

Robust Constrained-MDPs: Soft-Constrained Robust Policy Optimization under Model Uncertainty

Russel, Reazul Hasan; Benosman, Mouhacine; van Baar, Jeroen

TR2021-009 March 11, 2021

Abstract

In this paper, we focus on the problem of robustifying reinforcement learning (RL) algorithms with respect to model uncertainties. Indeed, in the framework of model-based RL, we propose to merge the theory of constrained Markov decision process (CMDP), with the theory of robust Markov decision process (RMDP), leading to a formulation of robust constrained-MDPs (RCMDP). This formulation, simple in essence, allows us to design RL algorithms that are robust in performance, and provides constraint satisfaction guarantees, with respect to uncertainties in the system's states transition probabilities. The need for RCMPDs is important for real-life applications of RL. For instance, such formulation can play an important role for policy transfer from simulation to real world (Sim2Real) in safety critical applications, which would benefit from performance and safety guarantees which are robust w.r.t model uncertainty. We first propose the general problem formulation under the concept of RCMDP, and then propose a Lagrangian formulation of the optimal problem, leading to a robust-constrained policy gradient RL algorithm. We finally validate this concept on the inventory management problem.

Advances in Neural Information Processing Systems (NeurIPS)-workshop

© 2021 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Robust Constrained-MDPs: Soft-Constrained Robust Policy Optimization under Model Uncertainty

Reazul Hasan Russel* Mouhacine Benosman Jeroen Van Baar
Mitsubishi Electric Research Laboratories (MERL)
Cambridge, MA 02139, USA
{rrussel, benosman, jeoren}@merl.com

Abstract

In this paper, we focus on the problem of robustifying reinforcement learning (RL) algorithms with respect to model uncertainties. Indeed, in the framework of model-based RL, we propose to merge the theory of constrained Markov decision process (CMDP), with the theory of robust Markov decision process (RMDP), leading to a formulation of robust constrained-MDPs (RCMDP). This formulation, simple in essence, allows us to design RL algorithms that are robust in performance, and provides constraint satisfaction guarantees, with respect to uncertainties in the system's states transition probabilities. The need for RCMDPs is important for real-life applications of RL. For instance, such formulation can play an important role for policy transfer from simulation to real world (Sim2Real) in safety critical applications, which would benefit from performance and safety guarantees which are robust w.r.t model uncertainty. We first propose the general problem formulation under the concept of RCMDP, and then propose a Lagrangian formulation of the optimal problem, leading to a robust-constrained policy gradient RL algorithm. We finally validate this concept on the inventory management problem.

1 Introduction

Reinforcement learning (RL) is a learning framework that addresses sequential decision-making problems, wherein an 'agent' or a decision maker learns a policy to optimize a long-term reward by interacting with the (unknown or partially known) environment. At each step, the RL agent obtains evaluative feedback (called reward or cost) about the performance of its action, allowing it to improve the performance of subsequent actions (Sutton & Barto, 1998). With the advent of deep learning, RL has witnessed huge successes in recent times (Silver et al., 2017). However, since most of these methods rely on model-free RL, there are several unsolved challenges, which restrict the use of these algorithms for many safety critical physical systems (Vamvoudakis et al., 2015; Benosman, 2018). For example, it is very difficult for most model-free RL algorithms to ensure basic properties like stability of solutions, robustness with respect to model uncertainties, etc. This has led to several research directions which study incorporating robustness, constraint satisfaction, and safe exploration during learning for safety critical applications. While safe exploration and robust stability guarantees are highly desirable, they are also very challenging to incorporate in RL algorithms. The main goal of our work is to formulate this incorporation into robust constrained-MDPs (RCMDPs), and derive the corresponding equations necessary to solve problems on RCMDPs.

Constrained Markov decision processes (CMDPs) can be seen as an extension of MDPs with expected cumulative cost constraints, e.g., (Altman, 2004). For such CMDPs, several solution methods have been proposed, e.g., linear programming-based solutions (Altman, 2004), surrogate-based methods (Chamiea et al., 2016; Dalal et al., 2018), Lagrangian methods (Geibel & Wysotzki,

* Russel is with the University of New Hampshire, and is spending an internship period at MERL. Preprint. Under review.

