 (which is finite for $0 \leq p < 1/2$), and (iii) the limiting behavior of X_t as $t \rightarrow \infty$.

A few examples help illustrate the important concepts. A useful model in modeling infectious diseases assumes that there are four possible states: Susceptible (S), Infected (I), Immune (A), Dead (R). Possible transitions are from S to I, S or R; from I to A or R; from A to A or R; from R to R only. The transition probabilities, from S to I, S to R and the loop S to S, must sum to one and can depend on characteristics of the individuals modeled, like age, gender, life style, etc. All individuals start in S, and move at each time unit (say a day). Given observations of the sequence of visited states (called trajectory) for a sample of individuals, with their personal characteristics, one can estimate the transition probabilities, by logistic regression, for example. This model assumes that the transition probability at time t from one state A to state B only depends on the state A and not on the trajectory that leads to A. This might not be realistic,

*Department of Biostatistics, University of Oslo & Norwegian Computing Centre, Oslo, Norway

†Department of Econometrics, Vrije Universiteit, Amsterdam, The Netherlands

as for example, remaining in the diseased state I over many days could increase the probability of transition to R. It is possible to model a system with longer memory, and thus leave the simplest setting of a Markov Chain (though one can formulate such a model still as a Markov Chain over a more complex state space which includes the length of stay in the current state). A second example refers to finance. Here we follow the daily value in Euro of a stock. The state space is continuous, and one can model the transitions from state x Euro to y Euro with an appropriate Normal density with mean $x - y$. The time series of the value of the stock might well show a longer memory, which one would typically model with some autoregressive terms, leading to more complex process again. As a further example, consider the set of all web pages on the Internet as the state space of a giant Markov chain, where the user clicks from one page to the next, according to a transition probability. A Markov chain has been used to model such a process. The transitions from the current web page to the next web page can be modeled as a mixture of two terms: with probability λ the user follows one of the links present in the current web page and among these uniformly; with probability $1 - \lambda$ the user chooses another web page at random among all other ones. Typically $\lambda = 0.85$. The longterm fraction of visits of the Markov chain to each state can then be interpreted as a ranking of the web page, as done basically by the Google page ranking algorithm. Again, one could discuss how correct the assumption is that only the current web page determines the transition probability to the next one. Finally, Markov Chain Monte Carlo (MCMC) algorithms are Markov chains, where at each iteration, a new state is visited according to a transition probability that depends on the current state. These stochastic algorithms are used to sample from a distribution on the state space, which is the distribution of the chain in the limit, when enough iterations have been performed. The modeler has to critically validate Markov and homogeneity hypothesis before trusting results based on the Markov chain model, or chains with higher order of memory. In general a stochastic process has the Markov property if the probability to enter a state in the future is independent of the states visited in the past given the current state.

In the literature the term Markov processes is used for Markov chains for both discrete- and continuous time cases, which is the setting of this note. Standard textbooks on Markov chains are [3, 4, 5, 6]. In this paper we follow [2] and use the term 'Markov chain' for the discrete time case and the term 'Markov process' for the continuous time case. General references on Markov chains are [7, 8, 9, 10, 11].

2 Discrete Time Markov Chains

Consider a sequence of random variables $X = \{X_t : t \in \mathbb{N}\}$ defined on a common underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with discrete state space (S, \mathcal{S}) , i.e., technically, X_t is $\mathcal{F} - \mathcal{S}$ -measurable for $t \in \mathbb{N}$. The defining property of a Markov chain is that the distribution of X_{t+1} depends on the past only through

the immediate predecessor X_t , i.e., it holds that

$$\mathbb{P}(X_{t+1} = x | X_0 = x_0, X_1 = x_1, \dots, X_{t-1} = x_{t-1}, X_t = y) = \mathbb{P}(X_{t+1} = x | X_t = y),$$

where x, y and all other x_i are elements of the given state space S . If $\mathbb{P}(X_{t+1} = x | X_t = y)$ does not depend on t , the chain is called *homogenous*. Provided that S is at most countable, the transition probabilities of a homogeneous Markov Chain are organised into a matrix $P = (p_{x,y})_{S \times S}$, where $p_{x,y} = \mathbb{P}(X_{t+1} = y | X_t = x)$ is the probability of a transition from x to y . The matrix P is called the *one-step transition probability matrix* of the Markov chain. For the random walk example above, the transition matrix is given by $p_{i,i+1} = p$, $p_{i,i-1} = p - 1$, for $i \geq 1$, $p_{0,1} = 1$ and otherwise zero. In order to fully define a Markov Chain it is necessary to assign an initial distribution $\mu = (\mathbb{P}(X_0 = s) : s \in S)$. The marginal distribution at time t can then be computed as

