Project Hermès: High-Resolution Modeling of the Earth System

MOPGA Kickoff Meeting

V. Balaji

Laboratoire des Sciences du Climat et de l'Environnement (LSCE)
Institut Pierre Simon Laplace (IPSL)
NOAA/Geophysical Fluid Dynamics Laboratory (GFDL) and Princeton University

1 February 2019

Outline

1. The computing challenge for high-resolution modeling
   - End of Dennard scaling
   - Computing transition to machine learning (ML)
2. Challenges posed by machine learning
   - Simulation versus understanding
   - Model calibration
   - "Model-free" methods
   - Should ML “learn” what we already know?
   - Limitations of training data
3. Project Hermès
   - Metamodels and supermodels
   - Supermodels: learning parameterizations from observations
   - Metamodels: model calibration in idealized and real settings
   - Strategy and timeline

Outline

1. The computing challenge for high-resolution modeling
   - End of Dennard scaling
   - Computing transition to machine learning (ML)
2. Challenges posed by machine learning
   - Simulation versus understanding
   - Model calibration
   - "Model-free" methods
   - Should ML “learn” what we already know?
   - Limitations of training data
3. Project Hermès
   - Metamodels and supermodels
   - Supermodels: learning parameterizations from observations
   - Metamodels: model calibration in idealized and real settings
   - Strategy and timeline

History of GFDL Computing

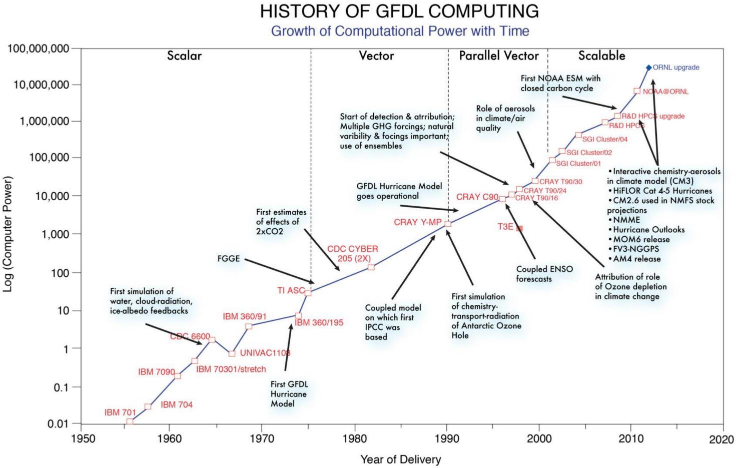

Courtesy V. Ramaswamy, NOAA/GFDL.

Moore's Law and End of Dennard scaling

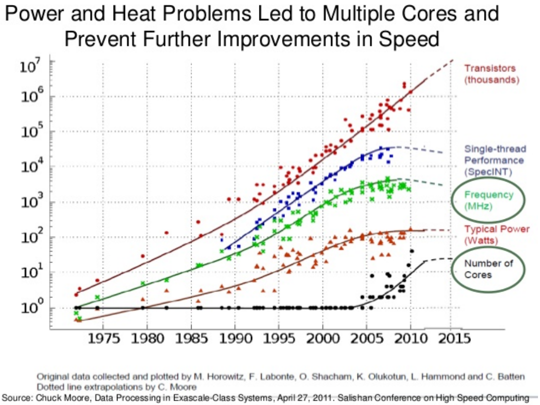

Figure courtesy Moore 2011: *Data processing in exascale-class systems.*

- Processor concurrency: Intel Xeon-Phi.
- Fine-grained thread concurrency: Nvidia GPU.

Deep Learning

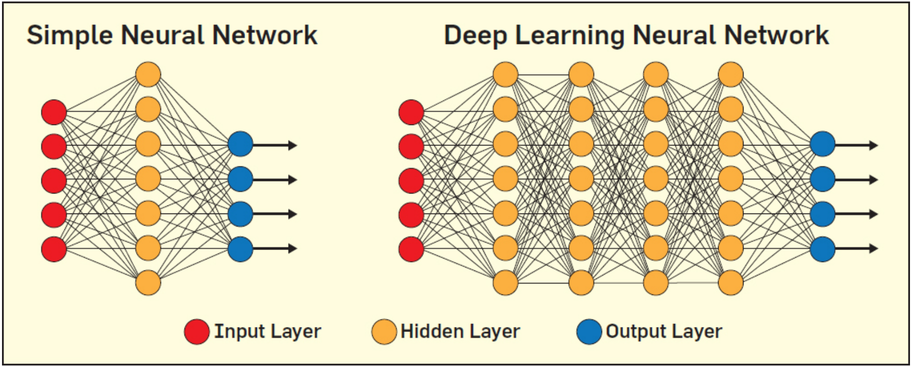

From Edwards (2018), ACM.

