Vijay Gadepally, Lincoln Laboratory, CSAIL
Lauren Milechin, EAPS
Jeremy Kepner, Lincoln Laboratory, CSAIL, Math

MIT SuperCloud for AI Algorithm Development

Patrick Brown, MIT Energy Initiative
Sam Elder, Math Department
Kritkorn Karntikoon, CSAIL
Jingzhao Zhang, LIDS

MIT SuperCloud Overview

- The MIT SuperCloud provides interactive, on-demand high performance computing to accelerate AI and Machine Learning research
- Combines traditional HPC resources with modern tools for AI and Big Data research
  - Jupyter Notebooks
  - Machine Learning Tools (ex: Tensorflow)
  - High Performance Databases
- Focus on productivity
  - User-friendly environment
  - Online courses
  - In-person tutorials

Resources located at the Massachusetts Green High Performance Computing Center, a state of the art facility in Holyoke, MA providing 93% carbon-free power.

Projects Using MIT SuperCloud

- Used in departments across MIT
- Some AI-related research on SuperCloud include:
  - Patrick Brown, MIT Energy Initiative: Use historical records of electricity prices, solar irradiance, and wind speed to estimate the economic value of distributed renewable energy and energy storage at tens of thousands of locations across the US and to evaluate long-distance complementarity between solar and wind resources (see upper right)
  - Kritkorn Karntikoon, LIDS: Reconstruct and analyze the global climate network via the geometric lens of Ricci curvature (see bottom right)
  - Sam Elder, Math Department: Investigating and evaluating better ways of doing machine learning model assessment and selection using new variants to standard cross-validation techniques
  - Jingzhao Zhang, LIDS: Interpret and design accelerated first order optimization algorithms
- Want to learn more? Visit https://supercloud.mit.edu

Economic Value of Renewable Energy in the United States

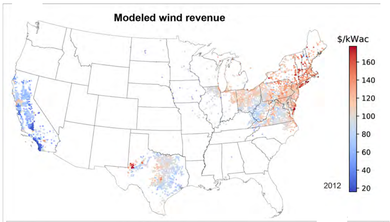

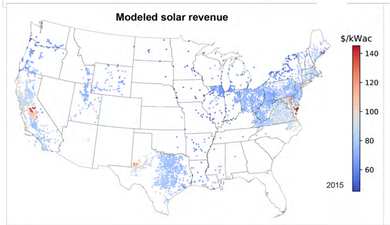

- The goal is to estimate economic value of renewable energy distribution and storage across the US
- Cost of wind and solar power is becoming competitive with fossil-fuel-based electricity; value is closely tied to the profile of electricity prices at their installation location
- Determine the locational value of electricity from wind and solar installations across the US using:
  - Locational electricity prices at ~14,000 transmission grid nodes
  - Publicly available wind and solar time series datasets from NREL*
  - Open-source generator models
- Supercloud computing resources accelerate these calculations: what would have taken months can now be done in days

*NREL: National Renewable Energy Laboratory

Climate Networks and Ricci Curvature

- The ultimate goal is to understand the climate
- We use a dataset with historical temperature time series for ten thousand positions over 70 years
- For our analysis, we need to:
  - Form a graph from the dataset
  - Compute the correlation between all 5×10^7 pairs of time series
  - Computing the correlation takes 6 hours on one processor
- With SuperCloud, we can compute correlations in parallel and make our analysis 60 times faster