Statistics and the Future of the Antarctic Ice Sheet

by Murali Haran, Won Chang, Klaus Keller, Robert Nicholas and David Pollard

Introduction

One of the enduring symbols of the impact of climate change is that of a polar bear drifting in the sea, alone on its own piece of ice. For those who are left untouched by the loneliness of drifting polar bears, images of partially submerged lands and the devastation wrought by storm surges showcase some potentially frightening impacts of sea level rise on human life. The threat of sea level rise, in turn, is linked to the melting of ice sheets. Ice sheets are therefore important to understanding our planet as well as learning about how our future may be impacted by climate change. A promising approach to improving our understanding of ice sheets and derive sound projections of their future is to combines ice sheet physics, statistical modeling, and computing.

First, what exactly is an ice sheet? It is an enormous mass of glacial land ice, more than 50,000 square kilometers in extent. The Antarctic ice sheet extends over 14 million square kilometers while the Greenland ice sheet extends over 1.7 million square kilometers. To put this in perspective, the area covered by the Antarctic ice sheet is comparable to the continental United States and Mexico combined. In fact, the Greenland and Antarctic ice sheets contain over 99% of the freshwater ice in the world and, roughly speaking, melting the entire Greenland ice sheet would result in sea level rise of around 7 meters (23 feet) while if the entire Antarctic ice sheet melted, it would result in sea level rise of around 57 meters (187 feet). It is easy to imagine how even partial melting of these giant ice sheets can potentially lead to large sea level rise, making, for instance, low-lying coastal regions more vulnerable to future storm surges. Hence, a number of high profile articles and documentaries have placed the melting of ice in polar regions squarely at the center of the discussion of the impacts of climate change. Knowing something about how ice sheets are changing has very practical consequences, for instance when making decisions about how and where to build infrastructure on the coasts, and how to assess risk to property due to climate change in the future. Risk is defined in terms of probabilities (risk of an event=expected loss under a probability distribution on that event). Hence, to carefully describe the risks associated with a climatic event, we need to estimate future probabilities. Studying the future of ice sheets in a statistically sound fashion is thus of interest both from a scientific as well as policy and decision-making perspective.