Outline

1. The computing challenge for high-resolution modeling
   - End of Dennard scaling
   - Computing transition to machine learning (ML)
2. Challenges posed by machine learning
   - Simulation versus understanding
   - Model calibration
   - "Model-free" methods
   - Should ML “learn” what we already know?
   - Limitations of training data
3. Project Hermès
   - Metamodels and supermodels
   - Supermodels: learning parameterizations from observations
   - Metamodels: model calibration in idealized and real settings
   - Strategy and timeline

NGGPS: Next-Generation Global Prediction System

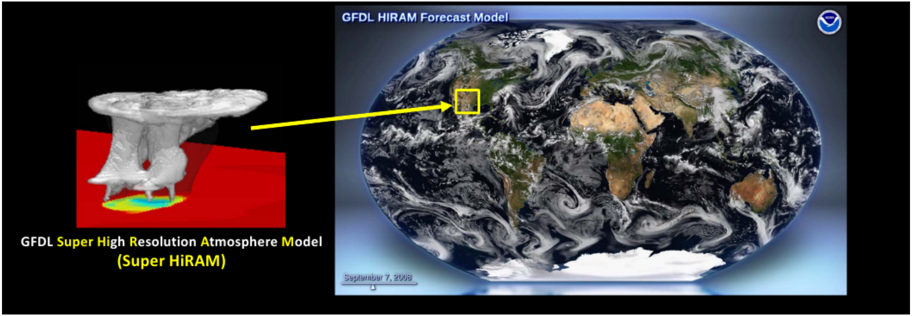

FV3 dynamical core from GFDL for the next-generation forecast model (target: 3 km non-hydrostatic in 10 years running at $\sim$ 200 d/d)

Passing the climate Turing test?

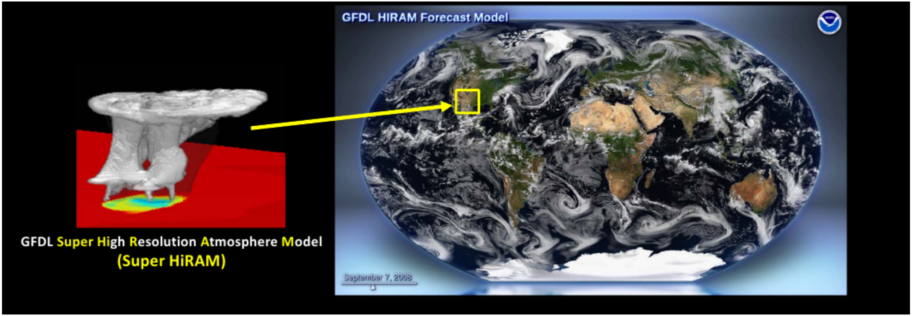

We may be able to simulate everything in great detail, but do we understand how it works?

The model hierarchy

- Molecular biology uses a hierarchy of "models": *E. Coli*, *C. Elegans*, fruit fly, mouse, *H. Sapiens*, ...
- We have a similar hierarchy: LES, CRM, AOGCM, ESM, ...
- and a hierarchy of idealized experiments: turbulent flow, radiative-convective equilibrium, aquaplanet, AMIP, OMIP, control, historical, ...
- Community must run common experiments at all levels of the hierarchy ("idealized MIPs")...
- "Verification" (or falsification) of idealized planet Earth? analysis must isolate underlying mechanisms even in complex models.

Adapted from Held (2005, 2014). Model Hierarchies Workshop, November 2016 in Princeton.

Model calibration

Model calibration or "tuning" consists of reducing overall model bias (usually relative to 20th century climatology) by modifying parameters. In principle, minimizing some cost function:

$$C(p_1, p_2, ...) = \sum_{1}^{N} \omega_i \|\phi_i - \phi_i^{obs}\|$$

- Usually the p must be chosen within some observed or theoretical range $p_{min} \leq p \leq p_{max}$.
- "Fudge factors" (applying known wrong values) generally frowned upon (see Shackley et al 1999 discussion on history of "flux adjustments". More on that later...)
- The choice of ω_i is part of the lab's "culture"!
- The choice of ϕ_i^{obs} is also troublesome:
  - overlap between "tuning" metrics and "evaluation" metrics.
  - "Over-tuning": remember "reality" is but one ensemble member!

