

MDPs in Medicine: Opportunities and Challenges

F. J. Díez M. A. Palacios M. Arias

Dept. Artificial Intelligence. UNED

Madrid, Spain

Abstract

In the last three decades hundreds of Markov models have been built for medical applications, but most of them fall under the paradigm of what we call *simple Markov models* (SMMs). Markov decision processes (MDPs) are much more powerful as a decision analysis tool, but they are ignored in medical decision analysis books and the number of medical applications based on them is still very small. In this paper we compare both types of models and discuss the challenges that MDPs must overcome before they can be widely accepted in medicine. We present a software tool, Open-Markov, that addresses those challenges and has been used to build a Markov model for analyzing the cost-effectiveness of the HPV vaccine.

1 Introduction

Markov models were introduced in the beginning of the 20th century by the Russian mathematician Andrei Andreyevich Markov [1906]. In the three decades passed since the pioneering work by Beck and Pauker [1983], hundreds of Markov models have been built for medical problems, mainly for therapy selection and prognosis. The most cited reference is the introductory paper by Sonnenberg and Beck [1993]. Nowadays, all the books on medical decision making devote at least one chapter to Markov models [Drummond *et al.*, 2005; Gold *et al.*, 1996; Haddix *et al.*, 2003; Hunink *et al.*, 2001; Kuntz and Weinstein, 2001; Levin and McEwan, 2001; Muenning, 2008; Parmigiani, 2002; Petitti, 2000; Sloan, 1995; Sox *et al.*, 1988]; two of them make a special emphasis on Markov models applied to cost-effectiveness analysis: [Briggs *et al.*, 2006] and [Gray *et al.*, 2010]

Independently, Markov decision processes (MDPs) were developed around the mid-20th century as a tool for planning, more specifically, for solving multistage decision problems in which the outcomes are partly random and partly under the control of a decision maker [Bellman, 1957]. The main limitation of those models was the assumption that the state of the system is always known with certainty, which is unrealistic in most cases. The relaxation of this assumption led to the emergence of partially observable Markov decision processes (POMDPs) [Åström, 1965], in which the

state of the system is not directly observable, but there is a variable that correlates probabilistically with it. POMDPs were developed in the field of automatic control as an extension of MDPs [Åström, 1965], but currently most of the research about them is carried out in artificial intelligence (AI), again as a tool for planning, especially in robotics [Ghallab *et al.*, 2004]. The main contribution of AI to this field comes from the area of probabilistic graphical models: Bayesian networks [Pearl, 1988] led to the development of dynamic Bayesian networks [Dean and Kanazawa, 1989], which generalize Markov chains and hidden Markov models [Murphy, 2002]. The idea of using several variables to represent the state of the system, instead of only one, led to factored MDPs [Boutilier *et al.*, 1995; 2000] and factored POMDPs [Boutilier and Poole, 1996], which can model efficiently many problems that were unmanageable with flat (i.e., non-factored) representations; correspondingly, there are new algorithms that can solve problems several orders of magnitude bigger than in the recent past [Hoey *et al.*, 1999; Poupart, 2005; Spaan and Vlassis, 2005].

In the rest of the paper, we use the acronym MDPs to refer to both fully observable and partially observable models (FOMDPs and POMDPs, respectively).

In medicine, FOMDPs have been applied to epidemic control [Lefebvre, 1981], anesthesia [Hu *et al.*, 1996], kidney transplantation [Ahn and Hornberger, 1996], spherocytosis [Magni *et al.*, 2000], liver transplantation [Alagoz *et al.*, 2007], HIV [Shechter *et al.*, 2008], type 2 diabetes [Denton *et al.*, 2009; Mason *et al.*, 2011], hepatitis C [Kırkızlar *et al.*, 2010], ovarian hyperstimulation [He *et al.*, 2010], etc.

