Statistical Methods for Ice Sheet Model Calibration

Murali Haran

Department of Statistics, Penn State University

November 13, 2020

STAMPS Lecture Series, Carnegie Mellon University

Ice Sheets and Sea Level Rise

(Courtesy of NASA)

Ice sheets contain an immense amount of water

- Antarctic ice sheet: equivalent of **58 m** sea level rise
- Greenland ice sheet: **6 m** sea level rise
- Greenland + Antarctic contain > 99% freshwater ice in the world

Hurricane Gustav, New Orleans (2008)

(Richard B. Alley, Penn State PA Environmental Resource Consortium)

Connecting Changing Temperatures to Sea Level Rise

Statistics and Ice Sheets

Formulating reasonable hypotheses regarding climatic change requires physical insight and ingenuity, but subsequently testing these hypotheses demands quantitative computation. – Edward N. Lorenz (1970)

- My translation: to study future climate we need physical models, data, and statistics and innovation on all three fronts
- Focus here: Studying the West Antarctic Ice Sheet (WAIS)

Talk Summary

- How can we understand the dynamics of the West Antarctic Ice Sheet and project its future behavior?
  - Ice sheet model: PSU3D-ICE (Pollard and DeConto, 2012)
  - Key input parameters governing model behavior are uncertain
  - Use model runs + observations to learn about these parameters
  - Model runs are computationally expensive
- Tradeoffs
  - model resolution
  - computational time
  - # of parameters to study
- I will give an overview and contrast two methods for this problem
  - Gaussian process emulation-calibration
  - Calibration via particle-based sequential Monte Carlo

Model of Ice Sheet Physics

Ice sheet dynamics are translated into a computer model
Several uncertain parameters drive the dynamics

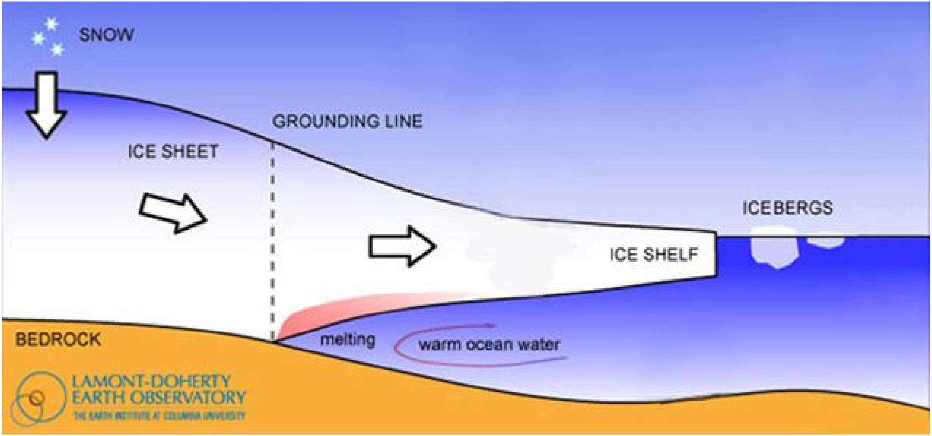

(Courtesy of the Earth Institute at Columbia University)

Uncertain Parameters

- Many parameters are key to ice sheet dynamics. They appear as constants in the computer model OCFACMULT, OCFACMULTASE, CRHSHELF, CRHFAC, ENHANCESHEET, ENHANCESHELF, FACEMELTRATE, TAUASTH, CLIFFVMAX, CALVLIQ, and CALVNICK
- Changing these parameters changes how the ice sheet evolves over time, and how it responds to the climate

Observations

- Modern ice sheet extent satellite (spatial) data
- Some parameters apply to processes that have occurred in the past and are expected in the future, but are not active today, for instance
  - Timescale of bedrock rebound (TAUASTH) under varying ice loads
- Or processes that are undergoing rapid change in recent decades, e.g.
  - Coefficients for oceanic melting at the base of floating ice shelves (OCFACMULT)
- Hence, utilize reconstructions of past ice sheet behavior, example:
  - From Last Interglacial period $\approx$ 115,000 to 125,000 years ago
  - Grounding line data since Last Glacial Maximum (LGM) $\approx$ 25,000 years ago (time series)

Why Use Probability Models Here?