Objective methods of tuning

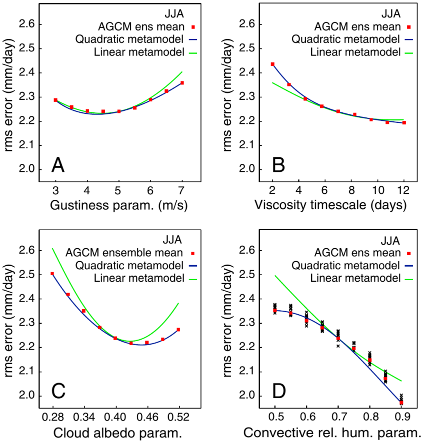

Neelin et al (2010) construct “metamodels” to aid in multi-parameter optimization. See also Zamboni et al.

Model-free prediction vs model augmentation

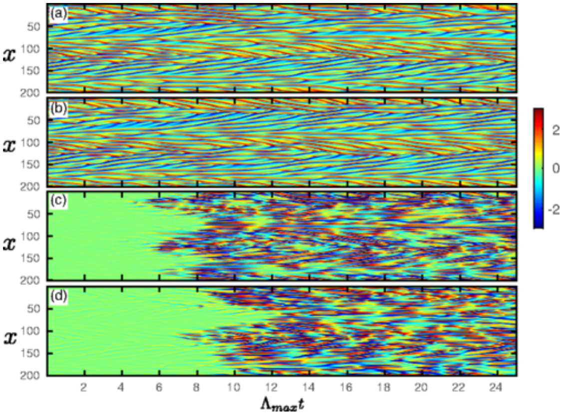

From Pathak et al, PRL (2018), *Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach*

Movie: Pathak's flame front in *Quanta*.

Making ML respect known physical constraints

See momentum conservation discussion in Bolton and Zanna (2018), *Applications of Deep Learning to Ocean Data Inference and Sub-Grid Parameterisation*.

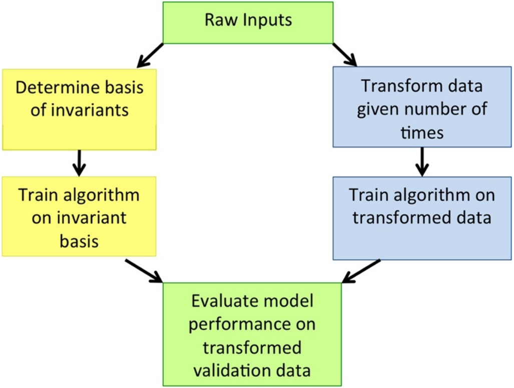

From Ling et al, JCP (2016), *Machine learning strategies for systems with invariance properties*

Distilling Free-Form Natural Laws from Experimental Data

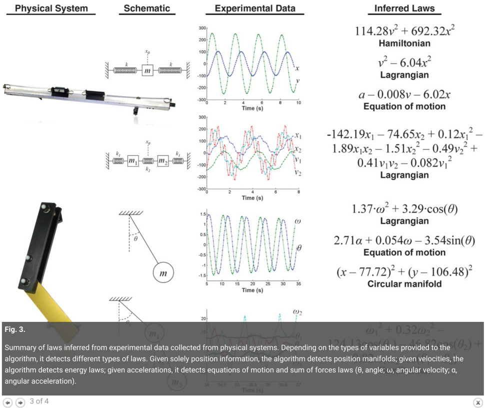

Fig. 3.

Summary of laws inferred from experimental data collected from physical systems. Depending on the types of variables provided to the algorithm, it detects different types of laws. Given solely position information, the algorithm detects position manifolds; given velocities, the algorithm detects energy laws; given accelerations, it detects equations of motion and sum of forces laws (θ, angle; ω, angular velocity; α, angular acceleration).

From Schmidt and Lipson, *Science*, 2009. My little *hommage*, Gaitán et al (2016), *Can we obtain viable alternatives to Manning's equation using genetic programming? Eureqa* software available under license.