In turn, POMDPs have been applied to liver transplantation [Tusch, 2000], ischemic heart disease [Hauskrecht and Fraser, 2000; Goulionis, 2009], colorectal cancer [Leshno *et al.*, 2003], anthrax [Izadi and Buckeridge, 2007] breast cancer [Maillart *et al.*, 2008; Ayer *et al.*, 2010], Parkinson's disease [Goulionis *et al.*, 2008; Goulionis and Vozikis, 2009], pneumonia-related sepsis [Kreke *et al.*, 2008], depression [Fard *et al.*, 2008], rehabilitation [Kan *et al.*, 2008], assistance of disabled people [Hoey *et al.*, 2008], diabetic foot syndrome [Goulionis *et al.*, 2010], epilepsy [Pineau *et al.*, 2009], prostate cancer [Zhang *et al.*, 2010], etc. Some of these models are reviewed briefly in so [Schaefer *et al.*, 2004; Alagoz *et al.*, 2010; Zhang *et al.*, 2011].

However, most of the people building Markov models in

medicine are completely unaware of even the existence of MDPs: in fact, none of the 13 books on medical decision making mentions them. In this paper we analyze the differences between simple Markov models (SMMs, defined below) and MDPs, and discuss the conditions required for a wide acceptance of MDPs in health sciences. In particular, we point out some challenges that the specific needs of medicine pose to the researchers of probabilistic graphical models.

2 Markov models vs. MDPs

2.1 Differences in the models

Decisions and policies

In this section, we introduce the term *simple Markov model* (SMM) to differentiate them from MDPs. (Other authors [Schaefer *et al.*, 2004; Alagoz *et al.*, 2010] use the term *Markov model* for what we call SMM, but we prefer to introduce a new term because, in our opinion, MDPs are also Markov models.)¹ The main similarity between both types of models is that they create several instances of each variable, one instance for each time slice (we will qualify this statement below). The main difference is that in SMMs **decisions** are atemporal, i.e., they **are made at the start of the process**, [Beck and Pauker, 1983; Briggs *et al.*, 2006; Gray *et al.*, 2010]; for example, the doctor decides which therapy to apply, and this decision is maintained throughout the process. There are more sophisticated models that include several chance variables and several policies, in which the decision may depend on some observed variables [Briggs *et al.*, 2006], but the basic idea is the same: there is **a very limited set of policies, established beforehand**, and the objective of the evaluation is to pick up the optimal one. Another possibility is to use decision trees with embedded Markov nodes [Sonnenberg and Beck, 1993], each one representing a Markov process in which, again, the decisions are made before the patient enters the sub-model, and for this reason such decision nodes are placed between the root and the Markov nodes in the tree, never between the Markov nodes and the leaves (the utilities). Also in this case the set of possible policies is very limited.

In contrast, MDPs contain **one decision per time slice**, and the policy always depends on the observed evidence, which may be the state itself (in the case of FOMDPs) or another variable that contributes indirect evidence about the state (in the case of POMDPs), and the set of policies evaluated is much higher.

Transition probabilities

In most medical problems, transition probabilities are clearly non-stationary: the probability of death and the probability of getting infected of sexually transmitted disease depends

¹Additionally, we must mention that in artificial intelligence, more specifically, in the area of probabilistic graphical models, “Markov model” is often a synonym of Markov random field, Markov network, or undirected graphical model, as opposed to models based on (acyclic) directed graphs, such as Bayesian networks and influence diagrams [Pearl, 1988]. In contrast, medical MMs, which in general are based on causal relations, are closer to directed graphical models than to undirected graphical models [Pearl, 2000].

strongly on the age of the patient, the probability that a prosthesis fails depends on how long it has been implanted, the toxic effects of a drug depend on the duration of the treatment, etc. Only in models that represent the progression of disease or in models for controlling the administration of a therapy for a short period of time—for example, when applying neurostimulation for the treatment of epilepsy [Pineau *et al.*, 2009]—is it reasonable to use time-independent probabilities, given the limited time span of the analysis.

Therefore, in medical SMMs it is usual to have **non-stationary transition probabilities**, modeled by a mathematical function; for example, by a Weibull distribution [Briggs *et al.*, 2006].