2005; Altman, 2004). We refer to these CMDPS as *non-robust* since they do not take uncertainties in the state transition probability into account, which is an important factor in real-life applications.

On the other hand, robustification of MDPs, w.r.t. model uncertainties, can be found in the context of robust MDPs (RMDPs) which generalize MDPs to the case where transition probabilities and/or rewards are not perfectly known, e.g., Nilim & Ghaoui (2004); Wiesemann et al. (2013). These RMDPs can be formulated and solved using so-called ambiguity or uncertainty sets, e.g., (Petrik & Russell, 2019; Petrik, 2012; Petrik & Luss, 2016). However, one noticeable point in all the RMDPs-based RL algorithms is the fact that they do not consider any safety constraints, i.e., expected cumulative cost constraints.

These safety constraints are important in real-life applications, where one cannot afford to risk violating some given constraints, e.g., in autonomous cars, there are hard safety constraints on the robot velocities and steering angles. Besides, in real applications, to mitigate the sample inefficiency of model-free RL algorithms, training often occurs on a simulated environment. The result is then transferred to the real world, typically followed by fine-tuning, a process referred to as Sim2Real. The simulation model is by definition uncertain with respect to the real world, due to approximations and lack of system identification. Domain randomization (van Baar et al., 2019) and meta-learning (Finn et al., 2017), aimed at addressing model uncertainty in transfer, offer no guarantees. Furthermore, for safety critical applications, a trained policy in simulation should offer certain guarantees on safety when transferred to the real world.

In light of these practical motivations, we propose to merge the two concepts of CMDPs and RMDPs, to ensure both safety and robustness. In this RCM DP concept, we propose to robustify both the performance cost minimization, as well as the safety constraints (via cumulative constraint costs). Indeed, robustness is equally (if not more) important in estimating the cumulative constraint costs along the whole trajectory in order to certify that there will be no unexpected violations, if the system is deployed in reality. That is, if deployed, the worst-case cumulative constrained-cost will not exceed a pre-determined safety budget.

The contribution of this paper is four-fold: 1) Intuited from the concepts of CMDP and RMDP, we formulate the concept of RCM DP, bridging the gap between constraints and robustness w.r.t. state transition probability uncertainties; 2) propose a robust soft-constrained Lagrange-based solution of the RCM DP problem; 3) derive associated gradient update rule and present a policy gradient algorithm; 4) illustrate the performance of the proposed algorithm on the inventory management problem under model uncertainties.

After our submission of this paper in the RWRL workshop and an earlier version in arXiv, we became aware of a recent arXiv paper (Mankowitz et al., 2020) which proposes similar robust constrained formulation for MDPs. They derive robust constraint bellman operator and present an actor-critic class of algorithm. Our focus in this paper remains within policy gradient class of algorithm and we handle robustness with explicitly constructed ambiguity sets for each state-action.

The paper is organized as follows: Section 2 describes the formulation of our Robust-CMDP problem and the objective we seek to optimize. A Lagrange based approach is presented in Section 3 along with required gradient update rules and a robust constrained policy gradient algorithm. We evaluate our algorithm in Section 4 and draw the concluding remarks in Section 5.

2 Problem Formulation

We consider a robust-MDP model with a finite number of states $\mathcal{S} = \{1, \dots, S\}$ and finite number of actions $\mathcal{A} = \{1, \dots, A\}$. Every action $a \in \mathcal{A}$ is available for the decision maker to take in every state $s \in \mathcal{S}$. After taking an action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$, the decision maker receives a cost $c(s, a) \in \mathbb{R}$ and transitions to a next state s' according to the *true* but *unknown* transition probability $p_{s,a}^* \in \Delta^{\mathcal{S}}$. p_0 is the distribution of an initial state. We further incorporate a constrained-MDP (Altman, 2004) setup into this robust-MDP model by introducing a constraint cost $d(s) \in [0, D_{\max}]$ and an associated constraint budget $d_0 \in \mathbb{R}_+$.