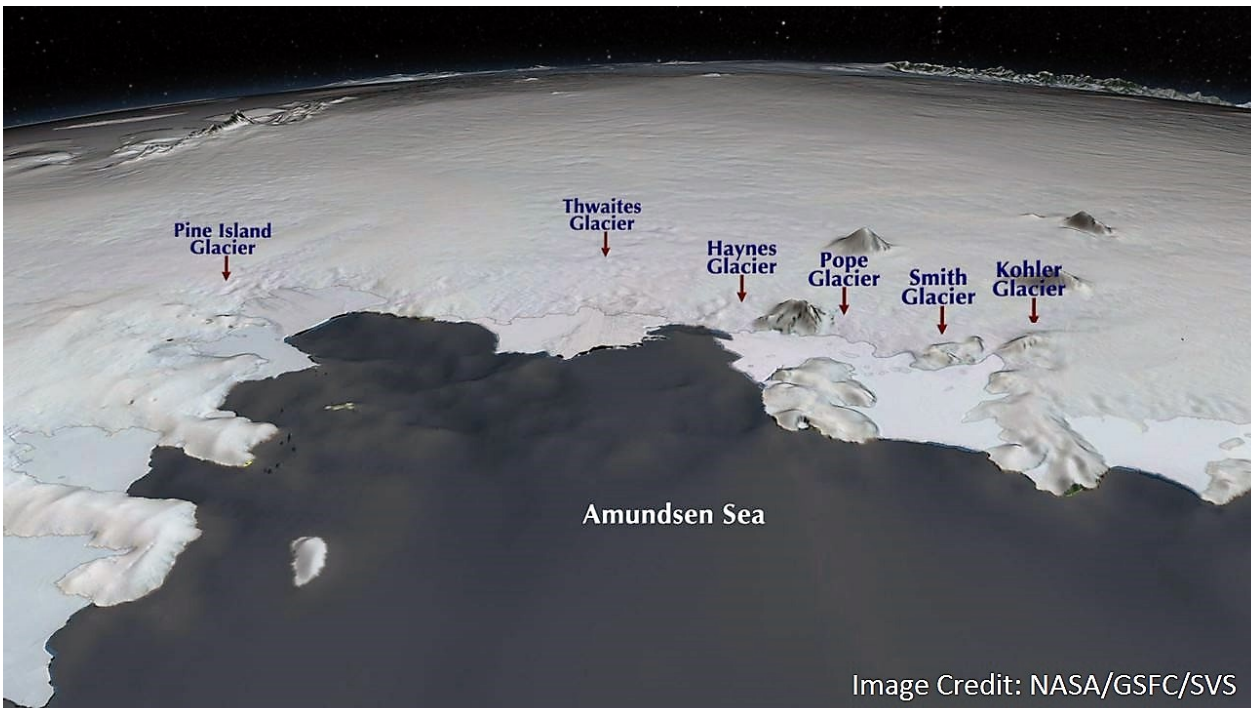

Figure 1: The West Antarctic Ice Sheet, viewed from the Amundsen Sea, with important glaciers highlighted

How do scientists study the future of ice sheets? Also, what role (if any) does statistical thinking play in studying ice sheets? A careful study of ice sheets involves four main pieces: (1) physics for modeling the behavior of the ice sheets over time, (2) computing and applied mathematics (mostly solving differential equations) for translating the model into computer simulation code, (3) data sets that are informative about the past and current state of the ice sheet and related climate variables, for instance ocean temperatures and snowfall accumulation, and (4) statistical methods that combine information from the physical model with observations of the ice sheet. This is an interesting and challenging area of research because sound scientific research requires an interdisciplinary collaboration between ice sheet modelers and statisticians. The statistical challenges involve combining information from disparate sources such as the physical model and observational data sets and the size and complexity of the data available requires careful attention to computing. This article is intended to give readers a taste of some of the interesting scientific questions surrounding ice sheets, the resulting statistical problems, as well as an outline of a statistical method that can be used to solve this problem. Our discussion is broadly targeted at ice sheets but we focus here predominantly on the West Antarctic Ice Sheet (WAIS) and PSU3D-ICE, Pollard and DeConto's ice sheet model.

The above four aspects involved in the study of the past, present and future of ice sheets are common to many other research areas in climate science where models, uncertain parameters and multiple sources of observations need to be brought together to understand the past, present and future state of the climate. In fact, similar statistical problems arise often in other scientific disciplines where complex dynamical models are used, and the applications of the statistical methods of emulation and calibration we describe here even extend to many manufacturing and engineering processes.

The Physics of Ice Sheets

Ice sheets are created by long-term snowfall accumulation. When snowfall exceeds snow melt each year, it builds layer upon layer of snow, the weight of which compresses the underlying snow to form ice. Over thousands of years, this has resulted in massive ice sheets that can be thousands of feet thick. The flow of the ice sheet is due to the height of the thick ice and snow.

Ice sheet experts have worked extensively on building physical models that describe how ice sheets flow and evolve over time. Figure 2 below provides a simplified view of the physics involved. We see, for instance, that the ice flows downslope from the highest central regions toward the edges of the ice sheet.

Figure 2 provides a sense of how the ice sheet rests on the continental crust, and how the ocean interacts with the ice sheet. The multiple parallel curves represent different ice flow lines, corresponding to different heights of ice (central regions of the West Antarctic Ice Sheet (WAIS) are over 2,000 meters high). Gravity is a fundamental driver of the flow, causing stresses and deformation that tend to flatten the sheet surface over tens of thousands of years.

The basic physical principle underlying the ice-sheet is the conservation of mass, which ensures that the local thickening or thinning of ice is balanced by ice added or removed. Ice can be added by snowfall, and is removed by *ablation*, the process by which snow or ice melts and flows away in stream-like channels and crevasses, evaporates, or is blown away by the wind.

The underlying topography as well as the slipperiness of the bedrock surface also influences the behavior of the ice sheet. Overall the ice sheet surface tends to be a smooth dome, but high peaks may protrude though the ice exposing the land in places.

The edge of the ice sheet is particularly vulnerable as it interacts with the ocean. Where the ice sheet meets or *abuts* the ocean it can form a vertical cliff, or can continue as an ice shelf, a floating table of ice hundreds of meters thick flowing out toward the open ocean. Sometimes part of the ice shelf breaks off (known as *calving*) to create floating icebergs. Individual calving events can be dramatic and spectacular especially if the ice cliff abutting the ocean is relatively tall.