It is common to use informal approaches to learn about parameters, e.g. find the parameter values that produce model output within x of observations (using some metric)? Why use probability models?

1. Posterior distributions of parameters easier to interpret: "Given all the data we have and our assumptions, $P(\theta > 1.4)$ is 0.1"
2. Easier to compare results across resolutions, kinds of data used, data aggregation etc. if they are in the form of probability distributions
3. Can model systematic model-data discrepancies (cf. Bayarri et al., 2007) + account for observation errors + approximation errors
4. Straightforward to use distributions of parameters to obtain model projections, also in the form of probability distributions
5. Can easily study relationships between parameters
6. In practice we find that results are sharper/more useful

Example of Value of Statistical Modeling

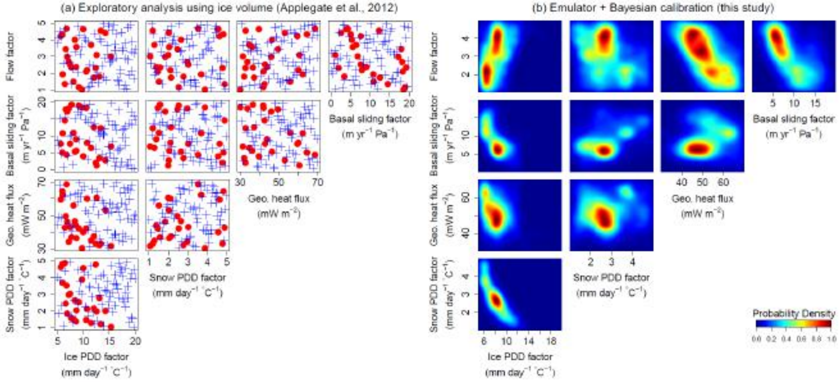

Left: informal approach. Right: statistical calibration
Chang, Haran, Olson, Keller (2014), *Annals of Applied Stats*

Model Complexity

- Ice sheet models vary in complexity
  - Key drivers: spatial and temporal resolution
- **Simple models** (cf. Shaffer, 2014; Bakker et al., 2016)
  - Simplify or exclude important physical processes
  - Run time $\approx$ few seconds
- **Complex models**, e.g. PSU3D-ICE (DeConto and Pollard, 2016) or Larour et al. (2012)
  - Better represent key ice dynamics; higher spatio-temporal resolutions
  - Can take hours or days ro run at each setting
  - Hence, difficult to study the behavior of the model
    - Challenging to study more than $\approx$ 4 - 6 parameters

Two Approaches

In both cases our group ran PSU3D-ICE, a model with a sophisticated representation of ice dynamics. The difference was in the resolution used

1. High resolution: horizontal resolution of 20 km
   - Takes several hours per model run
   - Consider only 4 parameters as uncertain
   - Methodology: handle computing by **model emulation** (approximation)
2. Coarse resolution: horizontal resolution of 80 km
   - Takes 10 - 15 minutes per model run
   - We want to use a better model for ice dynamics while still allowing for better exploration of the model
   - Can now consider 11 parameters
   - Methodology: handle computing costs by **particle methods and massive parallelization**

These approaches strike different compromises

Statistical Framework

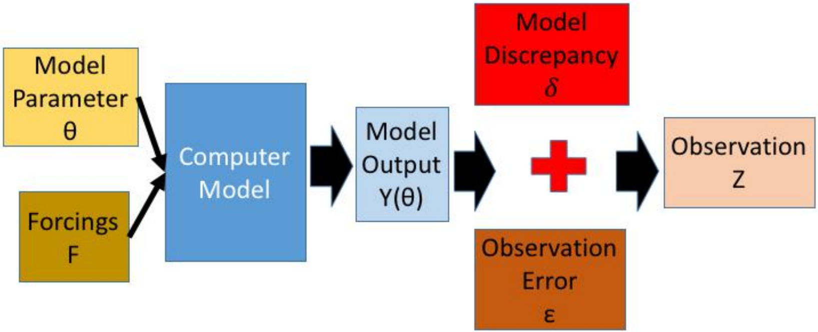

- External forcings on ice sheet model: e.g. global mean temperature
- Model Discrepancy (δ): Systematic difference between observations and model output around the "best" parameter settings