An ambiguity set $\mathcal{P}_{s,a}$, defined for each state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$, is a set of feasible transition matrices quantifying the uncertainty in transition probabilities. In this paper, we restrict our attention to s, a -rectangular ambiguity sets which simply assumes independence between different state-action pairs (Le Tallec, 2007; Wiesemann et al., 2013). We use in this paper L_1 -norm bounded ambiguity

sets around the nominal transition probability $\bar{p}_{s,a} = \mathbb{E}[p_{s,a}^* | \mathcal{D}]$, on some dataset \mathcal{D} , as:

$$\mathcal{P}_{s,a} = \{p \in \Delta^S : \|p - \bar{p}_{s,a}\|_1 \leq \psi_{s,a}\}$$

Where $\psi_{s,a} \geq 0$ is the budget of allowed deviations. This budget ψ can be computed using Hoeffding bound as (Russel & Petrik, 2019): $\psi_{s,a} = \sqrt{\frac{2}{n_{s,a}} \log \frac{SA2^S}{\delta}}$, where $n_{s,a}$ is the number of transitions in dataset \mathcal{D} originating from state s and an action a , and δ is the confidence level. Note that this is just one specific choice for the ambiguity set. Our method can be extended to any other type of ambiguity sets (e.g. L_∞ -norm, Bayesian, weighted etc.). We use \mathcal{P} to refer cumulatively to $\mathcal{P}_{s,a}$ for all states $s \in \mathcal{S}$ and actions $a \in \mathcal{A}$.

A stationary randomized policy $\pi(\cdot|s)$ for state $s \in \mathcal{S}$ defines a probability distribution over actions $a \in \mathcal{A}$, Π represents the set of all stationary randomized policies. We parameterize the randomized policy for state $s \in \mathcal{S}$ as $\pi(\cdot|s; \theta)$ where $\theta \subseteq \mathbb{R}^k$ is a k -dimensional parameter vector. Let $\xi = \{s_0, a_0, c_0, d_0, \dots, s_{T-1}, a_{T-1}, c_{T-1}, d_{T-1}, s_T\}$ be a sampled trajectory generated by executing a policy π from a starting state s_0 . Then the probability of sampling ξ is: $p^\theta(\xi) = p_0(s_0) \prod_{t=0}^{T-1} \pi(a_t|s_t; \theta) p(s_{t+1}|s_t, a_t)$. The total cost g for trajectory ξ is: $g(\xi) = \sum_{t=0}^{\infty} \gamma^t c(s_t, a_t)$ (Puterman, 2005). The value function is defined as the expected return: $v^\theta(s_0) = \mathbb{E}_p[g(\xi)]$. We define the robust value function \hat{v} as the expected return in the worst-case realization of the transition probability within \mathcal{P} as: $\hat{v}_\pi^\pi(s_0) = \max_{p \in \mathcal{P}} \mathbb{E}_p[g(\xi)]$. Similarly, the total constraint-cost for trajectory ξ is: $h(\xi) = \sum_{t=0}^{\infty} \gamma^t d(s_t, a_t)$. And the robust constraint value function \hat{u} is defined as: $\hat{u}_\pi^\pi(s_0) = \max_{p \in \mathcal{P}} \mathbb{E}_p[h(\xi)]$.

The robust Bellman operator $\hat{T}_\mathcal{P}$ for a state $s \in \mathcal{S}$ and an ambiguity set \mathcal{P} computes the best action with respect to the worst-case realization of the transition probabilities in \mathcal{P} as:

$$(\hat{T}_\mathcal{P}v)(s) := \min_{a \in \mathcal{A}} \max_{p \in \mathcal{P}_{s,a}} (c(s, a) + \gamma \cdot p^\top v)$$

The optimal robust value function \hat{v}^* , and the robust value function \hat{v}^π for a policy π are unique and satisfy $\hat{v}^* = \hat{T}_\mathcal{P}\hat{v}^*$ and $\hat{v}^\pi = \hat{T}_\mathcal{P}^\pi\hat{v}^\pi$ (Iyengar, 2005). Similarly, all these properties hold for the constrained robust value function \hat{u} as well.