$$\mathbb{P}(X_t = x) = \sum_{s \in S} p_{s,x}^{(t)} \mathbb{P}(X_0 = s),$$

where $p_{s,x}^{(t)}$ denotes the s, x element of the t -th power of the transition matrix. Given an initial distribution μ and a transition matrix P , the distribution of the Markov chain X is uniquely defined.

A Markov chain is said to be *aperiodic* if for each pair of states i, j the greatest common divisor of the set of all t such that $p_{ij}^{(t)} > 0$ is one. The random walk in our introductory example fails to be aperiodic as any path from starting in 0 and returning there has a length that is a multiple of 2.

A distribution $(\pi_i : i \in S)$ is called a *stationary distribution* of P if

$$\pi P = \pi.$$

A key topic in Markov chain theory is the study of the limiting behavior of X . X has limiting distribution ν for initial distribution μ if

$$\lim_{t \rightarrow \infty} \mu P^t = \nu. \tag{1}$$

Any limiting distribution is a stationary distribution. A Markov chain is called *ergodic* if the limit in (1) is independent of the initial distribution. Consequently, an ergodic Markov chain has a unique limiting distribution and this limiting distribution is also a stationary distribution. If P fails to be aperiodic, then the limit in (1) may not exist but can be replaced by the Cesaro limit

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=1}^t \mu P^k = \nu,$$

which always exists for finite Markov chains.

A Markov chain is called *irreducible* if for any pair of states $i, j \in S$, there exists a path from i to j that X can follow with positive probability. In other words, any state can be reached from any other state with positive probability.

An irreducible Markov chain is called *recurrent* if the number of steps from a state i to the first visit of a state j , denoted by $\tau_{i,j}$, is almost surely finite for all $i, j \in S$, and it is called *positive recurrent* if $\mathbb{E}[\tau_{i,i}] < \infty$ for at least one $i \in S$. Note that for $p = 1/2$ the random walk is recurrent and for $p < 1/2$ it is positive recurrent.

Any aperiodic, irreducible and positive recurrent Markov chain P possesses a unique stationary distribution π which is the unique probability vector solving $\pi P = \pi$ (and which is also the unique limiting distribution). This ergodic theorem is one of the central results and it has been established in many variations and extensions, see the references. Efficient algorithms for computing π are available, but fail for very large state-spaces.

An important topic is the estimation of the (one-step) transition probabilities. Consider a discrete time, homogeneous Markov chain with finite state space $S = \{1, 2, \dots, m\}$, observed at time points $0, 1, 2, \dots, T$ on the trajectory $s_0, s_1, s_2, \dots, s_T$. We wish to estimate the transition probabilities $p_{i,j}$ by maximum likelihood. The likelihood is

$$\mathbb{P}(X_0 = s_0) \prod_{t=0}^{T-1} \mathbb{P}(X_{t+1} = s_{t+1} | X_t = s_t) = \mathbb{P}(X_0 = s_0) \prod_{i=1}^m \prod_{j=1}^m p_{i,j}^{k(i,j)}$$

where $k(i, j)$ is the number of transitions from i to j in the observed trajectory. Ignoring the initial factor, the maximum likelihood estimator of $p_{i,j}$ is found to be equal to $\hat{p}_{i,j} = \frac{k(i,j)}{k(i,\cdot)}$, where $k(i, \cdot)$ is the number of transitions out from state i . Standard likelihood asymptotic applies, despite the data are dependent, as $k(i, \cdot) \rightarrow \infty$, which will happen if the chain is ergodic. The asymptotic variance of the maximum likelihood estimates can be approximated by $\text{var}(\hat{p}_{i,j}) \sim \hat{p}_{i,j}(1 - \hat{p}_{i,j})/k(i, \cdot)$. The covariances are zero, except $\text{cov}(\hat{p}_{i,j}, \hat{p}_{i,j'}) \sim -\hat{p}_{i,j}\hat{p}_{i,j'}/k(i, \cdot)$ for $j \neq j'$. If the trajectory is short, the initial distribution should be considered. A possible model is to use the stationary distribution $\pi(s_0)$, which depend on the unknown transition probabilities. Hence numerical maximization is needed to obtain the maximum likelihood estimates. In certain medical applications, an alternative asymptotic regime can be of interest, when many (k) short trajectories are observed, and $k \rightarrow \infty$. In this case the initial distribution cannot be neglected.