Navier-Stokes from data

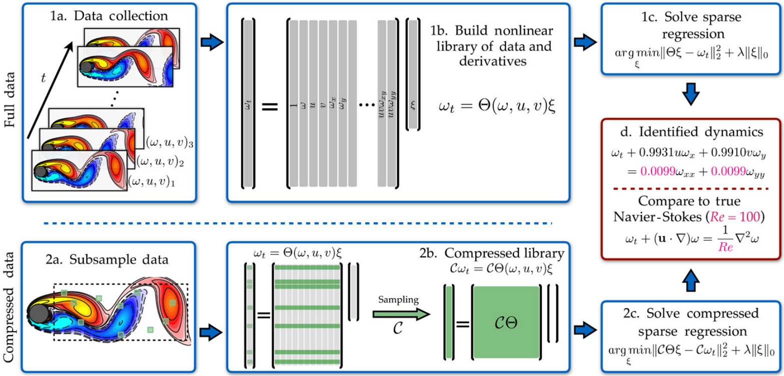

From Rudy et al (2017).

Limitations of training data

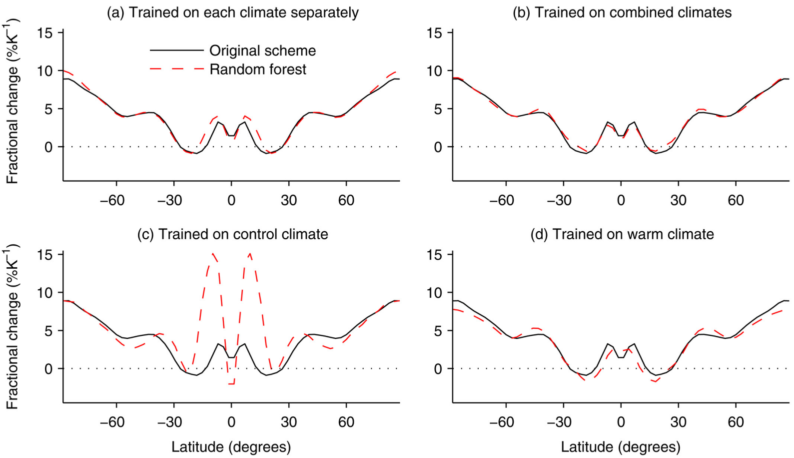

From O'Gorman and Dwyer, JAMES, 2018. Limitations of training on short non-stationary time series.

Error patterns associated with stationarity assumption

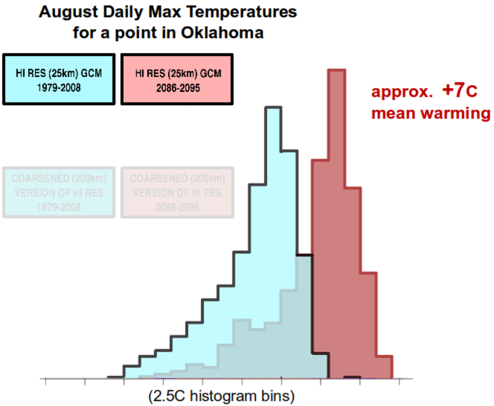

Errors can be traced with warming outside the temperature distribution of the training period. Caution needed at distribution tails ("extreme events"). Dixon et al (2016).

Where models and data are both weak...

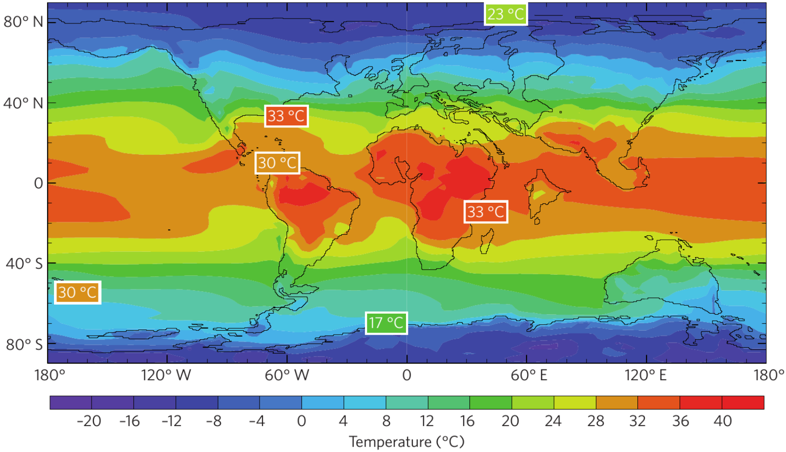

Fig 1 from Valdes (2011). GCMs are unable to simulate the Paleocene-Eocene climate of 55 My ago.

