DPP17-2017-001199

Abstract for an Invited Paper
for the DPP17 Meeting of
the American Physical Society

Contemporary machine learning: techniques for practitioners in the physical sciences[1]

BRIAN SPEARS, Lawrence Livermore Natl Lab

Machine learning is the science of using computers to find relationships in data without explicitly knowing or programming those relationships in advance. Often without realizing it, we employ machine learning every day as we use our phones or drive our cars. Over the last few years, machine learning has found increasingly broad application in the physical sciences. This most often involves building a model relationship between a dependent, measurable output and an associated set of controllable, but complicated, independent inputs. The methods are applicable both to experimental observations and to databases of simulated output from large, detailed numerical simulations. In this tutorial, we will present an overview of current tools and techniques in machine learning – a jumping-off point for researchers interested in using machine learning to advance their work. We will discuss supervised learning techniques for modeling complicated functions, beginning with familiar regression schemes, then advancing to more sophisticated decision trees, modern neural networks, and deep learning methods. Next, we will cover unsupervised learning and techniques for reducing the dimensionality of input spaces and for clustering data. We'll show example applications from both magnetic and inertial confinement fusion. Along the way, we will describe methods for practitioners to help ensure that their models generalize from their training data to as-yet-unseen test data. We will finally point out some limitations to modern machine learning and speculate on some ways that practitioners from the physical sciences may be particularly suited to help.

[1]This work was performed by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.