

# COMPUTATIONAL RESEARCH in BOSTON and BEYOND SEMINAR

## Enzyme: High-Performance, Cross-Language, and Parallel Automatic Differentiation

**WILLIAM MOSES**

Massachusetts Institute of Technology

**ABSTRACT:**

Automatic differentiation (AD) is key to training neural networks, Bayesian inference, and scientific computing. Applying these techniques requires rewriting code in a specific machine learning framework or manually providing derivatives. This talk presents Enzyme, a high-performance automatic differentiation compiler plugin for the low-level virtual machine (LLVM) compiler capable of synthesizing gradients of programs expressed in the LLVM intermediate representation (IR). Enzyme differentiates programs in any language whose compiler targets LLVM, including C/C++, Fortran, Julia, Rust, Swift, etc., thereby providing native AD capabilities in these languages with state-of-the-art performance. Unlike traditional tools, Enzyme performs AD on optimized IR. On a combined machine-learning and scientific computing benchmark suite, AD on optimized IR achieves a geometric mean speedup of 4.5x over AD on IR before optimization.

This talk will also include work that makes Enzyme the first fully automatic reverse-mode AD tool to generate gradients of existing GPU kernels. This includes new GPU and AD-specific compiler optimizations, and an algorithm ensuring correctness of high-performance parallel gradient computations. We provide a detailed evaluation of five GPU-based HPC applications, executed on NVIDIA and AMD GPUs.

**FRIDAY, DECEMBER 3, 2021  
12:00 PM – 1:00 PM**

<https://math.mit.edu/sites/crib/>

**ZOOM MEETING info:**

<https://mit.zoom.us/j/96155042770>

**Meeting ID: 961 5504 2770**