In contrast, most MDPs are stationary, also in medical applications. In our opinion, the reason for this is the difficulty of representing time-dependent transitions in this type of models. One way of representing this time dependency is to use a different transition matrix for each time slice; the problem of this approach lies in the difficulty of eliciting so many different matrices. Other possibility is to use the same transition matrix for all the time slices, but making it depend on a variable that represents time explicitly. However, most of the software tools for MDPs do not admit this possibility.

Given that most medical problems involve non-stationary transitions, as explained above, we may wonder why most of the literature on MDPs and almost all the software tools assume that transitions are stationary. In our opinion, the reason is that most of the research on MDPs has been done in the fields of automatic control and robotics, in which it is reasonable to assume that the underlying properties of the system do not change over time.

2.2 Differences in the evaluation algorithms

The differences between the models lead to applying different algorithms to evaluate them. In SMMs, in which there is a limited set of policies established beforehand, as explained above, the standard evaluation method consists of evaluating the model forwards, computing the probabilities of each time slice from those of the previous one, and summing the rewards, with the corresponding discount factor. This way we obtain the expected value of each policy, which allows us to select the optimal one. In other cases, the cost and effectiveness are summed separately, in order to perform an incremental cost-effectiveness analysis of the different policies [Briggs *et al.*, 2006; Gray *et al.*, 2010]. When the SMM is represented by a decision tree [Sonnenberg and Beck, 1993; TreeAge Software, Inc., 2009], the evaluation method consists of evaluating first the embedded Markov nodes by simulating a certain number of cycle iterations, and then applying the roll-back algorithm from the leaves to the root of the tree. When the SMM is represented in a spreadsheet, the model can be evaluated by programming in the own spreadsheet the macros that implement the computation of probabilities and utilities.

In contrast, the evaluation of MDPs is performed with well-known algorithms, such as value iteration and policy iteration, that return the optimal policy and the expected value.

2.3 Differences in software tools

We can also see that there are differences in the software tools and languages used for the two types of models. In the USA, medical decision trees are usually built and evaluated with the commercial package TreeAge² [TreeAge Software, Inc., 2009], which permits to evaluate embedded Markov nodes and to perform cost-effectiveness analyses (CEAs). However, some of the most active research groups in this area, such as the group led by Prof. Milton Weinstein at Harvard University and the one led by Prof. Mark Roberts at the University of Pittsburgh, use to build their models in C++, because TreeAge does not scale up properly for complex models. In the United Kingdom, mainly at the Universities of York and Oxford, and in other European countries that learned from them, it is usual to build Markov models for CEA with Excel [Briggs *et al.*, 2006], or exceptionally with the statistical language R [Hawkins *et al.*, 2005].

In contrast, most of medical MDPs have been built with general-purpose languages, such as C++, Matlab, or Java, often using general-purpose AI planning tools.

3 Opportunities for MDPs in medicine

From the above analysis we can conclude that the main contributions that MDPs can make to medical decision analysis is the possibility of building much more complex models:

- factored models, in which the relations between the variables are given by causal graphs and the transition probability between two time slices can be decomposed into a set of smaller conditional probability distributions; similarly, the utility function of each can usually be decomposed into smaller utility functions;
- some of those conditional probabilities and utility functions can be built using canonical models [Díez and Druzdzel, 2006] and ADDs [Hoey *et al.*, 1999], which reduces significantly the number of parameters required, facilitates the process of parameter elicitation, and speeds up inference significantly;
- the analysis of the model is not restricted to evaluating a very reduced set of policies established beforehand by the decision analyst; on the contrary, it suffices to build the model by introducing the transition probabilities and the utilities, because there are efficient algorithms that can evaluate the model, examine a huge number of policies, and return the optimal one.

Unfortunately, most medical decision makers are unaware of the possibilities offered by MDPs. One of the causes is the isolation of the different research fields: given the current specialization of science and the amount of publications, it is difficult for health professionals, who often have to divide their time between clinical practice and research, to keep up to date with the advances produced in other fields, especially in probabilistic artificial intelligence, whose contributions are usually expressed in a complicated mathematical language. This difficulty may be overcome if the researchers in AI and

planning make an effort to learn the language of health professionals, to attend conferences on medical decision making, and to publish papers and books that can be understood by doctors and health economists. This has been one of the priorities of our research group in the last 20 years: to bridge the gap between both fields. We try to bring to AI the most pressing problems faced by medical decision analysts (and to solve them, within the limits of our possibilities), and vice versa, to explain the new developments of AI in a language that health researchers can understand.