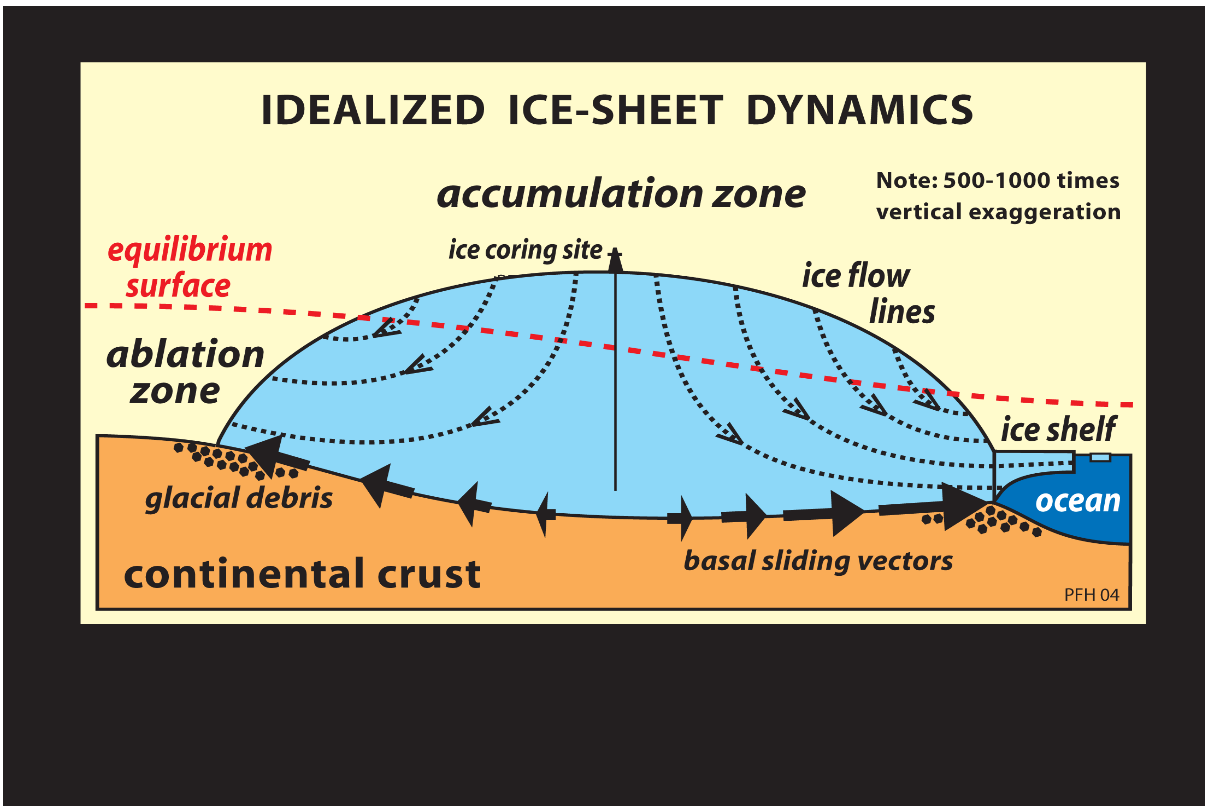

Figure 2 from <www.snowballearth.org> (courtesy Paul Hoffman)

The arrows in the figure illustrate the direction of ice flow. Ice accumulates on top through precipitation (snowfall), and flows downward due to gravity. Ablation means the melting or evaporation of ice. This cartoon illustrates how the ice sheet rests (and slides)

on the continental crust, and points out its important interaction with the ocean waters. Parameters (inputs) that determine how the ice sheet slides, and how it interacts with the surrounding ocean waters, are key to future projections of the ice sheet.

This brief outline of ice sheet dynamics shows that there are many inputs or parameters of the model that, when changed, can have a considerable influence on ice-sheet behavior. For instance, the slipperiness of the bedrock surface (the amount of friction between the ice sheet and the bedrock surface) affects how fast the ice sheet slides over it. The ocean melt coefficient is a parameter that describes the sensitivity of the ice sheet to temperature changes in the surrounding ocean. Hence, changes to this parameter will cause the ice sheet to react very differently to the changes in the surrounding ocean temperatures. Different parameter values will result in very different projections of the future of the ice sheet. Figuring out reasonable parameter values to use is therefore a very important research problem, and it makes sense to find parameter values that allow the ice sheet model to credibly reproduce both the past and current behavior of the ice sheet. In fact, parameter inference is precisely the problem we focus on here. Careful science requires that we not only provide “best” values of the parameters (point estimates) but that we also provide uncertainties about the parameter values.

