Author Response to RC3: "Computationally efficient subglacial drainage modelling using Gaussian Process emulators: GlaDS-GP v1.0"

Tim Hill, Derek Bingham, Gwenn E. Flowers, Matthew J. Hoffman

Reviewer comments are in black and we provide our responses in blue.

This paper describes a Gaussian Process emulator of the GlaDS subglacial drainage model and its testing on a synthetic ice-sheet margin setup. Modelling subglacial drainage is starting to become an important aspect of ice dynamics simulations as that system impacts the basal boundary condition significantly. However, subglacial drainage models are relatively costly to evaluate and in particular operate on different, shorter time scales compared to ice flow. Thus running coupled ice-flow drainage simulations is typically difficult and costly at the moment. Emulating the subglacial drainage model using a statistical representation is likely an important step in making these types of coupled models readily applicable.

Whilst emulations of GlaDS with neural network based emulators have been achieved over the last few years, this is the first time a Gaussian Process based emulator has been put forward. The advantage of GP emulators is their greatly reduced number of parameters to fit compared to a neural network as well as built-in capability to quantify uncertainties of the emulation.

The manuscript lays out the procedure to construct the GP emulator; of note is that this construction is relatively involved as it also entails, for instance, decomposition of the GlaDS training data into principal components, fitting of hyperparameters using Bayesian schemes, etc. The emulator is then tested extensively on a synthetic setup and the authors discuss the pros and cons relative to neural network based emulators.

The study and manuscript are carefully constructed. As I am not an expert in statistical emulators, I cannot judge the appropriateness and correctness of the approach to implement the GP emulator. The testing and assessment of the emulator is certainly fine and the discussion is interesting and relevant. Thus, with above caveat, I recommend to publish this manuscript in GMD with the minor corrections outlined below.

Thank you for the detailed and constructive review. We have responded to your comments individually below.

Comments

I think it would be useful to discuss a bit more how this emulator could be used for inversions or for coupled ice-flow & drainage simulations as, in my opinion, this are the most sought after usages of such tools. This can just be in the Discussion and/or Introduction, no need for more simulations or an implementation.

Thank you for the suggestion. We have expanded Section 6.6 "Applications and considerations" to describe Bayesian inference of subglacial drainage model parameters as an appealing direct extension of this work. We have described the steps needed to use the emulator for coupled ice-flow and subglacial drainage simulations and highlight some of the additional uncertainties related to the basal slip relationship and the ice flow law that could be addressed by extending the current work.

The construction of the emulator has many steps. Looking through the manuscript, I can see:

- training data construction using parameter design matrix

- running the simulations with GlaDS

- principal component decomposition and component selection or (reduction of variables to scalars)

- fit the GP emulator to the data using an MCMC scheme

Then using the GP in different ways for predictions and analysis is then yet another step. Would it make sense to somehow graphically represent this, flow-chart or some such? Or maybe a numbered list?

Thank you for suggesting ways to make the construction of the GPs more accessible. We have designed and added the following summary of the steps involved in the emulator construction.

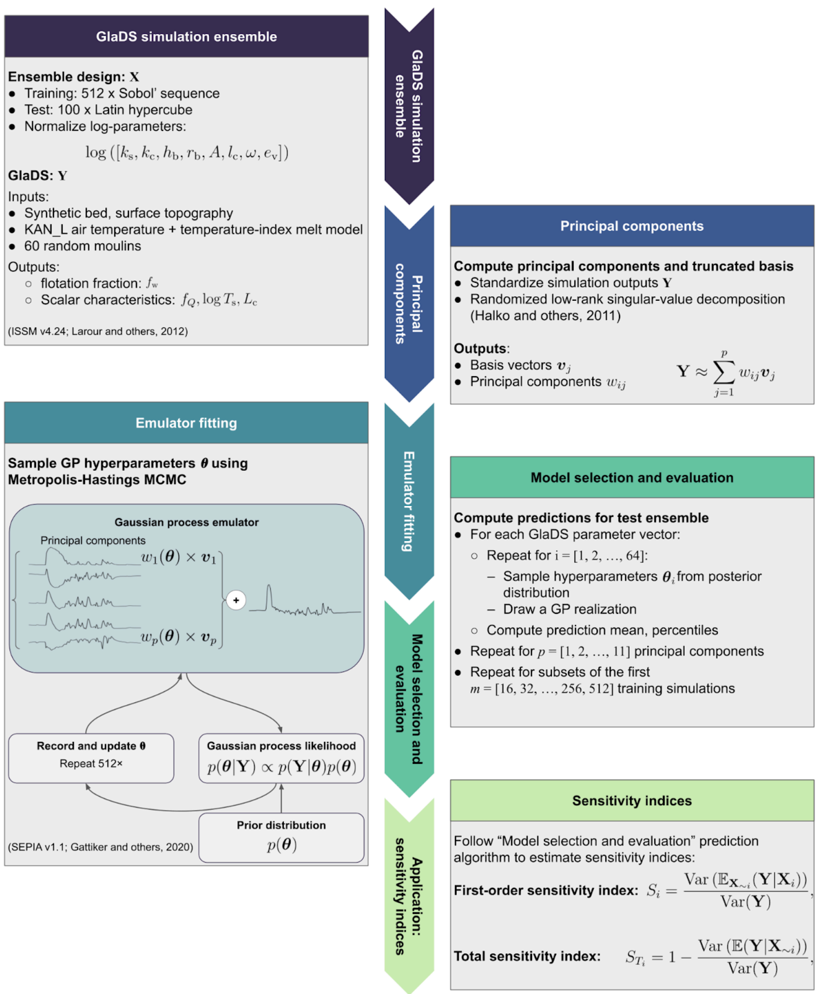

Figure: Overview of steps involved in constructing the Gaussian Process emulators. X is the design matrix of GlaDS parameters (defined in Table 2) with corresponding GlaDS outputs Y. The Gaussian Process emulator is constructed as a truncated linear combination of p principal components $w_i(\theta)$ and basis vectors v_j for i = 1, . . . , p, where θ are Gaussian Process hyperparameters that are inferred by Markov Chain Monte Carlo (MCMC) sampling. Emulators are fit using *m*-member subsets of the training data and constructed using different numbers of principal components *p*. The performance is evaluated on the independent set of 100 test simulations. The emulator is used to compute the sensitivity of model outputs to model parameters (Section 5).

Irrespective of the lack of such a graphical overview, I struggled to understand the GP emulator from the description. I am not sure whether I should expect to understand GP emulation from reading about it in such a publication or whether I should just need to go elsewhere to learn it. I see that the authors try to keep the reading smooth by moving quite a bit of the explanations to the appendix but I wonder whether that makes it even harder to understand as now the content is disjoint? Maybe if this layout is kept, then make it even more high level in the main text and have the full description in the appendix which then could be in one place; or, alternatively, move all into the main text? In fact, I think that would be my preferred option and, I think, would fit GMD well as this journal is mostly about methods and not science. As it is, I think it is a bit of a difficult split.

Thank you for highlighting that Section 2.2 was not as accessible to non-experts as we had intended and for suggesting improvement in the content and structure. We have expanded Section 2.2 to integrate the content from Appendix B so that the reader has all the information in one place. We have also expanded the high-level description of the terms in the equations and defined all statistical terminology. At the beginning of the GP section, we have also clarified our intention to provide a high-level overview with only the details necessary to understand our application of the method and the differences compared to other statistical models (e.g., neural networks):

> “This section briefly provides a high-level overview of the Gaussian Process (GP) model and the architecture that we use to emulate spatially and temporally resolved GlaDS outputs. For background on Gaussian Processes seeJones et al. (1998) and Rasmussen and Williams (2005), and see Higdon et al. (2008) for a complete description of the emulators constructed here.”