Objective Our objective is then to solve the RCMDP optimization problem below:

$$\begin{aligned} \min_{\pi \in \Pi} \hat{v}_\pi^\pi(s) \\ \text{s.t. } \hat{u}_\pi^\pi(s) \leq d_0 \end{aligned} \quad (1)$$

This objective resembles the objective of a CMDP, but with additional robustness integrated by the quantification of the uncertainty in the model.

3 Robust Constrained Optimization

A general approach for solving Equation (1) is to apply the Lagrange relaxation procedure (Chapter 3 of Bertsekas (2003)), which turns it into an unconstrained optimization problem:

$$\max_{\lambda \geq 0} \min_{\theta} \left(L(\theta, \lambda) = \hat{v}_\pi^\pi(s) + \lambda \left(\hat{u}_\pi^\pi(s) - d_0 \right) \right) \quad (2)$$

where λ is known as the *Lagrange multiplier*. The goal is then to find a saddle point (θ^*, λ^*) that satisfies $L(\theta, \lambda^*) \geq L(\theta^*, \lambda^*) \geq L(\theta^*, \lambda)$, $\forall \theta, \lambda$. This is achieved by descending in θ and ascending in λ using the gradients.

Without loss of generality, we rewrite Equation (2) for a fixed starting state s_0 and perform some algebraic manipulation:

$$\begin{aligned}
L(\theta, \lambda) &= \hat{v}_{\mathcal{P}}^{\pi}(s_0) + \lambda \left(\hat{u}_{\mathcal{P}}^{\pi}(s_0) - d_0 \right) \\
&= \max_{p \in \mathcal{P}} \mathbb{E}_p [g(\xi)] + \lambda \left(\max_{p \in \mathcal{P}} \mathbb{E}_p [h(\xi)] - d_0 \right) \\
&\stackrel{(a)}{=} \mathbb{E}_{\hat{p}_v^\theta} [g(\xi)] + \lambda \mathbb{E}_{\hat{p}_u^\theta} [h(\xi)] - \lambda d_0 \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) + \lambda \hat{p}_u^\theta(\xi) h(\xi) \right) - \lambda d_0
\end{aligned}$$

Here (a) follows with $\hat{p}_v^\theta = \operatorname{argmax}_{p \in \mathcal{P}} \mathbb{E}_p [g(\xi)]$ and $\hat{p}_u^\theta = \operatorname{argmax}_{p \in \mathcal{P}} \mathbb{E}_p [h(\xi)]$.

3.1 Gradient Update Rules

We now derive the gradient update rules with respect to θ as below:

$$\begin{aligned}
\nabla_{\theta} L(\theta, \lambda) &= \sum_{\xi} \left(\nabla_{\theta} \hat{p}_v^\theta(\xi) g(\xi) + \lambda \nabla_{\theta} \hat{p}_u^\theta(\xi) h(\xi) \right) \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) \nabla_{\theta} \log \hat{p}_v^\theta(\xi) + \lambda \hat{p}_u^\theta(\xi) h(\xi) \nabla_{\theta} \log \hat{p}_u^\theta(\xi) \right) \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) \nabla_{\theta} \log \left(p_0(s_0) \prod_{t=0}^{T-1} \hat{p}_v^\theta(s_{t+1}|s_t, a_t) \pi_{\theta}(a_t|s_t) \right) \right. \\
&\quad \left. + \lambda \hat{p}_u^\theta(\xi) h(\xi) \nabla_{\theta} \log \left(p_0(s_0) \prod_{t=0}^{T-1} \hat{p}_u^\theta(s_{t+1}|s_t, a_t) \pi_{\theta}(a_t|s_t) \right) \right) \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) \nabla_{\theta} \left(\log p_0(s_0) + \sum_{t=0}^{T-1} \log \hat{p}_v^\theta(s_{t+1}|s_t, a_t) + \log \pi_{\theta}(a_t|s_t) \right) \right. \\
&\quad \left. + \lambda \hat{p}_u^\theta(\xi) h(\xi) \nabla_{\theta} \left(\log p_0(s_0) + \sum_{t=0}^{T-1} \log \hat{p}_u^\theta(s_{t+1}|s_t, a_t) + \log \pi_{\theta}(a_t|s_t) \right) \right) \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) \sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) + \lambda \hat{p}_u^\theta(\xi) h(\xi) \sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) \right) \\
&= \sum_{\xi} \left(\hat{p}_v^\theta(\xi) g(\xi) + \lambda \hat{p}_u^\theta(\xi) h(\xi) \right) \sum_{t=0}^{T-1} \frac{\nabla_{\theta} \pi_{\theta}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)}
\end{aligned}$$