Computer Models for Studying Ice Sheets

In order to study how the ice sheet behaves under various parameter settings and the impact of external climate variables or external *forcings* (physics external to the system that impact the ice sheet) on the ice sheets, scientists create computer programs that incorporate the physics of the ice sheet as well as the various forces acting on it. These days using computer simulations to learn about the behavior of an ice sheet in response to internal and external conditions is common in the earth and atmospheric sciences and is often used in many science and engineering problems. In our work on the West Antarctic Ice Sheet (WAIS), we use the PSU3D-ICE model which strikes a balance between detailed physical modeling and computational efficiency. This balance allows it to produce realistic long-term behavior of the ice sheet without attempting to incorporate very high resolution physical modeling. This allows the long runs to be accomplished with a reasonable amount of computational effort. There are many decisions that need to be made about how to run the ice sheet model. For instance, an important choice is to determine how far back we start the ice sheet model to “spin it up” to the present time (we start it 40,000 years before present). The spin up phase of the model involves running it until it reaches a “steady state” that does not, hopefully, depend too much on the initial values chosen to run the model. Another choice is the kind of external forcings (physics external to the system that impact the ice sheet) to use on the ice sheet dynamics; we use well established data sets and models to provide the atmospheric and oceanic external forcings. The computer model output is in the form of a spatial grid. We therefore also need to determine the resolution at which we want model output, with a higher resolution typically taking more computational time. Here, we simply obtain information that is close to the same scale at which the observational data sets (described below) are available. Finally, and crucially, we need to determine a study design that suggests which parameter values to use when running the model since we are constrained by computational considerations.

Ice Sheet Data

Detailed modern observations of WAIS are constructed combining many different types of observations, including satellite altimetry, airborne and ground data surveys, and ground radar surveys. These data are very useful for learning about (referred to as *constraining* in the geosciences literature) important parameters of the model. However, in order to obtain better projections of WAIS on the scale of hundreds to thousands of years in the future, it is important to also use the long-term behavior of the ice sheet to learn about the parameters. The parameters we infer must be capable of producing realistic behavior of the ice sheet over much longer periods of time. Data from the distant past, going back hundreds of thousands of years or more, are reconstructions of the ice sheet's past. These are based on recent measurements such as sonar data about ocean floor features as well as shallow sediment cores, which have been used by researchers to provide maps of approximate grounding lines – the location where the ice sheet transitions from lying on bedrock to hanging over the ocean -- at 5,000 year intervals from 25,000 years ago to the present. Hence, these resulting data are in the form of time series.

What Makes This a Challenging Statistics Problem?

Hopefully it is already clear that this is a statistical problem. Multiple data sets are involved, after all, and there is an interest in inferring parameter values and making predictions. Let us consider two important scientific questions: (1) Given the recent satellite observations of the ice sheet and the paleo-reconstructed data about the ice sheet in the distant past, what are the likely values of the ice sheet model parameters? (2) What can we say about the future of the ice sheet based on what we know about the past? We can translate the first question into the language of statistics and probability as follows: Given the two data sets and what we know about the ice sheet model by running it at various parameter settings, what is our estimate of the probability distribution of the model parameters? The probability distribution captures our knowledge about the parameters given what we knew about the parameters (prior scientific information) and what information the model runs and the observations provide about the parameters. This fits naturally into the language of Bayesian inference which allows us to combine prior information with information from the data to obtain a posterior distribution of the parameters. An advantage of this approach is that once we have an estimate of the posterior probability distribution of the model parameters, this can be used to answer the second question. Essentially all we need to do is to see what the ice sheet model projections look like at various parameter settings, and weight the probability of these projections according to the posterior distribution of the parameters. To summarize, we have: (1) a deterministic computer model that describes the ice sheet behavior as a function of parameters, but we only have simulations of this model at a limited number of parameter settings, and (2) observations of the ice sheet, both modern satellite observations as well as paleo-reconstructions of the ice sheet from the distant past. We need to formulate a *statistical* model that combines all of this information, while allowing for measurement errors and imperfections in the computer model.