The authors state the principal component decomposition will make the representation necessarily smooth (line 170). Around the channels the hydraulic potential is often not smooth but has the channel as a kink, is that a problem (i.e. a spatial non-smoothness)? Also related to smoothness: in setups like the one presented, where there is no lateral variation in topography, channel position is not necessarily stable with parameter variation but they can jump around (and, for certain, channels move if the mesh is varied). Is that a problem for GP?
The perturbation in hydraulic potential (or more precisely flotation fraction for our work) near a channel is not a problem for the principal component-based GP. The spatial and temporal variations themselves do not need to be smooth since this complexity is encoded by the basis, which has no smoothness constraints. What the GP requires is that the principal components (w_{ij} in Eq. (9)) vary smoothly with respect to the GlaDS parameters. We have tried to explain this more clearly: “While the flotation fraction field need not be smooth in space and in time, the principal components w_{ij} (θ) tend to vary smoothly with

respect to the GlaDS parameters since the the spatiotemporal complexity is captured by the principal component basis.”

The comment about unstable channel positions is interesting. It’s possible that “boundaries” in parameter space that cause changes in channel position would be reflected as discontinuities in the principal components. This would show up as simulations with unusually high error when evaluating predictions on the test set. We have not seen evidence of such issues.

Line-by-line

L4: "the combination of the number" is not clearly formulated. Reword.
We have revised this sentence to read “While they are used to understand processes such as the relationship between surface melt and ice flow, the number of uncertain model parameters and the computational cost of running models makes it difficult to [..]”

L8: "daily representation" is not clear to me. Maybe "diurnally averaged"?
Thank you for the suggestion, we have updated the text as suggested since “diurnally averaged” is more accurate.

L66: I would cite the ISSM GlaDS implementation here too, I think that is Ehrenfeucht&al 2023.
Correct, we have added this citation.

L83: "see B" -> "see Appendix B"
Corrected.

L84: "fast predictions" is a bit sloppy, they are fast to run but not fast themselves.
That is correct, this sentence has been simplified to say: “Following tuning and evaluation of the emulators, we apply them to compute the sensitivity of model outputs to parameters.”

Table 2: r_b is not defined in the original GlaDS paper nor in this manuscript. Needs to be defined, at least in Appendix A.
Thank you for the suggestion. We have added the definition of the aspect ratio r_b as the ratio of the bump length l_b to the bump height h_b, such that the aspect ratio should be roughly >1, following equation (A2).

L108: state here that theta is what is fitted and maybe also state the (approximate) size of theta.
We have expanded the description of the GP hyperparameters: “The hyperparameters typically control the variance of the Gaussian Process and the sensitivity to each input, but their interpretation depends on the type of covariance function that is used. Gaussian Processes typically have a similar number of hyperparameters as inputs to the emulator. The hyperparameters must be optimized to obtain an accurate emulator.”

L110: "The second choice" really needs a clear statement above of what the first choice is (namely k), otherwise the reader will stumble over this.

Thank you for the suggestion, since we do not clearly articulate the covariance function as the first choice, we have removed the language about "the second choice". Instead, this paragraph begins with "We make the common choice to set the prior mean to zero everywhere…"

L124: x is not defined, or if its definition is "prediction input", then that is not clear enough.
We have defined $\boldsymbol{x}$ on line 88 as the vector of GlaDS model parameters. Since using both "model parameters" and "inputs" to refer to the same thing is confusing, we have referred to $\boldsymbol{x}$ as model parameters throughout (see also response to reviewer 1).

L127: the "posterior distribution" comes out of the blue here
Thank you for highlighting this. As detailed in our response to your third comment (clarity of the GP exposition), we have revised this section to more fully explain the posterior distribution and posterior predictions, keeping in mind to explain these statistical terms.

L130: are there d+1 hyperparameters for any k? Couldn't it be less as well? Or more?
The fact that there are d+1 hyperparameters is specific to how we have written the covariance function k. Since we have not yet introduced a particular covariance function k it is not appropriate yet to provide a specific length of the hyperparameter vector. We have revised this sentence to remove the precise number of hyperparameters: "the fact that the GP model is simple enough to allow for Bayesian inference is a key advantage compared to a neural network for uncertainty quantification."

L223: A negative floation fraction implies negative water pressure, right? But how can the water pressure go negative in the presented setting? I don't think it can drop below the value of the Dirichlet BC which corresponds to zero water pressure.
This is a good question, and the reviewer is correct that negative flotation fraction implies negative water pressure. The negative flotation fraction (and water pressure) happens during the melt season in response to a rapid drop in surface melt rates and only lasts for one to a few days. We have clarified that this is a transient issue: "We have found that broadening the parameter ranges results in numerous nonphysical simulations with nearly zero water pressure during the melt season, transient negative flotation fraction as low as f_w<-10 or extremely high flotation fraction as high as f_w >> 100 which degrade the performance of the principal component decomposition."

L274: RMSE is not defined yet. But then it gets defined in L296.
Thank you for highlighting, we have defined root mean square error at the first instance of RMSE.

Fig 2: state to which fields the PCs are encoding
The principal components (a) and basis vectors (b) represent the spatiotemporal flotation fraction field. We have added "flotation fraction" to the caption related to (b): "Width-averaged representation of the first seven **flotation fraction** f_w spatiotemporal principal component basis vectors [...]"

L283: It would be nice to have some snapshots of the PC fields and the GlaDS fields side by side (probably in the appendix). So similar to Fig 2 panel b, but not width averaged but instead just a few instances in time. This would allow to get a bit of a feel on how accurate the spatial fidelity of the PCs are.

This is a good idea, it will provide some intuition of how the PCs and GP behave. We have added the following figure to Appendix C. We would also like to clarify that Fig. 2b does not show the PC low-rank representation of the GlaDS-simulated flotation fraction. Fig. 2b shows the principal component basis, $\boldsymbol{v}_j$ in Eq. (9). We have referenced $\boldsymbol{v}_j$ and Eq. (9) in the Fig. 2 caption to minimize confusion about the quantity shown in (b).

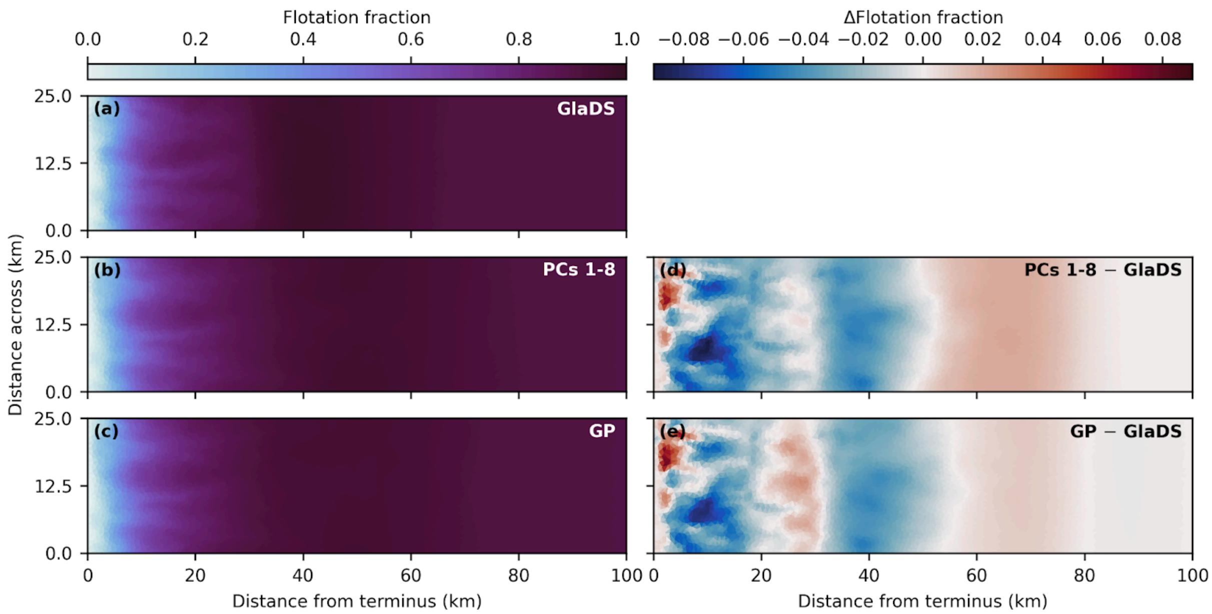

Figure: Principal component truncation error and GP prediction error. (a) GlaDS-simulated flotation fraction for the test simulation with median RMSE on 29 July, (b) corresponding principal component representation of the GlaDS flotation fraction using 8 PCs, and (c) Gaussian Process (GP) emulator prediction. Difference maps show the principal component representation (d) and the Gaussian Process emulator prediction (e) minus the GlaDS output.

