Replies to referee comments on "Investigating the thermal state of permafrost with Bayesian inverse modeling of heat transfer"

The Cryosphere, `10.5194/egusphere-2022-630`

RC: ***Referee's Comment***, AR: Authors' Response, ☐ Manuscript Text

2. Response to referee 2

We are grateful to the referee for their time and effort in providing valuable feedback to improve our work. The referee made several good points regarding the need for a sensitivity analysis as well as helpful suggestions regarding the description of the method, the clarify of the text, and references to alternative methods in the ensemble inversion literature. We respond to the referee's specific comments and concerns in-line below.

RC: ***I regret the absence of a more through sensitivity analysis for the parameters that were not included.***

AR: We agree that a parameter sensitivity analysis would be a valuable addition to our work. We are happy to report that we have re-run the simulations with a larger ensemble size ($N = 512$) and have produced sensitivity analysis plots for each site (see Fig. B1-B4) using the EASI method of Plischke (2010) as implemented by Dixit and Rackauckas (2022). We will include these figures along with an accompanying discussion in the appendix of the revised manuscript.

The only model parameters not included in our analysis are the geothermal heat flux (this is discussed further below), constituent material thermal properties (conductivities and heat capacities), and physical constants. We also exclude soil composition parameters in layers where there is strong a priori reason to do so (e.g. organic content is excluded from deeper layers). We will add a table summarizing these constant parameters and their values in the supplement of the revised manuscript.

Of these excluded parameters, the most influential are certainly the constituent thermal properties, and in particular, the constituent thermal conductivities. While, in principle, it would be better to include these parameters in the analysis, the current parameterization of the bulk thermal conductivity and heat capacity in terms of the constituents would lead to colinearity with the soil composition parameters (i.e. organic content, porosity, saturation, and excess ice) that are varied already in our analysis. This is both numerically problematic and functionally redundant since the bulk thermal properties can already be adjusted by varying the soil constituent fractions. Choosing only one of these two sets of parameters (i.e. fixing the constants and varying the composition, or vice versa) avoids this problem.

We should note, however, that this approach would not be appropriate if the primary goal of the inversion were to recover the correct thermal properties or composition parameters from temperature measurements. This would require a more careful analysis such as the hierarchical approach in Wang and Zabaras (2005) as well as higher resolution temperature measurements (we use only annual means in this study). In this work, however, we are primarily interested in the predictive densities and the corresponding relationship between temperature and latent heat rather than the exact parameter values.

RC: ***L65-73. I regret that other solutions to this issue are not discussed at all. If EKS is one possible avenue, others have been proposed, such as the combination of efficient multiple-chain McMC algorithm with reduced dimension representation of the parameter space (Laloy et al., 2018) , or bypassing the inverse problem by directly predicting the posterior distribution from simulation-based machine learning approaches (Thibaut et al., 2022 for example in hydrogeology dealing with the same type of prediction (temperature field)). Since this aspect is also included in the discussion, it could be interesting to expand the perspectives beyond the technique used in the paper (i.e. EKS).***

AR: We thank the referee for the additional references. In particular, we were not aware of the work of Thibaut et al. (2022) and Hermans et al. (2018) on "Bayesian Evidential Learning" which looks like a very promising alternative to the inversion method used in this work. We will add some discussion mentioning these methods as well as others such as the Null-space Monte Carlo method used by Harp et al. (2016) and the more general sampling method of Sequential Monte Carlo (Sisson et al., 2007; Doucet et al., 2001) which is commonly used in other similar data assimilation tasks (Kantas et al., 2014; Noh et al., 2011; Moradkhani et al., 2005).

RC: ***L133-134. The prior does not encode information about Y, it encodes information known about the unobserved quantities (X), before the data Y is actually collected and thus correspond to what we know and don't know about X before the experience.***

AR: This was a typo in the text which we have now fixed.