Statistical model: $Z = Y(\theta) + \delta + \epsilon$

Calibration

Model for observations Z

$$Z = Y(\theta) + \delta + \epsilon,$$

- $Y(\theta)$: Model output
- θ: Model parameter
- δ: Discrepancy term
- ϵ: Observational error w/ parameter σ^2

Inference is based on posterior distribution, $\pi(\theta, \delta, \sigma | Z, Y)$:

$$\boldsymbol{\pi(\theta, \delta, \sigma_z^2 | Z, Y) \propto \underbrace{\mathcal{L}(\theta, \delta, \sigma; Z, Y)}_{\text{Likelihood}} \times \underbrace{p(\theta, \delta, \sigma^2)}_{\text{Prior for } \theta, \delta, \sigma}}$$

(cf. Kennedy & O'Hagan, 2001)

Calibration via Markov chain Monte Carlo

- Inference via Markov chain Monte Carlo algorithm with $\pi(\theta, \delta, \sigma^2 | Z, Y)$ as its stationary distribution
- In principle, this would even work well for many parameters (high-dimensional θ)
- Problem: likelihood evaluation involves running the model: $Y(\theta)$ If model takes hours at each θ, this is prohibitively expensive

Approach # 1: emulation-based approach that replaces slow computer model with a fast stochastic approximation

Approach 1: Emulation-Calibration

Study only 4 parameters; remaining are fixed at values determined by experts and past studies. Then use emulation-calibration:

1. Emulation step: Find fast approximation for computer model using a Gaussian process (GP) (cf. Sacks et al., 1989)
2. Calibration step: Infer climate parameter using emulator and observations, while accounting for data-model discrepancy

Doing it in stages ("modularization") has computational and inferential advantages (e.g. Liu et al., 2009; Bhat, Haran, Olson, Keller, 2012)

Gaussian Process Emulation

- Idea: Fit flexible model relating parameter to model output
  - Model is simple (easy to evaluate/simulate)
  - Model allows for approximation uncertainty
  - Model is stochastic: useful for inference
- Run model at p parameter settings to obtain $(\theta_1, Y(\theta_1)), \ldots, (\theta_p, Y(\theta_p))$
- Fit Gaussian process (GP) to these pairs to obtain emulator
  - GP is infinite dimensional process with a positive definite covariance function s.t. every finite collection of random variables has a multivariate normal distribution

$$(Y(\theta_1), \ldots, Y(\theta_p))^T \sim N(\mu, \Sigma_\phi)$$

    where ϕ are covariance function parameters.
    Fitting the GP involves estimating μ, ϕ from model runs
  - Fitted model provides probability model for Y at any new value of θ: $\eta_\phi(\theta)$

Emulation Step

Simple example: model output is a scalar and continuous

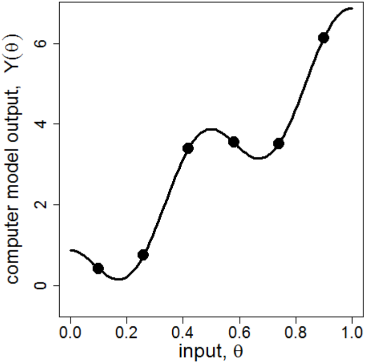

Computer model output (y-axis) vs. input (x-axis)