The formulation of the statistical problem above may seem pretty standard, except for one important twist. For Bayesian inference we need both a prior distribution of the parameters as well as a probability model that connects the observations with the parameters. More specifically, the probability model provides a probability distribution for the observations – in this case the satellite data and the paleo-reconstructed data – at each parameter value. This probability model is used to obtain a likelihood function, and then the rest of it is (modulo computational challenges), routine Bayesian inference. Here the only connection we have between the parameters and the observations is via the ice sheet model. This poses some challenges: (1) the model is deterministic, not probabilistic, so it does not provide a probability model on its own; (2) we only see the model output at a few (relatively small number of) parameter values; and (3) we know that the ice sheet model is an imperfect representation of the observations. The problem and an outline of how we can think about solving it, is summarized in Figure 3 below. What makes the problem challenging is the fact that the output from the model is high-dimensional and in the form of spatial or temporal data. These data are always not easily modeled using Gaussian models. New computationally efficient statistical methodology is therefore necessary for addressing these issues.

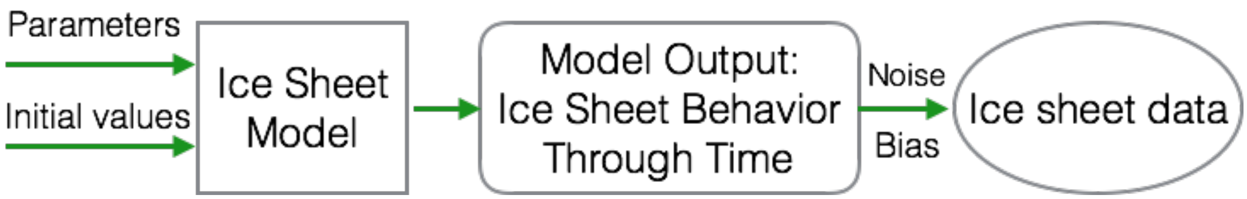

Figure 3: Parameters and initial values drive the ice sheet model. Its output describes the behavior of the ice sheet through time. Because this is an imperfect model, we account for noise (measurement error) and biases (missing processes in the model) to develop a model for the ice sheet observations. Note that the ice sheet model is a "black box" -- we only see model output for any given set of parameters. Example parameters include those that determine the basal sliding of the ice sheet, and the sensitivity of the ice sheet to the surrounding ocean water temperatures. Emulation approximates via a Gaussian process how the ice sheet model above maps parameters into model output. This approximate model combined with a model for error and bias is used as a statistical model for the ice sheet data (observations) on the far right.

Ice Sheet Model Emulation and Calibration

How do we solve this problem? We can do this in two main stages. We first approximate the ice sheet model with a statistical model, that is, develop a statistical model that can predict how the ice sheet model will behave at new parameter values. Think of this problem as one where we need a flexible regression-type approach: given lots of predictors (various parameter settings) and corresponding model output (responses), we can predict model output at new predictors (any new parameter setting), along with some uncertainty about the model output. This kind of uncertainty may be referred to as epistemic uncertainty, meaning that the uncertainty arises from our lack of knowledge (*episteme* is Greek for "knowledge") about what the model will do, *not* the fact that there is anything random associated with the model (it is deterministic!). This process of approximating the model is called **emulation**. Emulation results in a probability model that links the parameters of the ice sheet model to the output of the ice sheet model. The statistical model we use for emulation is a Gaussian process, a popular model in spatial statistics, which well suited to interpolating functions. Consider the simple example in Figure 4 where we consider a collection of random variables that are a function of a single parameter. There are of course an infinite number of these random variables on any given range of parameter values, say between 0 and 1. A Gaussian process model states that any finite collection of random variables, for example the six function values between 0 and 1 (black dots on Figure 4), has a joint normal distribution. Crucially, the dependence among the random variables decreases as a function of the distance between them, making two random variables that are close to each other (in parameter value) more dependent, and hence more alike. This suggests how Gaussian processes provide a nice approach for interpolation – the predicted value for a random variable at any parameter value, say a function value in between the six black dots, is more like (more dependent on) values that are close to it, and depends less on values that are far from it. The precise dependence between the random variables at various parameter values is controlled by a covariance function that describes how covariances change as a function of distance. Hence, Gaussian processes provide a simple and effective way to interpolate a function, using dependence, without the need to determine a specific form for the function. The idea we just outlined extends in principle to functions of multiple parameters as well. Figure 4 (right) shows what a Gaussian process interpolator produces for a toy example with only one input parameter.

