Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA
Published: April 8, 2025
DOI: 10.46254/AF6.202500400040

Stochastic Gradients via Conditional Monte Carlo and The Generalized Likelihood Ratio Method

Michael C. Fu
Robert H. Smith School of Business & Institute for Systems Research
University of Maryland
College Park, Maryland, 20742
USA
mfu@umd.edu

Abstract

Stochastic gradients are used in the training and optimization of many algorithms and techniques used in modern artificial intelligence (AI) and machine learning (ML), e.g., playing a central role in stochastic gradient descent used for training deep neural networks. Before such AI/ML applications, these techniques have been around for over half a century, arising in diverse areas of industrial engineering (IE) and operations research (OR) areas such as queueing networks, production/inventory and preventive maintenance systems, and quality control, as well as in financial applications. This talk will review the history of stochastic gradients, as well as survey the various approaches, both in the IE/OR and AI/ML contexts. More recent research has focused on discontinuous sample performance measures, for which conditional Monte Carlo and the generalized likelihood ratio method are two of the most successful approaches. IE/OR applications that require these approaches include problems that can be modeled as Markov decision processes (MDPs), for which it is often the case that optimal policies often take the form of a threshold policy. One healthcare example that will be presented has to do with optimal acceptance of a donated organ, e.g., a kidney patient with end-stage renal disease in need of a transplantation, which is modeled as an optimal stopping MDP. Another example that will be presented is estimating higher-order derivatives for queueing systems, focusing on the well-known single-server queue to serve as an illustrative example.

Keywords
stochastic gradients, conditional Monte Carlo, simulation optimization, sensitivity analysis, queueing

Biography

Michael C. Fu holds the Smith Chair of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, A. James Clark School of Engineering, all at the University of Maryland, College Park, where he has been since 1989. His research interests include stochastic gradient estimation, simulation optimization, and applied probability, with applications to manufacturing, supply chain management, and healthcare. He received a Ph.D. in applied math from Harvard in 1989 and degrees in math and EECS from MIT in 1985. He is the co-author of the books, *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, which received the INFORMS Simulation Society's 1998 Outstanding Publication Award, and *Simulation-Based Algorithms for Markov Decision Processes*, and has also edited/co-edited four volumes: *Perspectives in Operations Research*, *Advances in Mathematical Finance*, *Encyclopedia of Operations Research and Management Science* (3rd edition), and *Handbook of Simulation Optimization*. He received the INFORMS Simulation Society Distinguished Service Award in 2018, the

INFORMS Saul Gass Expository Writing Award in 2021, and the INFORMS George E. Kimball Medal in 2022. He has served as Program Director for the U.S. National Science Foundation (NSF) Operations Research Program, Stochastic Models and Simulation Department Editor for *Management Science*, and Simulation Area Editor for