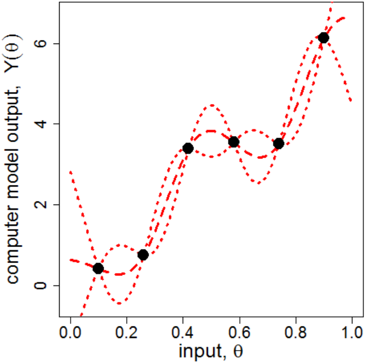

Emulation (approximation) of computer model using GP

Calibration: Inference by Approximate Likelihood

- Probability model for observations used to be

$$Z = Y(\theta) + \delta + \epsilon_\sigma.$$

Now approximate model for observations is

$$Z = \eta_\phi(\theta) + \delta + \epsilon_\sigma,$$

δ is discrepancy; ϵ is observation error
ξ, σ^2 are parameters for each process respectively

  - Discrepancy model needs to be flexible enough to adapt to systematic differences between observations and model but not so flexible that it causes identifiability issues. E.g. Gaussian process with strong priors

- Above leads to approximate likelihood, $\hat{\mathcal{L}}_\phi(\theta, \xi, \sigma^2; Z, Y)$
- Inference for θ using observations is now

$$\boldsymbol{\pi(\theta, \xi, \sigma^2 | Z, Y) \propto \underbrace{\hat{\mathcal{L}}_\phi(\theta, \xi, \sigma^2; Z, Y)}_{\text{Approximate likelihood}} \times \underbrace{p(\theta, \xi, \sigma^2)}_{\text{Prior for } \theta,\ \sigma^2, \xi}}$$

Simple example: model output, observations are scalars

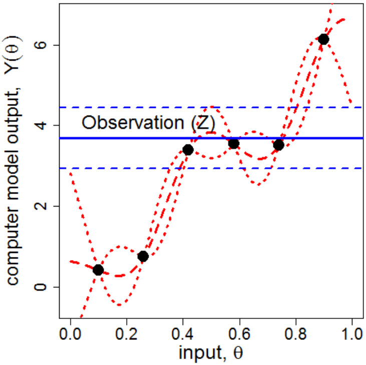

Combining observation and emulator

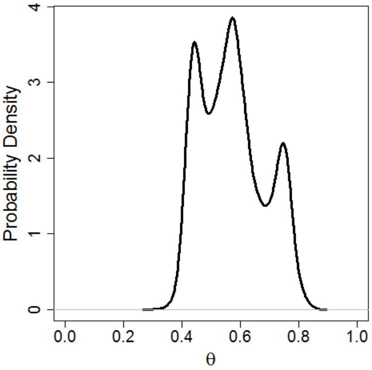

Posterior PDF of θ given model output and observations

Challenges in Emulation-Calibration Approach

- Very common for the model output and the observations to be high-dimensional and multivariate, especially spatial or time series
  - One approach: low-dimensional representation of the model output ($Y \rightarrow Y^R$) and data ($Z \rightarrow Z^R$), then carry out emulation and calibration using low-dimensional representation e.g. principal components (Higdon et al., 2008; Chang et al, 2014) or wavelets (Bayarri et al, 2007)
- Data are often non-Gaussian
  - Use spatial generalized linear mixed model for this (Chang et al., 2016)
  - For high-dimensional + non-Gaussian: principal components for non-Gaussian data
- Major challenge: as the # of model parameters increases, emulation deteriorates. We only consider 4; rest are fixed
  - How would adding in more parameters impact results/uncertainties?
  - This motivates the development of Approach # 2

Approach # 2

- Reduce the resolution of ice sheet model from 20 km to 80 km
- Use massively parallel sequential Monte Carlo approach

