

Snap ML – Accelerated Machine Learning for Big Data

Haris Pozidis

IBM Research Zurich, Switzerland

Abstract

Snap Machine Learning (Snap ML) is a new software library for training popular machine learning models, characterized by very high performance, scalability to TB-scale datasets and high resource efficiency. It continuously evolves and currently supports generalized linear models, decision trees, random forests and gradient boosting machines. Snap ML has been built to address the needs of business applications, which often have to deal with high-volume data, react fast to changing environments, and use resources efficiently to drive down cost. The high efficiency of Snap ML, in particular in dealing with big data, comes from innovations in distributed optimization, among other things. This talk will review the principles of the Snap ML library, explain how it achieves high speed and scalability, and present several cases of business workloads that demonstrate the benefits offered by Snap ML.

Haris Pozidis manages the Cloud Storage and Analytics group at IBM Research in Zurich, Switzerland. He was with Philips Research, Eindhoven, The Netherlands, before joining IBM. He has worked on read channel design for DVD and Blu-ray Disc at Philips, and played a key role in developing the first scanning probe-based data storage system at IBM, the “Millipede”. His current focus is on the development of Flash memory controllers for all-flash arrays, on phase change memory technology and system solutions, and on accelerated software libraries for machine learning. He holds over 120 US patents, has co-authored more than 120 publications, is an IBM Principal Research Scientist and Master Inventor, and a Senior Member of the IEEE.

2012 ACM Subject Classification Computing methodologies → Machine learning; Software and its engineering → Software libraries and repositories

Keywords and phrases Machine Learning, Big Data

Digital Object Identifier 10.4230/LIPICs.OPODIS.2019.3

Category Keynote Abstract



© Haris Pozidis;

licensed under Creative Commons License CC-BY

23rd International Conference on Principles of Distributed Systems (OPODIS 2019).

Editors: Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller; Article No. 3; pp. 3:1–3:1

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany