Markov models A Markov model is a nonlinear state space model subject to stochastic disturbances. It would appear that the introduction of noise would make these models much more complex than their deterministic analogs. In fact, the opposite is frequently true: By considering the evolution of the underlying distributions, rather than the state itself, a linear system is obtained. This is the basis of a development of Markov models that parallels the theory of linear state space models, with or without control. In particular, spectral theory and Lyapunov theory play a fundamental role in analysis and design.

These lectures are designed for an audience with an understanding of stochastic processes (no need for measure theory), and exposure to deterministic state space control at the first-year graduate level. Examples will be used throughout to illustrate the theory.

The two main references are

(1) S. P. Meyn. *Control Techniques for Complex Networks.* Cambridge University Press, Cambridge, 2007.

(2) S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability.* Cambridge University Press, Cambridge, second edition, 2009. Published in the Cambridge Mathematical Library.

1 An Introduction to Markov Models

Monday, 01/04/10 2:00 - 3:30

This first lecture lays out the goals for the week, and begins the theoretical development with structural theory of Markov models without control. In particular, we obtain a simple representation of the steady-state distribution, as well as solutions to dynamic programming equations, generalizing the Lyapunov equations that arise in linear state space theory.

Reading Sections A1-A3 of CTCN

Organization

(i) Motivation. Some applications to math and to the real world,

- Optimal Control [17], Risk Sensitive Optimal Control [1, 3]
- Approximate Dynamic Programming [4, 16]
- Finance
- Dynamic Economic Systems [6, 5]
- Large Deviations [1, 12, 11].
- Simulation [8, 13].
- Google (spectral graph theory and Markov spectral theory [9, 15, 14]).

(ii) Notation: Mainly discrete-space and discrete-time.
Transition matrix P, state space X.

Generator (discrete-time) $\mathcal{D} = P - I$,
Resolvent:

$$R_\alpha = [\alpha I - \mathcal{D}]^{-1} = \sum_{t=0}^{\infty} (1+\alpha)^{-t-1} R^t, \qquad \alpha > 0.$$

Resolvent equations,

$$\mathcal{D} R_\alpha = \alpha R_\alpha - I$$

Simple examples: State space model and the CRW queue.

(iii) φ-Irreducibility for general space chains:
The Markov chain $\boldsymbol{X}$ is called *φ-irreducible* if there exists a measure φ on $\mathcal{B}(\mathsf{X})$ such that, whenever $\varphi(A) > 0$, we have $R(x, A) > 0$ for all $x \in \mathsf{X}$.

(iv) Small sets and functions: $R \geq s \otimes \nu$. Potential kernel,

$$G = \sum_{t=0}^{\infty} (R - s \otimes \nu)^t$$

(v) Un-normlized invariant measure $\mu = \nu G$,

$$\mu \mathcal{D} = 0$$

$\pi(\,\cdot\,) = \mu(\,\cdot\,)/\mu(\mathsf{X})$.

(vi) Solution to Poisson's equation $h = G\tilde{c}$,

$$c + \mathcal{D} h = \eta \tag{1}$$

$\eta = \pi(c)$.

(vii) Vector space setting: Given $f \colon \mathsf{X} \to [1, \infty)$, function $c \colon \mathsf{X} \to \mathbb{R}$, and measure μ on X,

$$\|c\|_f = \sup_{x \in \mathsf{X}} \frac{c(x)}{f(x)} \qquad \qquad \|\mu\|_f = \sup_{g : \|g\|_f < \infty} \frac{\mu(g)}{\|g\|_f}$$

History of small sets (from [18]):

> "... the ideas for the proof of their [small sets] existence go back to Doeblin [7], although the actual existence as we have it here is from Jain and Jamison [10]. Our proof is based on that in Orey [21], where small sets are called C-sets. Nummelin [20] Chapter 2 has a thorough discussion of conditions equivalent to that we use here for small sets... "

2 Stochastic Stability and Dynamic Programming

Tuesday, 01/05/10 9:30 - 11:30

The f-Norm Ergodic Theorem of MCSS characterizes stability in terms of a particular Lyapunov drift condition. This result is proven, and illustrated with examples. For controlled Markov chains, a variant of the Lyapunov drift condition is the average-cost optimality equation of dynamic programming.

Reading Sections A4-A6 and Chs. 8 and 9 of CTCN. See also Part III of MCSS

Organization

(i) Basic drift criterion
$$\mathcal{D}V \leq -f + bs \tag{V4}$$

(ii) Discussion and examples:
How to construct a Lyapunov function? How to approximate a value function?

(iii) Comparison Theorem: $\pi(f) \leq b\pi(s)$.

(iv) Ergodic Theorem: $P^t c \to \pi(c)$ if $\|c\|_f < \infty$.

(v) Value function bounds: $\|h\|_V < \infty$ with $h = G\tilde{c}$, provided $\|c\|_f < \infty$.

(vi) Average-Cost Optimality Equation:
$$\min_u \big(c(x,u) + \mathcal{D}_u h^*\,(x) \big) = \eta^* \qquad \textit{Compare to } (1).$$

3 Approximate Dynamic Programming

Thursday, 01/07/10 9:30 - 11:00

In many applications we can construct approximations to the solution to a dynamic programming equation. In particular, in examples we have already seen that such approximations can be found by considering deterministic ODE models intended to approximate the stochastic model. In this lecture we show how to use this insight for the purposes of control. The TD-learning and Q-learning algorithms are introduced for this purpose, and illustrated with examples.