Outline

1. The computing challenge for high-resolution modeling
   - End of Dennard scaling
   - Computing transition to machine learning (ML)
2. Challenges posed by machine learning
   - Simulation versus understanding
   - Model calibration
   - "Model-free" methods
   - Should ML “learn” what we already know?
   - Limitations of training data
3. Project Hermès
   - Metamodels and supermodels
   - Supermodels: learning parameterizations from observations
   - Metamodels: model calibration in idealized and real settings
   - Strategy and timeline

Questions: metamodels and supermodels

- Supermodels: some components replaced by learning agents. Metamodels: low-dimensional emulators, “fast approximate models”.
- Fundamental questions still unanswered:
  - Are model-free methods useful?
  - How do we derive the invariant basis of a complex system?
  - Can we use ML to derive the functional form of a slow manifold?
  - Can we derive a useful model hierarchy?
  - Can this metamodel be used for parameter uncertainty exploration?
  - How much physical knowledge (e.g conservation laws) must be embedded in the ML? What if the embedded knowledge is incorrect? (“*It’s not what you don’t know, it’s what you know for sure that just ain’t so*”, Mark Twain never said.)
  - What happens to supermodels as the features of the training data evolve?

Learn parameterizations from observations

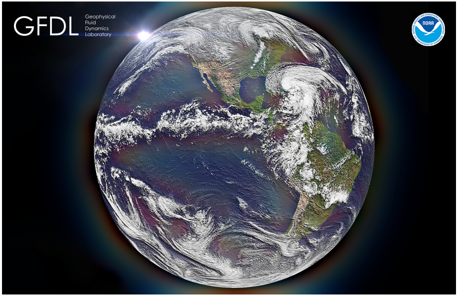

(Courtesy: S-J Lin, NOAA/GFDL).

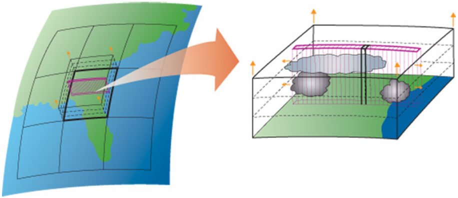

(Courtesy: D. Randall, CSU; CMMAP).

- Global-scale CRMs (e.g 7 km simulation on the left) and even super-parameterization using embedded cloud models (right) remain prohibitively expensive.
- Can we learn the statistical aggregate of small scales? See Schneider et al 2017, Gentine et al (2018), O'Gorman and Dwyer (2018), Bolton and Zanna (2018), ...

Learning sub-gridscale turbulence

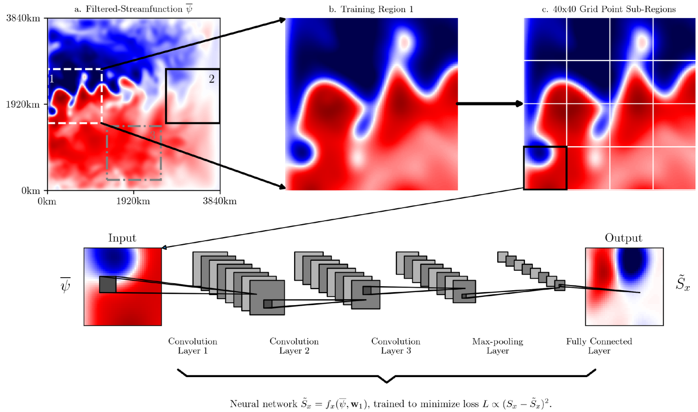

Fig 1 from Bolton and Zanna (2018), in review for JAMES.

Caltech/MIT Earth Machine

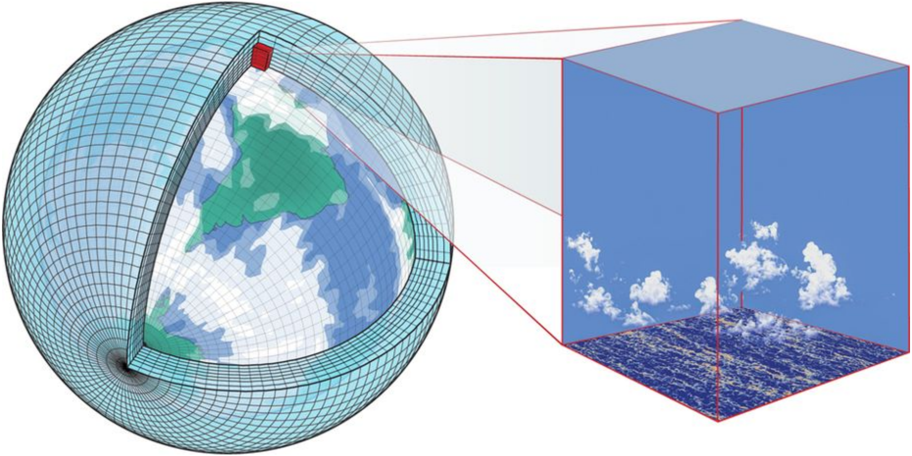

From Schneider et al 2017.

