Accelerating atomic-scale calculations with machine learning

Andrew Peterson, School of Engineering, Brown University

Chemicals, materials, and reactions are now being engineered from first principles. That is, the equations of quantum mechanics are solved (or rather, approximated) for collections of atoms, in order that we can predict their behavior and properties, and ultimately design atomic-scale systems. The fundamental abstraction in this field is the "potential energy surface", which simply gives the energy of a system of atoms as a function of their positions. The stationary points on this surface are of particular importance: local minima correspond to molecules, and (first-order) saddle points correspond to reactions.

However, we can only calculate points on this surface at a large computational expense, and this task is perhaps the largest global user of high-performance scientific computing. However, there is incredible redundancy in such calculations; with this we can learn approximate forms of the potential energy surface, given quantum mechanical calculations as input. However, to do so in an intelligent manner relies on abstracting the geometrical information about the atomic environment into feature vectors suitable for machine learning, collecting the output of several neural networks to compute a system energy, and calculating analytical derivatives of the output in order to predict, and train on, forces. In this talk, I'll illustrate the unique demands of this application, as well as highlight the successes and challenges encountered to date.