However, there are some challenges that are to be addressed before MDPs are widely accepted by the medical decision making community.

4 Challenges for MDPs in medicine

4.1 Graphical user interfaces

Most practitioners of POMDPs are computer scientists and engineers, used to building programs in C++ or Java. However, for a health professional it is very difficult, if not impossible, to use any of the open-source programs currently available. Clearly, it is essential to develop software tools with friendly graphical user interfaces (GUIs) that permit users with only a basic knowledge of probability and decision theory to build MDPs, in the same way that TreeAge permits to build decision trees and other tools permit to build Bayesian networks and influence diagrams with relative ease.

In our group we are developing an open source software tool, OpenMarkov (previously called Carmen [Arias and Díez, 2008]), that can be used to build and evaluate Bayesian networks, influence diagrams, SMMs, FOMDPs, POMDPs, etc. It is written in Java, so that it can run on any platform. Figure 1 shows a SMM for the human papilloma virus (HPV) vaccine that we are building in collaboration with the health agency of the regional government of Madrid. Our first purpose is to reproduce the results of the equivalent model that this agency has built using Excel. Then we will turn this SMM into a POMDP and evaluate it to find the optimal policy.

OpenMarkov may also be useful for building POMDPs in other domains, such as robotics, because—to our knowledge—there is no tool offering a GUI for building and editing those models.

4.2 Non-stationary models

As mentioned above, one of the main requirements of medical Markovian models is the possibility of introducing time-dependent transitions. The approach used in our tool is to encode in the model transition functions that depend on temporal variables. For example, the model in Figure 1 indicates that the probabilities of receiving a gynecologic screening, getting infected by the HPV, dying of cervical cancer (in the case of being infected), and dying of other causes depend on the age of the patient. Also the quality of life, represented by a utility node, depends on the age.

Another facility that OpenMarkov will offer is a potential called *WeibullDistribution*, aimed at representing the probability that a hazardous event Y (for example, the death of the

²www.treeage.com




Figure 1: A model for the HPV vaccine, as shown in OpenMarkov.

patient) occurs between time $t - 1$ and t ; the probability density, $f(t|\mathbf{x})$, depends on a temporal variable—such as the age of the patient or the “age” of a prosthesis—and on a set of other variables, $\mathbf{X}$, which represent the risk factors for Y .

4.3 Cost-effectiveness analysis

In medicine there is often a compromise between the health benefits (effectiveness) of an intervention and its cost. One approach used in many studies, including those based on Markovian models, is to try to maximize the effectiveness, disregarding the cost. However, there is an increasing awareness that resources are limited, and the need to optimize them has made mandatory to also pay attention to the economic costs. The problem is that these two dimensions, effectiveness and cost, are so different in nature that is not easy to combine them in a single scale. For this reason, it is more and more common to perform **cost-effectiveness analyses**, whose output consists of indicating what is the optimal intervention depending on the decisor’s *willingness to pay*, which is assumed to be unknown [Drummond *et al.*, 2005; Gold *et al.*, 1996].

In OpenMarkov, cost-effectiveness analysis is based on the distinction between two types of nodes. For example, the model shown in Figure 1 contains, for each time slice, a node that represents the quality of life (a measure of effectiveness) and three nodes that represent the cost of the vaccine, the cost of screening, and the cost of treatment. A way of evaluating the POMDP is to combine the cost and effectiveness into the net health value [Stinnett and Mullahy, 1998], an approach that has already been used in [Zhang *et al.*, 2010]. Another possibility is to develop a cost-effectiveness algorithm for POMDPs, similar to the one proposed in [Arias and Díez,

2011].

