Combinatorial Decision Making in Complex, Uncertain, and Highly Dynamic Environments

Stefano Ermon, Yexiang Xue, Jon Conrad, Carla Gomes, Bart Selman

Dept. of Computer Science,
Cornell University, Ithaca, NY 14853

Abstract. How can we effectively support large scale combinatorial decision making in complex, uncertain, and highly dynamic environments?

Keywords: MDP, Planning, Uncertainty, Risk Management

Computational sustainability problems range from wildlife preservation and biodiversity planning [2, 6], to large scale deployment and management of renewable energy sources [1], to the design of intelligent or "smart" control systems for energy-efficient buildings, data centers, vehicles, and appliances [5, 7, 8]. A key feature of these problems is that, in general, they require the consideration and balancing of complex and dynamic environmental, economic, and social interactions over long term horizons. As data is being captured today as never before, this gives rise to computational problems that are unique in scale, complexity, and richness, which involve combinatorial decision making in highly dynamic and uncertain environments.

As a first step towards the study of complex dynamical systems under uncertainty, in [3, 4] we introduce a class of Markov Decision Processes that arise as a natural model for many *renewable* resource allocation problems. A key and unique aspect of such a resource type is the fact that its stock is constantly replenished by an intrinsic growth process that is modeled as a dynamical system affected by uncertainty. As a new approach to deal with environmental risks, we tackle the management problems in a game theoretic framework, where the optimization problem is equivalent to a dynamic game played against nature. Upon extending results from the inventory control literature, we are able to prove some structural properties of the optimal management policy and we show how to exploit this structure to speed up its computation. In [3, 4], we consider the application of our new optimization framework to several sustainability problems arising in very different domains, such as forestry and pollution management problems, and we discuss in detail its application to the sustainable management of fisheries. Our framework is applied to a model based on real-world data for the Northern Pacific Halibut marine fishery, and we obtain a policy with a guaranteed worst-case performance that is structurally very different from the one currently employed. In particular, our approach suggests the use of periodic closures of the fishery, as opposed to current practices that involve a constant proportional harvesting every year. In our simulations, we show that our policy outperforms the historical ones.

Similar techniques can be applied to problems arising in very different domains, that yet share a similar mathematical structure. For example, electric vehicles, partially or fully powered by batteries, are one of the most promising directions towards more sustainable transportation systems [7]. However, the high costs, limited capacities, and long recharge times of batteries pose a number of obstacles to their widespread adoption. Multi-battery systems that combine a standard battery with supercapacitors are currently one of the most promising ways to increase battery lifespan and reduce operating costs. Intuitively, the idea is that the battery is good at holding the charge for long times, while the supercapacitor is efficient for rapid cycles of charge and discharge and can be used as a buffer. However, their performance crucially depends on how they are designed and operated, i.e. on the strategy used to charge and discharge the capacitor. Managing such systems is closely related to renewable resources management, because there is a mix of vehicle acceleration and regenerative braking (when power can be stored in the battery system) over time, and there is a constraint on the maximal charge the capacitor can hold. Intelligent management algorithms therefore need to analyze driving behavior and vehicle dynamics in order to make informed decisions on how to allocate the power demand over time. In [5], we formalize the problem of optimizing real-time energy management of multi-battery systems as an MDP, and we propose a novel solution based on a combination of optimization, machine learning and data-mining techniques. We evaluate the performance of our intelligent energy management system on various large datasets of commuter trips crowdsourced in the United States. We show that our policy significantly outperforms the leading algorithms that were previously proposed as part of an open algorithmic challenge.

References

1. J.M. Conrad. *Resource economics.* Cambridge Univ Pr, 2010.
2. M. Crowley and D. Poole. Policy gradient planning for environmental decision making with existing simulators. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence, Special Track on Computational Sustainability and AI.(AAAI-11)*, 2011.
3. S. Ermon, J. Conrad, C. Gomes, and B. Selman. Playing games against nature: optimal policies for renewable resource allocation. *UAI*, 2010.
4. S. Ermon, J. Conrad, C. Gomes, and B. Selman. Risk-sensitive policies for sustainable renewable resource allocation. In *IJCAI*, 2011.
5. S. Ermon, Y. Xue, , C. Gomes, and B. Selman. Learning policies for battery usage optimization in electric vehicles. *ECML-PKDD*, 2012.
6. C. Gomes, W.J. Van Hoeve, and A. Sabharwal. Connections in networks: A hybrid approach. *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pages 303–307, 2008.
7. M. Sachenbacher, M. Leucker, A. Artmeier, and J. Haselmayr. Efficient energy-optimal routing for electric vehicles. In *AAAI*, 2011.
8. P. Vytelingum, T.D. Voice, S.D. Ramchurn, A. Rogers, and N.R. Jennings. Agent-based micro-storage management for the smart grid. In *AAMAS*, 2010.