Fig 4: Eyeballing the convergence of the two error metrics (panel a,b,d,e), it looks like that the errors do not go to zero but approach some non-zero asymptote. Is that expected? If so, why? Maybe this could be briefly mentioned in the text.

This is correct and expected. There are errors in the GP predictions from two main sources:

- Basis truncation error: using 1–11 PCs obtains only an approximation of the full set of simulations (e.g., Fig. 1a)
- GP prediction error adds to the basis truncation error. The GP is only be expected to be a perfect predictor of the principal component representation of the data in the theoretical limit of infinite training runs

We have added a brief explanation where we present the results from Fig. 4 (now numbered Fig. 5): "Figure 5 seems to suggest that the RMSE and MAPE are converging to nonzero values. This is an expected outcome since the total error represents the sum of the basis truncation error associated with using at most 11 PCs (Fig. 3) and error in the GP predictions of the principal components"

L394: "supporting the interpretation of PC1 as representing water pressure in the absence of surface melt inputs": to me Fig2b1 shows that PC1 has a clear seasonal signal which the basal melt does not. So, I'm not sure this statement is correct or at least needs some more information.

Thank you for the question. While the first basis function (PC1 in Fig. 2b1) does indeed have a clear seasonal signal, Fig. 2b1 shows that PC1 “turns off” by being nearly 0 in the melt season and especially at lower elevations, so that PC1 does not contribute to the surface melt-forced drainage system. PC1 is consistently “turned on” with absolute values >1 during winter and above the maximum melt extent. From this, we argue that PC1 mostly represents subglacial drainage in the absence of surface melt. It turns off when surface melt dominates the drainage system, allowing other PCs to dominate at these times. We have expanded our explanation where we propose an explanation for PC1:
“Based on the first PC basis vector being nonzero in winter and upstream of the maximum surface melt extent (~80 km)**, and not contributing to the solution at low elevations during the melt season,** the first and most important PC in terms of its explained variance (80.3%) appears to control the baseline water pressure in the absence of surface melt inputs”

L444: Formulating more clearly what "in ice-flow modelling" means would be helpful
We have added the explanation: “if the emulated fields were used as part of the basal boundary condition for ice-flow modelling”.

Tab5: here the typesetting seems a bit off: in the fields spanning multiple lines, the line spacing should be less than between different rows.
We have had to adjust the formatting of Table 5 to force it to fit on a single page. We will ensure that the formatting is correct in the final typeset version of this table.

L552: ideally a DOI and stable archived version of SEPIA and ISSM should also be provided. At the very least the version of ISSM used needs to be stated.
We have used ISSM version 4.24, which is available as a release on GitHub (https://github.com/ISSMteam/ISSM/releases/tag/v4.24), but not with a DOI. We have added “v4.24” to the code and data availability statement. Since SEPIA and ISSM are not our code to archive, we have provided the best publicly accessible links that we can.

L554: the air-temp dataset needs to be clearly specified. The provided link points to very many datasets. Do note that this data-repository provides DOIs for each dataset.
The correct link to the Greenland weather station data is https://doi.org/10.22008/FK2/IW73UU and the text has been amended accordingly

Eq A2: I would expect r_b to feature here.
Thank you for highlighting this mistake, this equation (and the following equation for time-evolution of hydraulic potential) has been corrected to include the bed bump aspect ratio r_b instead of the bed bump length l_b:

$$\frac{\partial h_{\mathrm{s}}}{\partial t} = f_{\mathrm{b}} \frac{h_{\mathrm{b}} - h_{\mathrm{s}}}{h_{\mathrm{b}} r_{\mathrm{b}}} u_{\mathrm{b}} - \tilde{A} h_{\mathrm{s}} |N|^{n-1} N$$

The same correction has been made to Eq. (A3).

L613: "maximing" -> "maximising"
Corrected.