RC: ***Section 3.2. It was not directly clear to me that a Bayesian approach was applied to the trend analysis. Maybe it could be more explicit.***

AR: We will clarify the Bayesian aspects of the trend analysis in the revised text.

RC: ***L134-136. Is it? Since this integral is a constant for a given problem, the comparison of the likelihood (ratio) is sufficient to sample the posterior distribution (see rejection sampling or Metropolis sampling) and the integral is not such a problem. The main problem lies in the computation of the likelihood p(Y|X) which requires to solve the forward problem, generally through numerical approximation of partial derivative equations.***

AR: In the original text, we were referring to the historical difficulties of applying Bayesian methods that predated the development of numerical sampling based inference algorithms. These methods were not, to our knowledge, widely used for Bayesian inference until the work of Geman and Geman (1984) and then were further developed by Gelman et al. (1995) and Neal (2011) among others.

As discussed later in Sec. 3.5, the primary impediment to the application of modern sampling methods to simulation-based problems like the one in this work is the cost of the likelihood as described by the referee. We have rephrased this section to make this more clear:

> Difficulties in the practical applications of Bayesian methods have historically arisen from the intractability of the integral in the denominator of (2), often referred to as the *marginal likelihood* $p(Y)$. However, advances ~~in numerical sampling methods~~ over the last few decades ~~[...]and the~~ in numerical sampling methods such as Markov Chain Monte Carlo (MCMC) [...]which sidestep the need to compute the marginal likelihood, as well as a general increase in available computing power, have made Bayesian methods significantly more accessible.

RC: ***L190-191. I am wondering about the effect of this fixed boundary. Given the effort for modelling the uncertainty on the upper BC, why not also considering uncertainty on the bottom one? This flux is certainly not known for sure and it could contribute to significant uncertainty at depth.***

AR: It is true that the geothermal heat flux at the lower boundary can have substantial effect in some cases when (i) it is very large, e.g. in a volcanic area, (ii) the simulation period is very long (centuries or longer), or (iii) when the total depth-wise thickness of the modeled volume is relatively shallow, i.e. $< 100\,\mathrm{m}$. As we state in the text, however, we can safely assume the effect of the lower boundary to be negligible in our simulations since our simulation period is relatively short (40 years) and the depth of the lower boundary condition is very deep ($1\,\mathrm{km}$) so there simply is not enough time for such a small energy flux to propagate all the way to the upper $50\,\mathrm{m}$. Hermoso de Mendoza et al. (2020) found that, even when applying the average continental heat flux (which we use here) at a relatively shallow depth of $42\,\mathrm{m}$, the effect on the ground temperature at $2.86\,\mathrm{m}$ was only $(0.04 \pm 0.01)\,\mathrm{K}$ per $20\,\mathrm{mW\,m^{-2}}$ change in the geothermal heat flux. Given the much deeper depth of the lower boundary in our model, we would expect this effect to be still 1-2 orders of magnitude smaller. Thus, we chose to simply neglect it in the interest of limiting the dimensionality of the parameter space under study.

RC: ***Section 3.4. is empty.***

AR: This was due to a LaTeX typo which we have now fixed in the revised version.

RC: ***L218-220. This is a weird formulation. Any Bayesian approach will include some prior uncertainty on model parameters, and if modelling error is often neglected, it is generally included in the observation error from the likelihood. Maybe what is specific to your approach is that the target X and the predictor Y are actually the same (temperature) and that you have first to estimate the distribution of parameter phi?***

AR: It is common for applied Bayesian methods to neglect certain variables or parameters in the model. A common example is in regression problems, where the posterior distribution $p(\phi|X, Y)$ is typically sought. A fully Bayesian treatment of this problem would assign priors to the predictors, i.e: $p(\phi|X,Y) \propto p(Y|X,\phi)p(X,\phi)$. This is often referred to as a "generative" model since integrating over ϕ yields the joint data distribution $p(Y,X)$. However, $p(X)$ is often neglected which simply corresponds to an implicit uniform prior in the resulting probability model.