4.4 Sensitivity analysis

When building qualitative models, it is usually difficult to estimate with precision the parameters of the model, unless we have a lot of data. In some fields, such as robotics, it is not necessary to find the exact values of the parameters: it suffices to obtain a good approximation that allows the robot to achieve its objective with a near optimal use of the resources. In contrast, in medicine there are ethical and legal reasons that make it necessary to make sure, as much as possible, that the uncertainty in the model does not affect the conclusions of the analysis. For this reason, **sensitivity analysis** is compulsory in medical decision analysis [Claxton *et al.*, 2005]. Surprisingly, this issue has received little attention when applying artificial intelligence to medical problems.

In OpenMarkov the uncertainty about the parameters is given by second order probabilities; the format ProbModelXML [Arias *et al.*, 2011], developed by our group as the default format for OpenMarkov, determines which probability distributions can be used in each case. Thus, the second order probability distribution for a probabilistic parameters can be a uniform, a triangular, a Beta, or a Dirichlet. In the case of utility parameters, the candidate distributions are the normal, the logNormal, and the Gamma. If the transition probability for a variable is a Weibull distribution, the second-order distribution can be given by a covariance matrix; OpenMarkov will generate a second-order probability distribution and use it to sample new potential values for sensitivity analysis [Briggs *et al.*, 2006]. This process is based on the Cholesky decomposition of the covariance matrix, but the user does not need to care about it: his/her efforts can concentrate on building the

model and interpreting the results.

With this tool we are building a Markov model for analyzing the cost-effectiveness of the human papilloma virus (HPV) vaccine. Our first goal is to reproduce the results obtained from an Excel model which evaluated only a limited set of policies, pre-selected by its designers. Then, we will turn that model into a POMDP; in order to convert this bi-criteria model into a unicriterion problem,

5 Conclusion

In this paper we have seen that the Markov models typically used in medicine, that we have called *simple Markov models*, have very limited expressive power. MDPs have several advantages over SMM, but they have been the exception rather than the rule in medical Markov models. In recent years the number of FOMDPs and POMDPs built for medical problems is increasing exponentially fast (see the publication years for the references in the introduction), but there are still some challenges that must be solved before MDPs are widely accepted as a tool for medical decision analysis. We have shown how our tool OpenMarkov is addressing those challenges, in an attempt to make MDPs affordable to health professionals and medical decision makers. We are currently using this tool to build a Markov model for analyzing the cost-effectiveness of the HPV vaccine.

Acknowledgments

This work has been supported by grants TIN2006-11152 and TIN2009-09158, of the Spanish Ministry of Science and Technology, and by FONCICYT grant 85195. We thank the anonymous reviewers for many useful comments.