Computing Costs and Methods

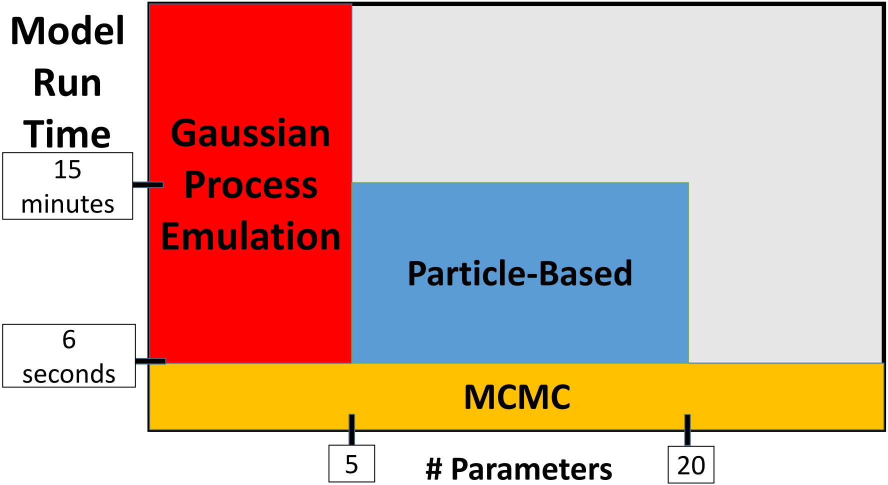

Sketch of Particle-Based Calibration

1. Sequential Monte Carlo with mutation
    - Adaptive target distributions
    - Automated stopping rules
2. Massive parallelization: 2,000+ processors on NCAR's Cheyenne supercomputer
3. Considerably reduces sequential computer model runs

Our approach is designed for

- Computer models with moderate run times ($\approx$ 6 sec to 15 min)
- $\approx 5 - 20$ model parameters
- High performance computing systems

Sampling Importance Resampling

- Idea: Sample from $q(\theta)$ + resample according to target $\pi(\theta)$
- Based on importance sampling approximation to $E_\pi g(\theta)$. Because

$$E_\pi\left[g(\theta)\right] = E_q\left[g(\theta)\frac{\pi(\theta)}{q(\theta)}\right] = E_q\left[g(\theta)w(\theta)\right]$$

for $q(\cdot)$ s.t. $\pi(\theta) > 0 \Rightarrow q(\theta) > 0$ the Monte Carlo approximation is

1. Simulate $\theta^{(1)}, ..., \theta^{(J)} \sim q(\cdot)$ and generate weights $w^{(j)}$ such that:

$$\frac{1}{J}\sum_{j=1}^{J} w^{(j)} g(\theta^{(j)}) \to E_\pi\left[g(\theta)\right]$$

2. Resample $\theta_1, ..., \theta_J \sim \{\theta_1, ..., \theta_J\}$ with weights $\{w^{(1)}, ..., w^{(J)}\}$

$$\boxed{\hat{\pi}(\theta) = \frac{1}{J}\sum_{j=1}^{J} \delta_{\theta^{(j)}}(\theta) \approx \pi(\theta)}$$

Sequential Monte Carlo with Mutation

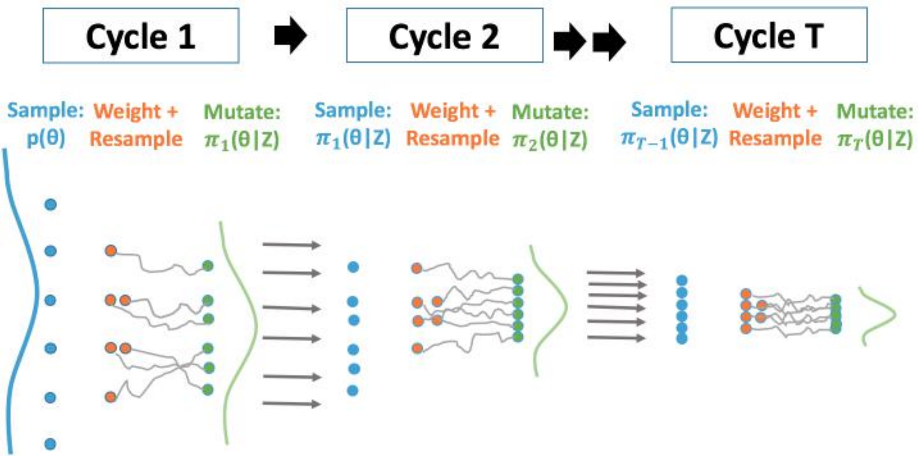

(cf. Jasra et al., 2011; Del Moral et al., 2006; Schaffer and Chopin, 2013)