Notice that the constraint budget d_0 does not play any role in the policy optimization. Also, the expectations over the cost $g(\xi)$ and constraint cost $h(\xi)$ are with respect to $\hat{p}_v^\theta(\xi)$ and $\hat{p}_u^\theta(\xi)$, respectively. However, the costs and constraint costs are coupled together in reality, meaning that the two trajectories would not diverge. So one of $\hat{p}_v^\theta(\xi)$ or $\hat{p}_u^\theta(\xi)$ can be chosen depending on the priorities toward robustness of cost or constraint cost, and both of the expectations can be evaluated with that common probability measure $\hat{p}^\theta(\xi)$. The gradient update rule then becomes:

$$\nabla_{\theta} L(\theta, \lambda) = \sum_{\xi} \hat{p}^\theta(\xi) \left(g(\xi) + \lambda h(\xi) \right) \sum_{t=0}^{T-1} \frac{\nabla_{\theta} \pi_{\theta}(a_t|s_t)}{\pi_{\theta}(a_t|s_t)}$$

And the gradient update rule with respect to λ as below:

$$\begin{aligned}
\nabla_{\lambda} L(\theta, \lambda) &= \nabla_{\lambda} \left(\sum_{\xi} \hat{p}_v^\theta(\xi) g(\xi) + \lambda \left(\sum_{\xi} \hat{p}_u^\theta(\xi) h(\xi) - d_0 \right) \right) \\
&= \sum_{\xi} \hat{p}_u^\theta(\xi) h(\xi) - d_0
\end{aligned}$$

4 Empirical Study

In this section, we empirically study the performance of our policy gradient algorithm on several problem domains. All the experiments are run with confidence parameter $\delta = 0.9$, discount factor $\gamma = 0.99$, and $n_{s,a} = 100$ number of samples drawn for each state-action from the underlying true transition distribution $p_{s,a}^*$. For domains emitting reward signals instead of costs, we simply consider their negative magnitude keeping our cost-based formulation intact. We implement several combinations of different settings: i) non-robust unconstrained, ii) robust-unconstrained and iii) robust-constrained. We note that the non-robust unconstrained setting is the general policy gradient algorithm (Sutton & Barto, 1998). The robust-unconstrained version deals only with robustness without any constraint cost or associated budget. The robust constrained version deals with both constraints and robustness as described in Algorithm 1.

Inventory Management We evaluate the policy gradient method on the classic inventory management problem (Behzadian et al., 2019; Puterman, 2005; Zipkin, 2000). The state space is discrete and is represented by the level of inventory. The goal is to order products from a supplier in order to meet demands. Demand for a product is random and comes from a normal distribution. Figure 1 shows the return distributions for the inventory management problem. Here the non-robust unconstrained variant has slightly higher expected return compared to the robust counterparts. This is expected because the robust version particularly deals with the worst-case situation. However, the unconstrained variant is not guaranteed to enforce any constraint. The plot also shows that the robust-constrained version outperforms the robust-unconstrained variant with a higher expected return.