References

- [Ahn and Hornberger, 1996] J. H. Ahn and J. C. Hornberger. Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective. *Management Science*, 42:629–641, 1996.
- [Alagoz et al., 2007] O. Alagoz, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. Choosing among living-donor and cadaveric livers. *Management Science*, 53:1702–1715, 2007.
- [Alagoz et al., 2010] O. Alagoz, H. Hsu, A. J. Schaefer, and M. S. Roberts. Markov Decision Processes: A tool for sequential decision making under uncertainty. *Medical Decision Making*, 30:474–483, 2010.
- [Arias and Díez, 2008] M. Arias and F. J. Díez. Carmen: An open source project for probabilistic graphical models. In *Proceedings of the Fourth European Workshop on Probabilistic Graphical Models (PGM’08)*, pages 25–32, Hirtshals, Denmark, 2008.
- [Arias and Díez, 2011] M. Arias and F. J. Díez. Cost-effectiveness analysis with sequential decisions. Technical Report CISIAD-11-01, UNED, Madrid, Spain, 2011.
- [Arias et al., 2011] M. Arias, F. J. Díez, and M. P. Palacios. Probmodelxml. a format for encoding probabilistic graphical models. Technical Report CISIAD-11-02, UNED, Madrid, Spain, 2011.
- [Åström, 1965] K. J. Åström. Optimal control of Markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965.
- [Ayer et al., 2010] T. Ayer, O. Alagoz, N. K. Stout, and E. S. Burnside. Designing a new breast cancer screening program considering adherence. In *INFORMS Annual Meeting*, Austin, TX, 2010.
- [Beck and Pauker, 1983] J. R. Beck and S. G. Pauker. The Markov process in medical prognosis. *Medical Decision Making*, 3:419–458, 1983.
- [Bellman, 1957] R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- [Boutilier and Poole, 1996] C. Boutilier and D. Poole. Computing optimal policies for partially observable decision processes using compact representations. In W. J. Clancey and D. Weld, editors, *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pages 1168–1175, Portland, OR, 1996. AAAI Press.
- [Boutilier et al., 1995] C. Boutilier, R. Dearden, and M. Goldszmidt. Exploiting structure in policy construction. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 1104–1111, Montreal, Canada, 1995.
- [Boutilier et al., 2000] C. Boutilier, R. Dearden, and M. Goldszmidt. Stochastic dynamic programming with factored representations. *Artificial Intelligence*, 121(1-2):49–107, 2000.
- [Briggs et al., 2006] A. Briggs, K. Claxton, and M. Sculpher. *Decision Modelling for Health Economic Evaluation*. Oxford University Press, New York, 2006.
- [Claxton et al., 2005] K. Claxton, M. Sculpher, C. McCabe, et al. Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. *Health Economics*, 14:339–347, 2005.
- [Dean and Kanazawa, 1989] T. Dean and K. Kanazawa. A model for reasoning about persistence and causation. *Computational Intelligence*, 5:142–150, 1989.
- [Denton et al., 2009] B. T. Denton, M. Kurt, N. D. Shah, S. C. Bryant, and S. A. Smith. Optimizing the start time of statin therapy for patients with diabetes. *Medical Decision Making*, 29:351–367, 2009.
- [Díez and Druzdzel, 2006] F. J. Díez and M. J. Druzdzel. Canonical probabilistic models for knowledge engineering. Technical Report CISIAD-06-01, UNED, Madrid, Spain, 2006.
- [Drummond et al., 2005] M. F. Drummond, M. J. Sculpher, G. W. Torrance, B. J. O’Brien, and G. L. Stoddart. *Methods for the Economic Evaluation of Health Care Programmes*. Oxford University Press, third edition, 2005.
- [Fard et al., 2008] M. M. Fard, J. Pineau, and P. Sun. A variance analysis for POMDP policy evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1056–1061, Chicago, IL, 2008.

