

CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday - May 29, 2020

DiffTaichi: Differentiable Programming for physical simulation

Leonard Gleyzer

We present DiffTaichi, a new differentiable programming language tailored for building high-performance differentiable physical simulators. Based on an imperative programming language, DiffTaichi generates gradients of simulation steps using source code transformations that preserve arithmetic intensity and parallelism. A light-weight tape is used to record the whole simulation program structure and replay the gradient kernels in a reversed order, for end-to-end backpropagation. We demonstrate the performance and productivity of our language in gradient-based learning and optimization tasks on 10 different physical simulators. For example, a differentiable elastic object simulator written in our language is $4.2\times$ shorter than the hand-engineered CUDA version yet runs as fast, and is $188\times$ faster than the TensorFlow implementation. Using our differentiable programs, neural network controllers are typically optimized within only tens of iterations