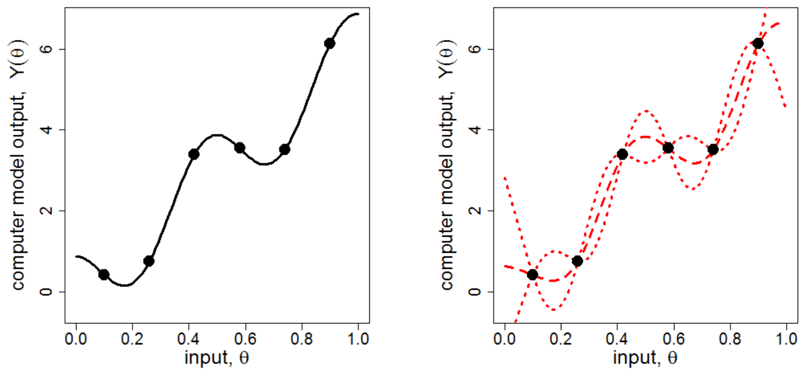

Figure 4: Emulation for a toy example. Left: Black dots correspond to input parameters for which the computer model was run. Right: Red dashed lines are interpolations by a Gaussian process – they provide approximate computer model output at every parameter value. The dotted red curves correspond to uncertainties; there is greater uncertainty as we get further from places where we have data.

We need the model for the observations to allow for the fact that the ice sheet model is not a perfect representation of the observations of the ice sheet. For this, we add a component to the model for errors (variability in the observations) and sources of systematic biases, called a model-data "discrepancy" term. Once we put these pieces together, we have a model that is potentially useful for observations of the ice sheet; this serves as the probability model for the observations given the parameters. Figure 5 shows how calibration works for a toy example where the model output is just a scalar and the observation consists of just a scalar as well.

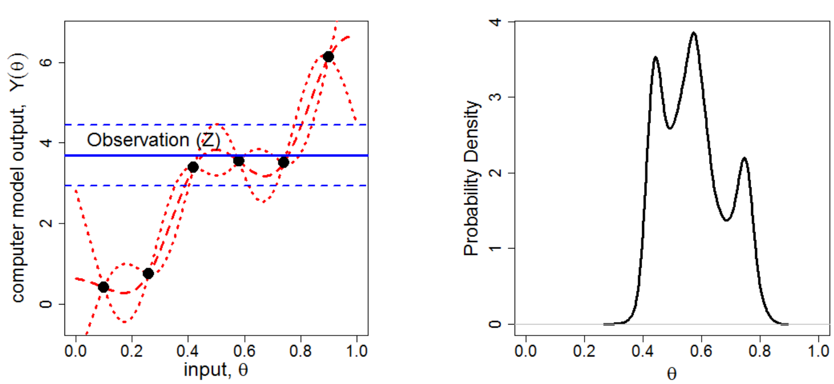

Figure 5: Calibration for a toy example. Left: The blue horizontal line represents a single data point. Calibration attempts to find parameter values that are compatible with that observation, while also taking into account uncertainties due to variability represented by the blue dotted lines). Right: Bayesian inference provides the (posterior) distribution on the right, which summarizes what we know about the parameter. Notice that there are three peaks in this density, corresponding to three black dots (left figure) closest to the observations.

We can summarize the entire approach as follows.

(1) **Generate an ensemble of model runs**: Run the ice sheet model at various parameter settings. This gives us pairs of parameters and model output, just like in a regression problem.

(2) **Emulate the ice sheet model**: Use a statistical model to approximate the relationship between the parameters and the model output. This is similar to fitting a flexible regression model, except the response is multivariate, spatial (satellite data) and temporal (paleo-reconstructed data).

(3) **Construct a model for the observations**: This is the fitted Gaussian process model + a model for errors and biases. We only specify the form of the errors and biases, their parameters still need to be inferred from the data (from Step 4 below).