- [Ghallab *et al.*, 2004] M. Ghallab, D. Nau, and P. Traverso. *Automated Planning: Theory and Practice*. Morgan Kaufmann, San Francisco, CA, 2004.
- [Gold *et al.*, 1996] M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 1996.
- [Goulionis and Vozikis, 2009] J. E. Goulionis and A. Vozikis. Medical decision making for patients with Parkinson disease under Average Cost Criterion. *Australia and New Zealand Health Policy*, 6:15, 2009.
- [Goulionis *et al.*, 2008] J. E. Goulionis, B. K. Koutsoumaris, and D. J. Stengos. Dynamic decision making for clinical patient management with Parkinson's disease. *JP Journal of Biostatistics*, 2:217–235, 2008.
- [Goulionis *et al.*, 2010] J. E. Goulionis, A. Vozikis, V. K. Benos, and D. Nikolakis. On the decision rules of cost-effective treatment for patients with diabetic foot syndrome. *ClinicoEconomics and Outcomes Research*, 2:121–126, 2010.
- [Goulionis, 2009] J. E. Goulionis. Medical treatments using partially observable Markov decision processes. *JP Journal of Biostatistics*, 3:77–97, 2009.
- [Gray *et al.*, 2010] A. M. Gray, P. M. Clarke, J. Wolstenholme, and S. Wordsworth. *Applied Methods of Cost-effectiveness Analysis in Healthcare*. Oxford University Press, New York, 2010.
- [Haddix *et al.*, 2003] A. C. Haddix, S. M. Teutsch, and P. S. Corso. *Prevention Effectiveness. A Guide to Decision Analysis and Economic Evaluation*. Oxford University Press, New York, second edition, 2003.
- [Hauskrecht and Fraser, 2000] M. Hauskrecht and H. Fraser. Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artificial Intelligence in Medicine*, 18:221–244, 2000.
- [Hawkins *et al.*, 2005] N. Hawkins, M. Sculpher, and D. Epstein. Cost-effectiveness analysis of treatments for chronic disease: Using R to incorporate time dependency of treatment response. *Medical Decision Making*, 25:511–519, 2005.
- [He *et al.*, 2010] M. He, L. Zhao, and W. B. Powell. Optimal control of dosage decisions in controlled ovarian hyperstimulation. *Annals of Operations Research*, 178:223–245, 2010.
- [Hoey *et al.*, 1999] J. Hoey, R. St-Aubin, A. Hu, and C. Boutilier. SPUDD: Stochastic planning using decision diagrams. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI'99)*, pages 279–288, Stockholm, Sweden, 1999. Morgan Kaufmann, San Francisco, CA.
- [Hoey *et al.*, 2008] J. Hoey, P. Poupart, C. Boutilier, and A. Mihailidis. POMDP models for assistive technology. In *AAAI 2008 Fall Symposium on AI in Eldercare*, Washington, DC, 2008.
- [Hu *et al.*, 1996] C. Hu, W. S. Lovejoy, and S. L. Shafer. Comparison of some suboptimal control policies in medical drug therapy. *Operations Research*, 44:696–709, 1996.
- [Hunink *et al.*, 2001] M. Hunink, P. Glasziou, J. Siegel, J. Weeks, J. Pliskin, A. Elstein, and M. Weinstein. *Decision Making in Health and Medicine*. Cambridge University Press, Cambridge, UK, 2001.
- [Izadi and Buckeridge, 2007] M. T. Izadi and D. L. Buckeridge. Optimizing anthrax outbreak detection methods using reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1781–1786, Vancouver, BC, 2007.
- [Kan *et al.*, 2008] P. Kan, J. Hoey, and A. Mihailidis. Automated upper extremity rehabilitation for stroke patients using a partially observable Markov decision process. In *AAAI 2008 Fall Symposium on AI in Eldercare*, Washington, DC, 2008.
- [Kirkızlar *et al.*, 2010] E. Kirkızlar, D. M. Faissol, P. M. Griffin, and J. L. Swann. Timing of testing and treatment for asymptomatic diseases. *Mathematical Biosciences*, 226:28–37, 2010.
- [Kreke *et al.*, 2008] J. E. Kreke, M. D. Bailey, A. J. Schaefer, M. S. Roberts, and D. C. Angus. Modeling hospital discharge policies for patients with pneumonia-related sepsis. *IEEE Transactions*, 40:853–860, 2008.
- [Kuntz and Weinstein, 2001] K. M. Kuntz and M. C. Weinstein. Modelling in economic evaluation. In M. F. Drummond and A. McGuire, editors, *Economic Evaluation in Health Care*, chapter 7, pages 141–171. Oxford University Press, New York, 2001.
- [Lefebvre, 1981] C. Lefebvre. Optimal control of a birth and death epidemic process. *Operations Research*, 29:971–982, 1981.
- [Leshno *et al.*, 2003] M. Leshno, Z. Halpern, and N. Arber. Cost-effectiveness of colorectal cancer screening in the average risk population. *Health Care Management Science*, 6:165–174, 2003.
- [Levin and McEwan, 2001] H. M. Levin and P. J. McEwan. *Cost-Effectiveness Analysis*. Sage Publications, Thousand Oaks, CA, second edition, 2001.
- [Magni *et al.*, 2000] P. Magni, S. Quaglini, M. Marchetti, and G. Barosi. Deciding when to intervene: A Markov decision process approach. *International Journal of Medical Informatics*, 60:237–253, 2000.
- [Maillart *et al.*, 2008] L. M. Maillart, J. S. Ivy, S. Ransom, and K. Diehl. Assessing dynamic breast cancer screening policies. *Operations Research*, 56:1411–1427, 2008.
- [Markov, 1906] A. A. Markov. Rasprostranenie zakona bol'shih chisel na velichiny, zavisyaschie drug ot druga. *Izvestiya Fiziko-matematicheskogo obschestva pri Kazanskom universitete*, 2-ya seriya, 15:135–156, 1906.
- [Mason *et al.*, 2011] J. E. Mason, D. A. England, B. T. Denton, et al. Optimizing statin treatment decisions for diabetes patients in the presence of uncertain future adherence. *Medical Decision Making*, In press, 2011.

