Mach Learn (2013) 92:177–194
DOI 10.1007/s10994-013-5378-z

Learning policies for battery usage optimization in electric vehicles

Stefano Ermon · Yexiang Xue · Carla Gomes · Bart Selman

Received: 18 November 2012 / Accepted: 2 May 2013 / Published online: 22 May 2013

Abstract The high cost, limited capacity, and long recharge time of batteries pose a number of obstacles for the widespread adoption of electric vehicles. 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. However, their performance crucially depends on how they are designed and operated.

In this paper, we formalize the problem of optimizing real-time energy management of multi-battery systems as a stochastic planning problem, 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. Further, we show how to extend our approach to an incremental learning setting, where the policy is capable of improving and adapting as new data is being collected over time.

Keywords Reinforcement learning · Computational sustainability · Regression

Editors: Tijl De Bie and Peter Flach.

S. Ermon (✉) · Y. Xue · C. Gomes · B. Selman
Department of Computer Science, Cornell University, Ithaca, 14850 NY, USA
e-mail: ermonste@cs.cornell.edu

Y. Xue
e-mail: yexiang@cs.cornell.edu

C. Gomes
e-mail: gomes@cs.cornell.edu

B. Selman
e-mail: selman@cs.cornell.edu