

Hackathon Speeds Progress Toward Climate Model Collaboration

Climate scientists collaborated in a nationwide event to analyze and compare archived Earth system model simulations and to generate input for the IPCC's upcoming climate change report.



Nathan Collier, a computational Earth system scientist at Oak Ridge National Laboratory, evaluates CMIP6 model soil carbon turnover times while groups of researchers at other locations work on their own analyses during a hackathon event this past summer. Credit: Forrest M. Hoffman

By [Wilbert Weijer](#), Forrest M. Hoffman, Paul A. Ullrich, Michael Wehner, and Jialin Liu  13 hours ago

This past summer, scientists from the U.S. Department of Energy (DOE) gathered at six hubs across the United States to participate in a climate model comparison “hackathon.” They pooled computing resources and expertise, and they collaborated in person and via videoconferencing. By joining forces, these scientists got results more quickly, reduced duplication of efforts, and spent less time solving software problems than they would have if they had worked on their own.

Studies with computational Earth system models provide important complementary information to observational data because they provide insights into future environmental conditions.

Their findings will contribute to a sweeping report, issued every 6 years or so by the Intergovernmental Panel on Climate Change (IPCC). This report reviews the state of climate change science, documents its socioeconomic implications, and identifies viable response strategies. The IPCC has produced five assessment reports so far, and the [Sixth Assessment Report](https://www.ipcc.ch/assessment-report/ar6/) (AR6) is currently in preparation.

Analyses of the Earth system based on observational data from sensors on the ground, in the oceans, and in space form an important basis for these reports. But studies with computational Earth system models (ESMs) provide important complementary information because they provide insights into future environmental conditions and help attribute observed changes to specific causes.

Each model (and there are many) incorporates its own body of source data, assumptions, and algorithms. Thus, the best overall picture of Earth’s climate emerges when results from several models are compared, taking note of the strengths and limitations of each. However, this type of comparison poses challenges to individual researchers.

An Ensemble Cast of Models

ESMs capture the essential processes within the Earth system in complex computer codes. Despite researchers’ best efforts, no model is perfect because scientists must make difficult choices to balance physical realism with reasonable [computational performance](https://eos.org/research-spotlights/reframing-sensitivity-analysis-in-earth-system-models). It is therefore dangerous to base important conclusions about climate change on any single model.

Models can be used to explore different narratives of future human responses to climate change, from scenarios of high greenhouse gas emissions to aggressive mitigation policies.

Recognizing this, the World Climate Research Programme organized the Coupled Model Intercomparison Projects ([CMIPs](https://www.wcrp-climate.org/wgcm-cmip)) in the mid-1990s. In these projects, modeling centers run their models according to a common [simulation protocol](https://eos.org/science-updates/toward-standardized-data-sets-for-climate-model-experimentation) using a common set of [forcings](https://www.climate.gov/maps-data/primer/climate-forcing) (e.g., prescribed carbon dioxide concentrations in the atmosphere).

Scientists soon realized that they learn more by comparing model behaviors with the observed climate as well as with other models. The range of future outcomes generated by various models using different inputs provides

estimates of both the internal variability of Earth's climate system and the models' structural uncertainties. Valid models can be used to explore different narratives of future human responses to climate change, from scenarios of high greenhouse gas emissions to aggressive mitigation policies that reduce emission rates to near zero (<https://eos.org/meeting-reports/achieving-a-near-zero-carbon-emissions-energy-system>).

CMIP6: A Work in Progress

Currently, CMIP is in its sixth phase. Central to CMIP6 is a set of idealized simulation protocols called Diagnostic, Evaluation and Characterization of Klima (DECK), which are designed to enhance understanding of the climate system's response to increasing amounts of greenhouse gases in the atmosphere.

Modeling centers are asked to contribute DECK simulations first because DECK is the most basic simulation protocol. Afterward, they are free to participate in any of the other 23 other MIPs that have been approved as part of CMIP6. When CMIP6 is complete, more than 20 modeling centers around the world will have provided tens of petabytes of data produced by more than 40 models and distributed over millions of files.

The DOE's Office of Science has significant investments in ESM development, validation, and analysis. Its Earth and Environmental System Modeling (EESM) program is contributing CMIP6 simulations with its own model, the Energy Exascale Earth System Model ([E3SM](https://e3sm.org/) (<https://e3sm.org/>)). This model, one of just a few completely new Earth system models introduced in recent years, saw its first release in 2018. Another program element of EESM, Regional and Global Model Analysis (RGMA), focuses on model evaluation, diagnosis, and analysis of ESMs; analysis of CMIP6 model results is an important mandate for RGMA-funded projects.

The Multimodel Analysis Challenge

For individual scientists, performing multimodel analysis can be a daunting task. First, an analyst needs access to local storage that can hold many terabytes of data, analysis and visualization software that can handle large data sets, and a powerful computing platform to perform complex computations. Second, many data processing tasks are tedious and time-consuming: identifying required data files in an online catalog, downloading and inspecting the data, noting each model's idiosyncrasies, preprocessing the data consistently through tens of models (e.g., extracting data for the time window and region of interest and calculating annual averages or anomaly time series), computing relevant metrics, and visualizing outcomes in meaningful ways. These tasks are repeated by analysts all over the world—sometimes even by multiple colleagues in the same group—representing a duplication of efforts and, often, time wasted.

If commonly used CMIP6 data were accessible to a large group of collaborators, directly connected to a powerful computational platform, then our teams could focus on producing science from the start.

During spring 2019, the authors of this article, all RGMA-funded scientists, decided that CMIP analysis could be greatly accelerated if these technical bottlenecks were managed as a group. If commonly used CMIP6 data were

accessible to a large group of collaborators, directly connected to a powerful computational platform with preinstalled and tailored analysis and visualization software, then our teams could focus on producing science from the start.