Similarly, here we mean to emphasize that the model states, $\mathbf{x}_{1:N}$, on which temperature is functionally dependent, are a deterministic function of the parameters, ϕ, in the probability model described by Eq. 7. A fully Bayesian solution of the problem would account for uncertainty in the model states, $\mathbf{x}_{1:N}$, by treating them also as stochastic quantities. This is precisely the formulation used by Hidden Markov Models (HMM) and their continuous equivalents, state space models (SSM) and stochastic differential equations (SDE). Since our forward model is a deterministic function that does not include such stochasticity, this distribution is implicitly a Dirac density, as we describe in this paragraph.

It is true that model error is implicitly included in the observation error from the likelihood. The limitations of EKS in this regard are also discussed in Sec. 5.6. We will update the text in this paragraph to mention this as well.

RC: ***L224-225. It sounds like a classical McMC approach wouldn't work. Any method sampling the posterior can solve the problem, right?***

AR: MCMC is not feasible here due to the fact that the random walk is sequential, and a very large number of samples (typically thousands) is required in order to produce high quality posterior samples for nontrivial models. More advanced, gradient-based methods such as Hamiltonian Monte Carlo (HMC) are more efficient in this regard but require computation of the gradient which is equally (if not more) costly in a dynamical model such as the one used in this work.

But in principle, any method that can sample the posterior could work, notwithstanding computational limitations. Other particle-based methods such as sequential Monte Carlo (SMC) or even importance sampling would also, in principle, be viable.

RC: ***Computational time for one forward model?***

AR: Each model run takes between 5 and 15 minutes for the full 40 year simulation period on a 5-node compute cluster, each with 24 x Intel® Xeon® Gold 6128 CPUs @ 3.40GHz and 200 GiB of RAM. The compute time for each individual simulation depends on the size of the calculated maximum timesteps which in turn depends on the parameter settings. Heavier tailed freeze curve configurations (e.g. for silty or clay-like soils) tend to be slower due to the strong nonlinearity induced by the freeze curve over a much wider range of subzero temperatures.

RC: ***L253-254. Assuming uncorrelated noise in time and space might be one of the unrealistic assumptions, depending of the type of sensors of course. Maybe mention it in the discussion?***

AR: We will add an additional sentence to the last point of section 5.6.

RC: ***L258-259. It is also a requirement for any Bayesian inference. The posterior is directly related to the prior, so the prior should reflect the actual knowledge about the site. McMC is sampling from the prior distribution as well.***

AR: We will add a sentence here emphasizing this in the revised manuscript.

RC: ***L284-291. I wonder about the validity of deducing trend with such short data sets. In the discussion, comparison with longer trend is introduced, but I feel this could be strengthened.***

AR: We use a Bayesian trend model with semi-informative priors and a robust (Student-t) likelihood for exactly this reason (see Appendix B1 for details). The wider tails of the Student-t likelihood makes the model much more "skeptical" of outliers and thus more conservative in estimating trends. The magnitude of the posterior slopes, despite this robust formulation, suggests that the observed trends are not spurious, although it is of course impossible to extrapolate them into the past without additional data.

As previously suggested by the referee, we will revise the corresponding method section to make it more clear how the trend analysis was performed.

RC: ***What is the reason? Is this biased visible consistenly throughout the year. If yes, all the models of the posterior must have a high misfit, what could indicate a lack of consistency between the prior and the data set (e.g., Lopez-Alvis et al., 2019).***

AR: The reasons for these biases are discussed in Sec. 5.2 of the original manuscript. Please also see our response to the related comment below along with the supplementary results presented in Fig. R1.

RC: ***Close to what? I am not sure the sentence makes sense. Please clarify.***

AR: We assume that the referee is referring to this sentence on line 331: "Both warmer sites also show more overall variability across the ensemble in both latent heat and thaw depth trends with standard deviations close to or sometimes more than double those of the two colder sites."