3 Markov Chains and Markov Processes

Let $\{X_t : t \geq 0\}$ denote the (continuous time) Markov process on state space (S, \mathcal{S}) with transition matrix $P(t)$, i.e.,

$$(P(t))_{ij} = \mathbb{P}(X_{t+s} = j | X_s = i), \quad s \geq 0, \quad i, j \in S.$$

Under some mild regularity conditions it holds that the *generator matrix* Q , defined as

$$\left. \frac{d}{dt} \right|_{t=0} P(t) = Q,$$

exists for $P(t)$. The stationary distribution of a Markov process can be found as the unique probability π that solves $\pi Q = 0$, see [1]. A generator matrix Q is called uniformizable with rate μ if $\mu = \sup_j |q_{jj}| < \infty$. While any finite dimensional generator matrix is uniformizable a classical example of a Markov process on denumerable state space that fails to have this property is the M/M/ ∞ queue. Note that if Q is uniformizable with rate μ , then Q is uniformizable with rate η for any $\eta > \mu$. Let Q be uniformizable with rate μ and introduce the Markov chain P_μ as follows

$$[P_\mu]_{ij} = \begin{cases} q_{ij}/\mu & i \neq j \\ 1 + q_{ii}/\mu & i = j, \end{cases} \quad (2)$$

for $i, j \in S$, or, in shorthand notation,

$$P_\mu = I + \frac{1}{\mu}Q,$$

then it holds that

$$P(t) = e^{-\mu t} \sum_{n=0}^{\infty} \frac{(\mu t)^n}{n!} (P_\mu)^n, \quad t \geq 0. \quad (3)$$

Moreover, the stationary distribution of P_μ and $P(t)$ coincide. The Markov chain $\mathcal{X}_\mu = \{X_n^\mu : n \geq 0\}$ with transition probability matrix P_μ is called the *sampled chain*. The relationship between \mathcal{X} and \mathcal{X}_μ can be expressed as follows. Let $N_\mu(t)$ denote a Poisson process with rate μ , then $X_{N_\mu(t)}^\mu$ and X_t are equal in distribution for all $t \geq 0$. From the above it becomes clear that the analysis of the stationary behavior of a (uniformizable) continuous time Markov chain reduces to that of a discrete time Markov chain.

References

- [1] W. Anderson: *Continuous-time Markov Chains, An Applications Oriented Approach*, Springer, New York, 1991.
- [2] M. Iosifescu: *Finite Markov Processes and Their Applications*, Wiley, New York, 1980.
- [3] M. Kijima: *Markov Processes for Stochastic Modelling*, Chapman & Hall, London, 1997.
- [4] S. Meyn and R. Tweedie: *Markov Chains and Stochastic Stability*, Springer, London, 1993.
- [5] E. Nummelin: *General Irreducible Markov Chains and Non-negative Operators*, Cambridge University Press, Cambridge, 1984.
- [6] D. Revuz: *Markov Chains* Second Edition, North-Holland, Amsterdam, 1984.

- [7] W. Feller: *An Introduction to Probability Theory and its Applications*. Vol. I. Third edition. Wiley, New York, 1968.
- [8] W. Gilks, S. Richardson and D. Spiegelhalter (editors): *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, 1995.
- [9] O. Haeggstroem: *Finite Markov Chains and Algorithmic Applications*, London Mathematical Society Student Texts (No. 52), 2002.
- [10] J. Kemeny and J. Snell: *Finite Markov Chains*, originally published by Van Nostrand Publishing Company, 1960, Springer Verlag, 3rd printing, 1983.
- [11] E. Seneta: *Non-negative Matrices and Markov Chains*, originally published by Allen & Unwin Ltd., London, 1973 Springer Series in Statistics, 2nd rev. ed., 2006.