From this realization, the idea for a hackathon arose. Hackathons are common events in the software engineering world in which programmers gather to collaborate intensively on a specific task. We intended for our hackathon to be an opportunity to make rapid progress in processing and analyzing CMIP6 data. We had the following goals: (1) assemble a common data cache that is quickly accessible (low latency) to many scientists from a powerful analysis platform, (2) build a common analysis environment capable of handling large data volumes, and (3) build a community of scientists collaborating toward the common goal of producing policy-relevant science.

1. We wanted to assemble a common data cache because most of the participants in this effort placed a high value on having easy access to CMIP6 data. We worked with the staff at the National Energy Research Scientific Computing Center ([NERSC \(https://www.nersc.gov/\)](https://www.nersc.gov/)), the primary scientific computing facility for the DOE Office of Science, who provided us with a 2.25-petabyte disk array. Then we began downloading to this large disk array a subset of model outputs from the DECK and Shared Socioeconomic Pathway (SSP) simulations from the CMIP6 data archive that is hosted on the Earth System Grid Federation ([ESGF \(https://esgf.llnl.gov/\)](https://esgf.llnl.gov/)).

We downloaded data around the clock for several months and staged on the disk array a variety of observational and climate reanalysis data for use in evaluating, validating, and benchmarking the performance of CMIP6 models. We distributed a survey in which participants could indicate their data needs, allowing us to prioritize the data downloads.

2. To build a common analysis environment, we promoted the use of Community Data Analysis Tools ([CDAT \(https://cdat.llnl.gov/\)](https://cdat.llnl.gov/)) software developed by our colleagues in the Program for Climate Model Diagnosis and Intercomparison ([PCMDI \(https://pcmdi.llnl.gov/\)](https://pcmdi.llnl.gov/)). In the lead-up to the hackathon, we provided several training opportunities to enable participants to become familiar with the analysis software, culminating in a 4-hour tutorial session presented by CDAT developer Charles Doutriaux. Even so, it seemed that the lead time was too short for many participants to familiarize themselves with the unique capabilities of CDAT, so in the end, most analysts stuck with the tools they were most familiar with: Python, Matlab, NCL (National Center for Atmospheric Research Command Language), and others.
3. By fostering interactions among scientists from different projects, we hoped to facilitate exchanges of useful information and possibly initiate new collaborations.

With regard to building a community of scientists, several RGMA-funded projects routinely perform model intercomparisons, but other teams have less experience. By fostering interactions among scientists from different RGMA projects, we hoped to facilitate exchanges of useful information and possibly initiate new

collaborations. Several days before the hackathon, we organized a teleconference during which scientists could present and discuss their analysis plans. This session gave people the opportunity to learn about their colleagues' plans, coordinate analysis tasks, and request help. This event initiated several new collaborations.

A Successful Hackathon

The hackathon was held from 31 July through 6 August 2019. Scientists worked in a collaborative and focused setting from six hubs distributed across the country. Videoconferencing capabilities enabled the hubs to remain in contact around the clock. Participants exchanged information using the messaging software Slack, and they exchanged analysis scripts using the software development platform GitHub.

Roughly 50 scientists participated in the hackathon at any given time, and about 100 RGMA scientists signed up for access to the tutorial materials, our ongoing Slack discussions, the GitHub code repository, and the data cache that we established at NERSC. Each day at 1 p.m. Eastern time, hackathon participants discussed their progress and challenges, one hub at a time. This daily check-in led to constructive discussions and suggestions for improving analyses and graphical diagnostics.

The coordinated CMIP6 analysis activity enabled many users to make rapid progress on meaningful science. Within 3 months, the first analyses that were initiated during the hackathon were completed.

The coordinated CMIP6 analysis activity enabled many users to make rapid progress on meaningful science using the CMIP6 archive and high-performance computing resources. Within 3 months, the first analyses that were initiated during the hackathon were completed and submitted for publication, and findings were shared with lead authors of the IPCC AR6. Many more papers are in preparation. One study, led by the lead author of this article, documents how CMIP6 models unanimously project a significant slowdown of the Atlantic Meridional Overturning Circulation (<https://eos.org/research-spotlights/theoretical-models-advance-knowledge-of-ocean-circulation>), an important player in the climate system, by 2100. Other teams are studying processes like hurricanes and cyclones; monsoons; extreme precipitation events, droughts, and heat waves; Arctic sea ice; tropical forests; and the carbon cycle.

Participant feedback showed that the hackathon and associated activities were very well received among those who took part. One participant noted that the event “helped me analyze CMIP6 data more efficiently and solve software and programming issues quickly.” Another noted that the hackathon “saved a lot of my time that I would have spent otherwise learning by myself.” After the clear success of this event, we hope to organize similar analysis and data synthesis activities in the future.

Author Information

Wilbert Weijer (wilbert@lanl.gov), Los Alamos National Laboratory, N.M.; Forrest M. Hoffman, Oak Ridge National Laboratory, Tenn.; Paul A. Ullrich, University of California, Davis; and Michael Wehner and Jialin Liu, Lawrence Berkeley National Laboratory, Berkeley, Calif.

Citation: Weijer, W., F. M. Hoffman, P. A. Ullrich, M. Wehner, and Ji. Liu (2019), Hackathon speeds progress toward climate model collaboration, *Eos*, 100, <https://doi.org/10.1029/2019EO137735>. Published on 23 December 2019.

Text not subject to copyright.

Except where otherwise noted, images are subject to copyright. Any reuse without express permission from the copyright owner is prohibited.