

Differentiable, ML-augmented Earth system modeling for skillful prediction, online learning, and uncertainty quantification

Patrick Heimbach, Professor, Department of Earth and Planetary Sciences

Abstract: I aim to develop the computational foundations for a new-generation of differentiable, ML-enabled ocean and climate models that comprehensively learn from data. The work is centered around the GPU-native Julia-based ocean model `ClimaOcean.jl`, for which a prototype configuration has been differentiated successfully using automatic differentiation. The framework enables comprehensive calibration of physical parameters, online or end-to-end training of embedded neural networks within a hybrid system, and equation discovery from heterogeneous observational data or high-fidelity simulations. A hierarchy of model setups will be developed, from idealized pole-to-pole single-basin baroclinic to eddy-permitting global configurations on a tripolar grid. These configurations will support gradient-based calibration of key subgrid closures and air–ice–ocean flux parameterizations, to reduce model drift, as well as the replacement of complex closure schemes with equivariant neural networks trained to capture subgrid-scale processes. The project will also introduce stochastic parameterizations into the `ClimaOcean` framework. In parallel, it will explore stochastic variational inference to approximate full posterior distributions over control variables for scalable uncertainty quantification. In collaboration with partners at AWI, PIK, TUM, and Oxford, the work will extend to an adjoint-enabled coupled ocean–seaice–atmosphere system using `SpeedyWeather.jl`. This system will be explored for coupled data assimilation and for hands-on geophysical fluid dynamics teaching in idealized configurations. Overall, the project will lay the foundations a framework for differentiable hybrid ocean–sea-ice–atmosphere model that can rigorously assimilate diverse data, perform advanced Bayesian inversion and UQ, and exploit ML-purposed HPC architectures, positioning it as a NASA’s next-generation ocean-ice state estimation system.