Lorenz 96, a nice abstraction

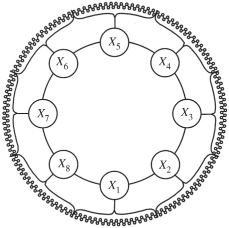

$$\frac{dX_k}{dt} = -X_{k-1}(X_{k-2} - X_{k+1}) - X_k + F - \frac{hc}{b}\sum_{j=1}^{32} Y_{j,k} + f \tag{1}$$

$$\frac{dY_{j,k}}{dt} = -cbY_{j+1,l}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{hc}{b}X_k \tag{2}$$

A nice abstraction of a system with fast and slow modes, whose coupling strength can be varied... maybe too interesting? See metastability issues in Schneider et al, GRL (2017).

Lorenz96 in perfect model setting

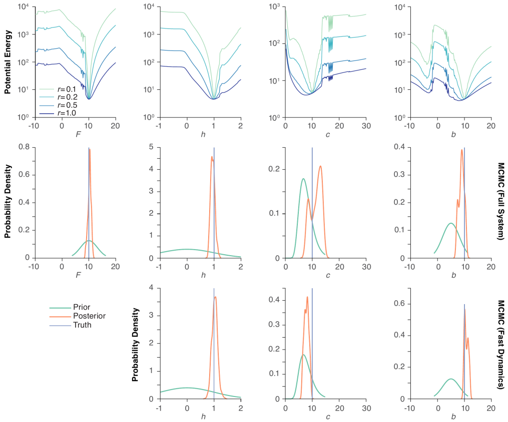

From Schneider et al 2017. Learn Lorenz96 parameters F, h, c, b from prior run.

Project Hermès: Strategy

- Project Hermès will study learning methods for metamodels and supermodels, for atmosphere and ocean, across a hierarchy of models (e.g 1D and 3D, idealized, LES, GCM, ...)
- Project Hermès will aim to foster collaboration in the emerging field of Climate/ML: between the ML and climate communities, theoretical and applied science, between institutions in EU/US/...
- Project Hermès will aim to build a community of interdiscplinary scientists equally at home in machine learning and Earth System science.
- Project Hermès will aim to be open-minded and opportunistic: this is a nascent field and there will be unexpected twists and turns!
- All work will be shared with the community via articles in open-access journals, open-source software, open data. No commercial or proprietary interests. Articles will list LSCE affiliation first and acknowledge MOPGA funding.

Collaborations beginning under Project Hermès

Initial presentations:

- Presentation of the ML challenge to the community: IPSL (Dec 2018), LSCE (Jan 2019).

Beginning collaborations:

- Extension of Bolton-Zanna approach using high-resolution ocean models
  - In collaboration with LOCEAN, Uni Grenoble, Oxford, Princeton
  - MOPGA postdoc (LSCE) under recruitment.
- Application of ML to model calibration
  - In collaboration with LMD, Univ Exeter, École Normale Supérieure, ANR High Tunes project.
  - Doctoral student (ENS) under recruitment.
- Detection of features (e.g tropical cyclones) in high-resolution climate data.
  - In collaboration with IPSL.

Proposed timeline

- Year 1:
  - Recruitment for subprojects (see above).
  - Monthly Journal Club starting Feb 2019.
  - Presentation of the Climate/ML challenge to the community: IPSL (Dec 2018), LSCE (Jan 2019), SAMA IA-Climat (Feb 2019).
  - Invited presentation at LEFE/MANU Journée Thematique à Rennes (Feb 2019)
  - Invited keynote presentation at EGU Assembly Vienna (April 2019).
  - Articles in preparation:
    - *Metamodels and supermodels*: ideas and challenges from machine learning in Earth System Science.
    - *The biology analogy*: will *in silico* science become like *in vitro*?
- Year 3:
  - Demonstration of supermodeling approach in at least one aspect of IPSL model.
  - Demonstration of ML application in calibration of IPSL model.
- Year 5:
  - Hybrid (ML/physics-based) model in production.