We meant to say that the standard deviations of the latent heat and thaw depth trend slopes computed across the ensemble are generally much larger than those for the two colder sites, i.e. nearly double or more than double. Fortunately, our updated results with a larger ensemble have made this distinction more clear:

> Both warmer sites also show more overall variability across the ensemble in both latent heat and thaw depth trends with standard deviations ~~close to or sometimes~~ of the trend slopes more than double those of the ~~two colder~~ cold sites.

RC: ***L374-375. Is this significant enough to state that the difference is not only due to the depth of the sensor? If both had some sensors deeper, this would largely reduce the uncertainty and has likely nothing to do with the fact that they are warmer, don't you think?***

AR: Following the suggestion of the other referee, we repeated the Samoylov simulations with all sensors below $10\,\mathrm{m}$ excluded from the inference procedure. The results showed that this does indeed increase the spread of the predicted temperature, as we might expect, but does not affect the uncertainty related to changes in latent heat and active layer thickness. We have revised the text and moved it to section 5.3:

The wider spread in the ensemble thaw depth and latent heat change at both Bayelva and Parson's Lake (Fig. 5 and Fig. 6) indicates that there is more uncertainty in modeling the thermal dynamics of warmer permafrost. This is likely due to higher sensitivity to soil properties, initial conditions, and changes in surface conditions and soil water content. We suspect that this sensitivity is in large part attributable to the nonlinear effects of the freeze curve on the thermal dynamics, especially in deep permafrost (i.e. near the depth of zero annual amplitude) where thermal gradients are smaller and heat diffuses more slowly.

RC: ***Why a reference to the median suddenly? The discrepancy is for all models, not only the median or the mean, all predictions are wrong.***

AR: We assume that the referee is referring to this sentence: "For Samoylov, the ensemble median of the modeled annual temperature range of the permafrost layers is slightly too narrow..."

This is actually a mistake; we have now corrected this line in the revised text:

For Samoylov, the ensemble ~~median of the modeled~~ annual temperature range of the permafrost layers is slightly too narrow...

In general, we often refer to the median elsewhere in the text because, as a quantile, it is robust against distributional skew which does occur in some of the ensemble predictive distributions (most dramatically for Bayelva).

RC: ***L386-387. Have you tried enlarging the prior? Do you observe a similar bias for other years?***

AR: We generally tried to select priors for the n-factors and soil properties that were informative enough to be useful but not so informative where they would be restrictive. Ideally we would perform a robust sensitivity analysis with a wide range of priors, but given the complexity of the downstream modeling task, this is prohibitively difficult. While we did perform a number of ad-hoc experiments with wider or narrower priors, this did not affect this specific issue for Samoylov. The bias is also persistent across years (Fig. R1) with the exception of 2016 in which the observed temperatures are unusually warm due to early season snowfall.

It is very likely that the low temporal-resolution (bidecadal) of the n-factor parameterization is the source of this bias since it cannot capture interseasonal (or even interannual) variation in snow cover effects. It is only a very coarse approximation of snow cover effects to allow the model to capture long term trends.

RC: ***What do you mean by "smooth regularizers". One of the advantage of Bayesian inversion is to use realistic prior and avoid regularization (such as smoothing). Smoothing is normally only visible in the mean which is not necessarily a sample of the posterior.***

AR: The intent here was to relate the role of the prior in Bayesian inference to the role of regularization in classical optimization problems (Calvetti and Somersalo, 2018). It is a well known result that there is a correspondence between certain regularization