- [Muenning, 2008] P. Muenning. *Cost-Effectiveness Analysis in Health*. Jossey-Bass, San Francisco, CA, second edition, 2008.
- [Murphy, 2002] K. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, Computer Science Division, University of California, Berkeley, 2002.
- [Parmigiani, 2002] G. Parmigiani. *Modeling in Medical Decision Making*. John Wiley and Sons, Chichester, UK, second edition, 2002.
- [Pearl, 1988] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [Pearl, 2000] J. Pearl. *Causality. Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.
- [Petitti, 2000] D. B. Petitti. *Meta-Analysis, Decision Analysis, and Cost-Effectiveness Analysis*. Oxford University Press, New York, second edition, 2000.
- [Pineau *et al.*, 2009] J. Pineau, A. Guez, R. Vincent, G. Panuccio, and M. Avoli. Treating epilepsy via adaptive neurostimulation: A reinforcement learning approach. *International Journal of Neural Systems*, 19:227–240, 2009.
- [Poupart, 2005] P. Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. PhD thesis, Dept. of Computer Science, University of Toronto, Canada, 2005.
- [Schaefer *et al.*, 2004] A. J. Schaefer, M. D. Bailey, S. M. Shechter, and M. S. Roberts. Modeling medical treatment using Markov decision processes. In M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, editors, *Operations Research and Health Care: A Handbook of Methods and Applications*, pages 593–612. Kluwer Academic Publishers, Boston, MA, 2004.
- [Shechter *et al.*, 2008] S. M. Shechter, M. D. Bailey, A. J. Schaefer, and M. S. Roberts. The optimal time to initiate HIV therapy under ordered health states. *Operations Research*, 56:20–33, 2008.
- [Sloan, 1995] F. A. Sloan, editor. *Valuing Health Care*. Cambridge University Press, Cambridge, MA, 1995.
- [Sonnenberg and Beck, 1993] F. A. Sonnenberg and J. R. Beck. Markov models in medical decision making: A practical guide. *Medical Decision Making*, 13:322–339, 1993.
- [Sox *et al.*, 1988] H. C. Sox, M. A. Blatt, M. C. Higgins, and K. I. Marton. *Medical Decision Making*. Butterworth-Heinemann, Woburn, MA, 1988.
- [Spaan and Vlassis, 2005] M. T. J. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.
- [Stinnett and Mullahy, 1998] A. A. Stinnett and J. Mullahy. Net health benefit: A new framework for the analysis of uncertainty in cost-effectiveness analysis. *Medical Decision Making*, 18:S68–S80, 1998.
- [TreeAge Software, Inc., 2009] TreeAge Software, Inc. *TreAge Pro User's Manual*. TreeAge, Williamstown, MA, 2009.
- [Tusch, 2000] G. Tusch. Optimal sequential decisions in liver transplantation based on a POMDP model. In *Proceedings of the European Conference on Artificial Intelligence*, pages 186–190, Berlin, Germany, 2000.
- [Zhang *et al.*, 2010] J. Zhang, B. T. Denton, H. Balasubramanian, N. D. Shah, and B. A. Inman. Optimization of PSA screening policies: A comparison of the patient and societal perspectives. Submitted to *Medical Decision Making*, 2010.
- [Zhang *et al.*, 2011] J. Zhang, J. E. Mason, B. T. Denton, and W. P. Pierskalla. Applications of operations research to the prevention, detection, and treatment of disease. In J. J. Cochran *et al.*, editors, *Wiley Encyclopedia of Operations Research and Management*. Wiley, 2011. In press.