(4) **Calibration**: Fit the above model to the observations. This gives us a distribution on the parameters, while also providing some information on the errors and biases.

(5) **Project the future of the ice sheet**: Use the posterior distribution on the parameters to run the model forward and provide the future of the ice sheet in the form of a ("posterior predictive") distribution.

Of course, here the model output is quite a bit more complicated than a standard regression response because the model output is a map of the current ice sheet (this is a spatial data set) along with information about the ice sheet's past over time (this is a time series data set). Also, the relationship between the parameters and the model output can be quite complicated. There are additional complications because the data tend not to be Gaussian, for example the ice sheet data are modeled as presence-absence. We use a spatial generalized linear model version of a Gaussian process, which allows us to approximate the deterministic model with non-Gaussian output by a probabilistic model. The high-dimensionality of the data also necessitates some dimension-reduction approaches. We use a principle components analysis-based approach. The Gaussian process methodology is remarkably flexible, allowing us to emulate the ice sheet model quite well. How well it does can be studied by using cross-validation, for instance by leaving ten percent of the model runs (parameter settings) out when fitting the Gaussian process model to the ice sheet model runs, then looking at what the Gaussian process model predicts for the parameter settings that were left out. If it resembles what the model actually outputs at those parameter settings, it suggests that emulation is working well. Figure 6 illustrates this well.

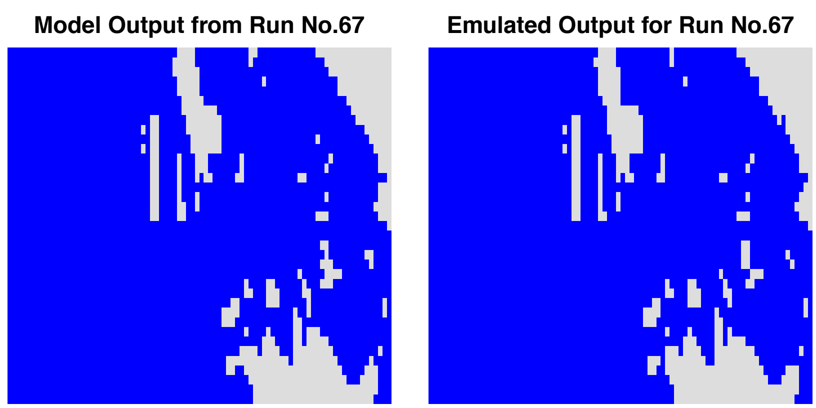

Figure 6: Comparison of actual model output (left) with emulated model output (right). Blue corresponds to "no ice sheet" and white corresponds to "ice sheet presence". The emulator is able to mimic the model run very well.

Results

When we use emulation and calibration methods to these data and models we obtain parameter estimates and resulting probability distributions for future projections. This is summarized in Figure 7, which shows the distribution of potential sea level rise due to the melting of WAIS in 500 years. Calibration with both the modern and paleo data results in different sea level rise projections (red curve, “current approach”) when compared to projections with calibration using just the modern data (Dash-dot blue curve “modern obs only”). In particular, using both sources of data eliminates any possibility of there being no sea level rise, that is, the value 0 is included in the distribution for just the modern data, while it is essentially excluded when both data are used.