Comments on Particle-based Calibration

- Easy to parallelize. E.g. each particle on a different processor
- Mutation via Metropolis-Hastings is primary driver of cost so we use an automated rule to efficiently control this cost
- Automate # resampling steps and stopping rule
- Verified via simulated examples that this method works well. For simple examples similar results to MCMC
- For analysis: priors are selected based on physical knowledge or past data; also conducted prior sensitivity analysis
- We find calibration information from the Pliocene era (5.3 to 2.6 mil years ago) impacts parameters CALVLIQ and CLIFFMAX that affect important process called marine ice cliff instability (MICI).
    - MICI: Subaerial ice cliffs exceeding 90 m in height are likely to collapse under their own weight, and could lead to runaway ice sheet retreat

Projections

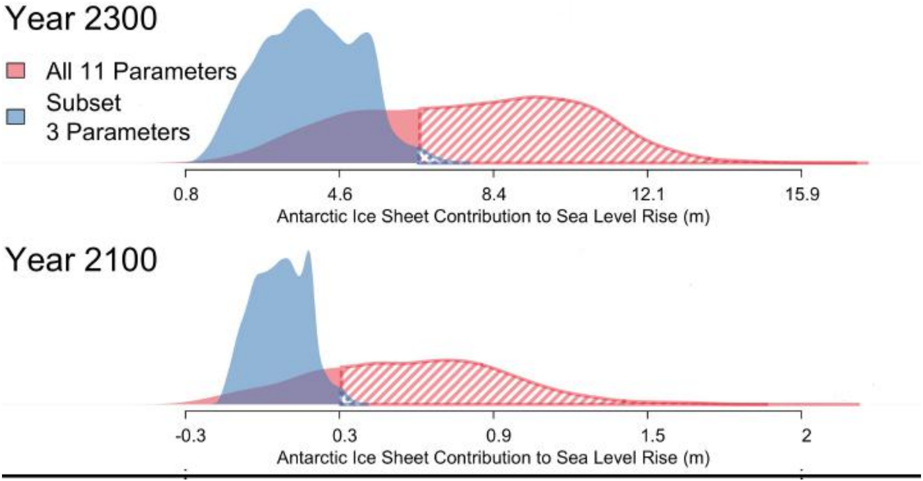

Lots of Caveats

- We focused on parametric uncertainty for the ice sheet model. There are many other uncertainties
  - Climate forcings
  - Impact of including or excluding various observations
- Model ("structural") uncertainty about the ice sheet model itself
  - Impact of resolution
  - The effect of changing the time frame

Translation: our conclusions depend heavily on the model

Summary

- Emulation-calibration
  - Approximation allows for complex, computationally expensive models
  - Emulation limits the number of parameters we can consider
- Fast particle-based approach for computer model calibration
  - Reduces wall times through massive parallelization + stopping rules
  - Can expand the number of parameters considered (11)
- Would like methods to increase the number of parameters considered + increase complexity of the model
- Caution: increased complexity may not always result in better models!

References

- Chang, W., Haran, M., Olson, R., and Keller, K. (2016) Fast dimension-reduced climate model calibration, *Annals of Applied Statistics, arXiv:1303.1382*
- Lee, B., Haran, M., Fuller, R.W., Pollard, D., and Keller, K. (2020) A Fast Particle-Based Approach for Calibrating a 3-D Model of the Antarctic Ice Sheet, Annals of Applied Statistics, 14, 2, 605-634.
- Jasra, A., et al. (2011) Inference for Levy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian J of Stat*
- Kennedy, M. C. and O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society (B)*

This work was partially supported by

- Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507.
- US Department of Energy: Office of Science, Biol & Env Research Program, Earth and Environmental Systems Modeling de-sc0016162

Thanks to Ben Seiyon Lee (George Mason University)