DIFFUSION MONTE CARLO USING MACHINE LEARNING POTENTIAL ENERGY SURFACES

RYAN J. DIRISIO, MARK A. BOYER, JACOB M FINNEY, FENRIS LU, ANNE B McCOY, *Department of Chemistry, University of Washington, Seattle, WA, USA*.

Diffusion Monte Carlo (DMC) is a technique for obtaining the ground-state solution to the vibrational time-independent Schrödinger equation based on a stochastic sampling of an electronic potential energy surface (PES). Ideally, the electronic energies used in DMC are calculated at the CCSD(T) level of theory. Recently, Tom Miller's group at Caltech has developed a technique for calculating accurate CCSD(T) corrections to the electronic energy at the cost of a standard Hartree-Fock calculation using Gaussian process regression, a machine learning technique[ab]. DMC is a well-suited technique to validate these PESs, referred to as MOB-ML surfaces, since DMC samples the entire region of the PES in which the vibrational ground state wave function has amplitude. Despite the speed of these MOB-ML PESs relative to an *ab initio* calculation, calculating the Hartree-Fock energy for every one of thousands of DMC configurations over tens of thousands of time-steps makes the use of these surfaces computationally challenging.

To this end, we combined two techniques to bring these calculations into computational feasibility. First, we implemented a DMC algorithm to take advantage of massively parallel high performance computing environments. Second, we use this DMC code to collect training data for a neural network (NN). We use a generic algorithm to train the NN to learn the MOB-ML surface based on the points sampled by the DMC. Using TensorFlow, the evaluation of the NN can then easily be done using graphics processing units (GPUs). We have performed small-scale DMC calculations on $H2O$ and $CH5+$ using these MOB-ML surfaces, and have also run large-scale NN+MOB-ML DMC simulations to obtain a high resolution ground state vibrational wave function and measure of the zero-point energy. Due to the generality of both the MOB-ML surface generation and the NN fitting workflow, we can easily extend this work to larger systems, such as large ion-water clusters or a Criegee intermediate, where running the DMC with traditional electronic structure would be intractable.

---

[a] Welborn, M. , Cheng, L., Miller, T. M. JCTC. 2018 14 (9), 4772-4779

[b] Cheng, L., Kovachki, N. B., Welborn, M., Miller, T. M. JCTC. 2019 15 (12), 6668-6677