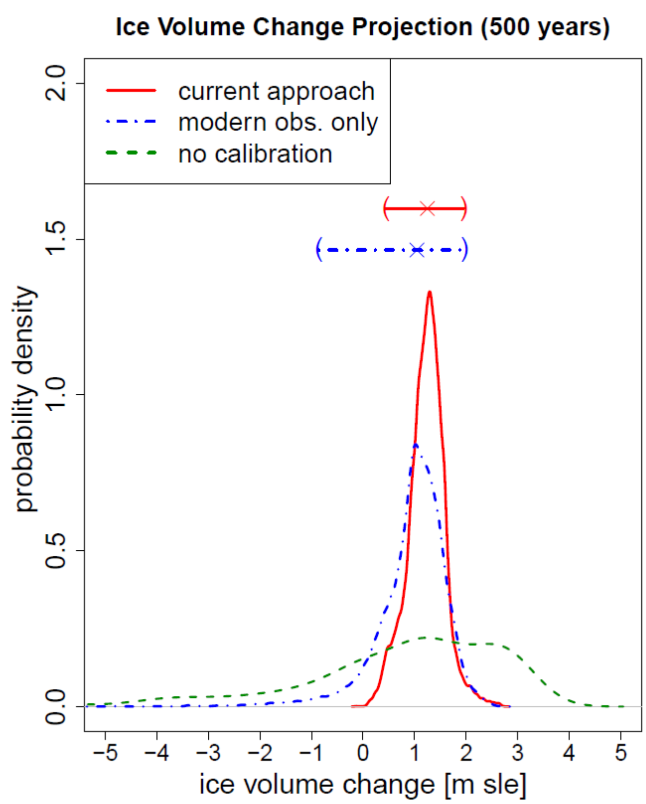

Figure 7: Posterior predictive distribution of projected ice sheet contribution to sea level rise. Adding the paleoclimate data results in a much sharper projection (red curve) than

when only modern satellite data are used (blue curve). In particular, the possibility of no (zero) sea level rise due to ice volume change is virtually eliminated in the red curve. (Reproduced from Chang et al. (2016) *The Annals of Applied Statistics*)

Our research shows that sea level rise is inevitable, though our results are relatively conservative in stating that it is most likely to be around 2 meters. Even 2 meters of sea level rise will leave many low-lying regions in the world completely submerged, and would put many more areas at high risk of potentially devastating storm surge damage, for instance the Netherlands and the Maldives, and future storm surges are likely to cause much greater devastation through flooding. Recently-developed models that incorporate a few additional features of the ice sheet dynamics suggest that sea level rise may be even more dramatic.

Caveats

With all the complicated sources of information that have gone into this research, we have to be cautious about our conclusions. The ice sheet model does not include all the processes that affect the ice sheet. There are uncertainties in the paleo data that have not been accounted for. Furthermore, the ice sheet model will behave differently for different initial values; ideally, we would incorporate uncertainties due to this variation too. Similarly, there are a number of different ways in which external forcings, climate variables that are external to the ice sheet, may change over time. These also impact how the ice sheet behaves. Incorporating all these uncertainties is daunting largely because of the computational challenges involved. Hence, whatever we say about the behavior of the ice sheets in the future is necessarily imperfect. Having said that, the information we have summarized here incorporates cutting edge physics, multiple observation data sets, and pieces the information together in a principled manner. Hence, in spite of all these caveats, we have made progress! To quote Einstein, all our science, measured against reality, is primitive and childlike -- and yet it is the most precious thing we have.

Further Reading:

(1) General reading on ice sheets: From the National Snow and Ice Data Center: https://nsidc.org/cryosphere/sotc/ice_sheets.html

(2) General reading on Antarctic glacier physics: http://www.antarcticglaciers.org/modern-glaciers/introduction-glacier-mass-balance/

(3) Much of the statistical methodology outlined here is based on this paper: Chang, W., Haran, M., Applegate, P. and Pollard, D. (2016) "Improving Ice Sheet Model Calibration Using Paleoclimate and Modern Data," Annals of Applied Statistics, *10, 4, 2274--2302*.

(4) The relationships between ice-sheet projections, flood risks, and decision-making are discussed, for example, in Wong, T., A. Bakker, and K. Keller (2017) Impacts of Antarctic Fast Dynamics on Sea-Level Projections and Coastal Flood Defense. 144, 2, 347-364; Bakker, A., D. Louchard, and K. Keller (2017) Deep uncertainties surrounding sea-level projections: Sources and Implications. Climatic Change Letters, 140, 3, 339–347; Diaz, D., and K. Keller (2016) A potential disintegration of the West Antarctic Ice Sheet: Implications for economic analyses of climate policy. American Economic Review, 106, 5, 607-611.

About the Authors:

Murali Haran is a professor in the Department of Statistics at Penn State University.

Won Chang is an assistant professor in the Department of Mathematical Sciences at the University of Cincinnati.

Klaus Keller is a professor in the Department of Geosciences, Director of SCRiM (Sustainable Climate Risk Management) at Penn State University, and an adjunct professor at the Department of Engineering and Public Policy, Carnegie Mellon University.

Rob Nicholas is a research associate at the Earth and Environmental Systems Institute and Managing Director of SCRiM at Penn State University.

David Pollard is a senior scientist at the Earth and Environmental Systems Institute at Penn State University.