Reading Lecture notes from *ECE 555 Stochastic Systems*.
Section 11.5 of [17], and recent publications [4, 16]:
`netfiles.uiuc.edu/meyn/www/spm_pubs.html`
See also [19, 22, 2]

Organization

(i) Goal: A Riccati Equation for Markov models.

(ii) TD-Learning: $\min_\theta \|h - h^\theta\|_\pi^2$.

(iii) Q-Learning: $\min_\theta \|Q - Q^\theta\|_\pi^2$ (Q-function or Hamiltonian)

4 Spectral Theory and Model Reduction for Markov Models

Friday, 01/08/10 - CCDC Winter 2010 Seminars - 3:00 - 4:00PM

We survey recent approaches to model reduction for Markov models based on spectral theory. The approximating process is described as a hidden Markov model (HMM) with finite state space.

These approximations are possible under a Lyapunov drift condition introduced by Donsker-Varadahn in their treatment of large deviations for diffusions. In the special case of diffusions, we find that eigenfunctions provide a decomposition of the state space into "almost invariant" sets. The diffusion mixes rapidly in each of these subsets prior to exiting, and the exit time from one of these sets is approximately exponential.

Reading Lectures based on joint with with Kontoyiannis, Huisinga, and Schuette, `http://decision.csl.illinois.edu/ meyn/pages/PhaseTransitions/PhaseTransitions.html`

References

[1] S. Balaji and S. P. Meyn. Multiplicative ergodicity and large deviations for an irreducible Markov chain. *Stoch. Proc. Applns.*, 90(1):123–144, 2000.

[2] V. S. Borkar and S. P. Meyn. The O.D.E. method for convergence of stochastic approximation and reinforcement learning. *SIAM J. Control Optim.*, 38(2):447–469, 2000. (also presented at the *IEEE CDC*, December, 1998).

[3] V. S. Borkar and S. P. Meyn. Risk-sensitive optimal control for Markov decision processes with monotone cost. *Math. Oper. Res.*, 27(1):192–209, 2002.

[4] W. Chen, D. Huang, A. Kulkarni, J. Unnikrishnan, Q. Zhu, P. Mehta, S. Meyn, and A. Wierman. Approximate dynamic programming using fluid and diffusion approximations with applications to power management. Accepted for inclusion in the 48th IEEE Conference on Decision and Control, December 16-18 2009.

[5] I.-K. Cho and S. P. Meyn. Efficiency and marginal cost pricing in dynamic competitive markets. Under revision for Theo. Economics, 2006.

[6] In-Koo Cho and S. Meyn. Efficiency and marginal cost pricing in dynamic competitive markets with friction. *Decision and Control, 2007 46th IEEE Conference on*, pages 771–778, Dec. 2007.

[7] W. Doeblin. Eléments d'une théorie générale des chaînes simples constantes de Markov. *Annales Scientifiques de l'Ecole Normale Supérieure*, 57(III):61–111, 1940.

[8] S. G. Henderson, S. P. Meyn, and V. B. Tadić. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems: Theory and Applications*, 13(1-2):149–189, 2003. Special issue on learning, optimization and decision making (invited).

[9] W. Huisinga, S. Meyn, and C. Schütte. Phase transitions and metastability in Markovian and molecular systems. *Ann. Appl. Probab.*, 14(1):419–458, 2004. Presented at the 11TH INFORMS Applied Probability Society Conference, NYC, 2001.

[10] N. Jain and B. Jamison. Contributions to Doeblin's theory of Markov processes. *Z. Wahrscheinlichkeitstheorie und Verw. Geb.*, 8:19–40, 1967.

[11] I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13:304–362, 2003. Presented at the INFORMS Applied Probability Conference, NYC, July, 2001.

[12] I. Kontoyiannis and S. P. Meyn. Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes. *Electron. J. Probab.*, 10(3):61–123 (electronic), 2005.

[13] I. Kontoyiannis and S. P. Meyn. Computable exponential bounds for screened estimation and simulation. *Ann. Appl. Probab.*, 18(4):1491–1518, 2008.

[14] A. N. Langville and C. D. Meyer. A survey of eigenvector methods for Web information retrieval. *SIAM Rev.*, 47(1):135–161 (electronic), 2005.

[15] Amy N. Langville and Carl D. Meyer. *Google's PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press, Princeton, NJ, USA, 2006.

[16] P. Mehta and S. Meyn. Q-learning and Pontryagin's Minimum Principle. Accepted for inclusion in the 48th IEEE Conference on Decision and Control, December 16-18 2009.

[17] S. P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, Cambridge, 2007.

[18] S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, Cambridge, second edition, 2009. Published in the Cambridge Mathematical Library. 1993 edition online: `http://black.csl.uiuc.edu/~meyn/pages/book.html`.

[19] C.C. Moallemi, S. Kumar, and B. Van Roy. Approximate and data-driven dynamic programming for queueing networks. Preprint available at http://moallemi.com/ciamac/research-interests.php, 2008.

[20] E. Nummelin. *General Irreducible Markov Chains and Nonnegative Operators*. Cambridge University Press, Cambridge, 1984.

[21] S. Orey. *Limit Theorems for Markov Chain Transition Probabilities*. Van Nostrand Reinhold, London, 1971.

[22] J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Trans. Automat. Control*, 42(5):674–690, 1997.