5 Conclusion

In this paper, we studied the problem of MDPs under constraints, and model uncertainties. We proposed to merge together the concepts of constrained MDPs and robust MDPs, leading to the concept of robust constrained MDPs (RCMDPs). Indeed, by doing so, one can take advantage of the safety guarantees given by the CMDP formulation, as well as the robustness guarantees w.r.t. model uncertainties, given by the RMDP formulation. We then proposed a robust soft-constrained Lagrange-based solution to the RCMDP problem, and a corresponding policy gradient algorithm. Next work will focus on extending the proposed approach to continuous domains, and validate the performance of this RCMDP formulation on more safety critical examples, e.g., robotics test-beds.

References

- Altman, E. *Constrained Markov Decision Processes*. 2004.
- Behzadian, B., Russel, R. H., and Petrik, M. High-Confidence Policy Optimization: Reshaping Ambiguity Sets in Robust MDPs. 2019.
- Benosman, M. Model-based vs data-driven adaptive control: An overview. *International Journal of Adaptive Control and Signal Processing*, 2018.
- Bertsekas, D. P. *Nonlinear programming*. Athena Scientific, 2003.
- Borkar, V. S. Stochastic Approximation: A Dynamical Systems Viewpoint. *International Statistical Review*, 2009.
- Chamiea, M. E., Yu, Y., and Acikmese, B. Convex synthesis of randomized policies for controlled markov chains with density safety upper bound constraints. In *IEEE American Control Conference*, pp. 6290–6295, 2016.
- Dalal, G., Dvijotham, K., Vecerik, M., Hester, T., Paduraru, C., and Tassa, Y. Safe exploration in continuous action spaces, 2018.
- Finn, C., Yu, T., Zhang, T., Abbeel, P., and Levine, S. One-shot visual imitation learning via meta-learning. 2017.
- Geibel, P. and Wysotzki, F. Risk-sensitive reinforcement learning applied to control under constraints. *Journal of Artificial Intelligence Research*, 2005.

- Iyengar, G. N. Robust dynamic programming. *Mathematics of Operations Research*, 2005.
- Le Tallec, Y. *Robust, Risk-Sensitive, and Data-driven Control of Markov Decision Processes*. PhD thesis, MIT, 2007.
- Mankowitz, D. J., Levine, N., Jeong, R., Shi, Y., Kay, J., Abdolmaleki, A., Springenberg, J. T., Mann, T., Hester, T., and Riedmiller, M. Robust Constrained Reinforcement Learning For Continuous Control With Model Misspecification. 2020.
- Nilim, A. and Ghaoui, L. E. Robust solutions to Markov decision problems with uncertain transition matrices. *Operations Research*, 53(5):780, 2004.
- Petrik, M. Approximate dynamic programming by minimizing distributionally robust bounds. In *International Conference of Machine Learning*, 2012.
- Petrik, M. and Luss, R. Interpretable Policies for Dynamic Product Recommendations. In *Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Petrik, M. and Russell, R. H. Beyond Confidence Regions: Tight Bayesian Ambiguity Sets for Robust MDPs. In *Advances in Neural Information Processing Systems*, 2019.
- Puterman, M. L. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 2005.
- Russel, R. H. and Petrik, M. Beyond confidence regions: Tight Bayesian ambiguity sets for robust MDPs. *Advances in Neural Information Processing Systems*, 2019.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 2017.
- Sutton, R. S. and Barto, A. *Reinforcement learning*. 1998.
- Vamvoudakis, K., Antsaklis, P., Dixon, W., Hespanha, J., Lewis, F., Modares, H., and Kiumarsi, B. Autonomy and machine intelligence in complex systems: A tutorial. 2015.
- van Baar, J., Sullivan, A., Corcodel, R., Jha, D., Romeres, D., and Nikovski, D. N. Sim-to-real transfer learning using robustified controllers in robotic tasks involving complex dynamics. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- Wiesemann, W., Kuhn, D., and Rustem, B. Robust Markov decision processes. *Mathematics of Operations Research*, 2013.
- Zipkin, P. H. *Foundations of Inventory Management*. 2000. ISBN 0256113793.