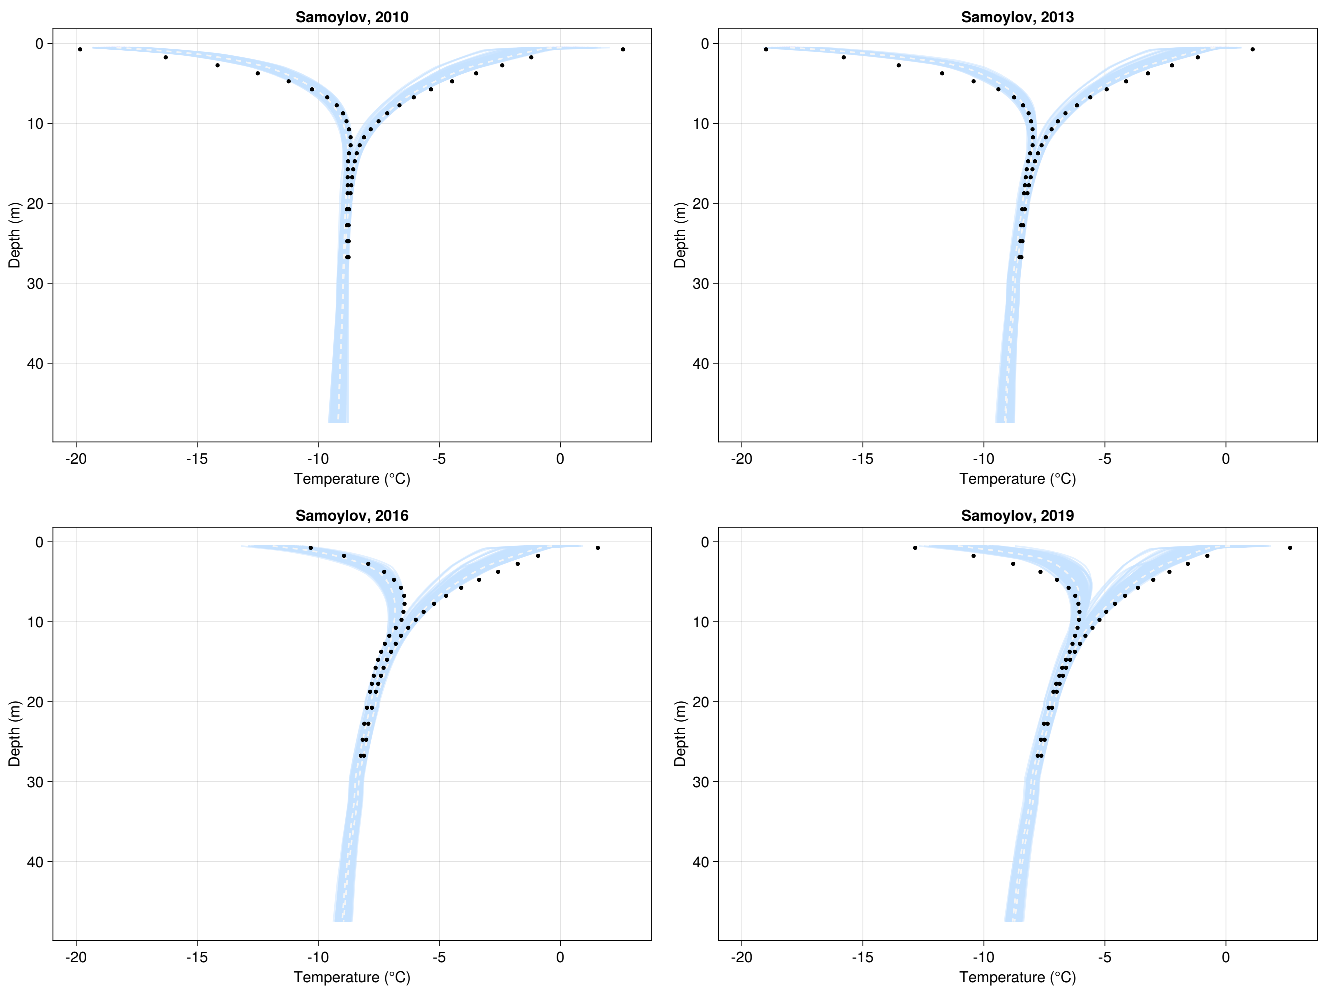

Figure R1. Samoylov ensemble trumpet curves (annual min vs max temperature) for four selected years with data available: 2010, 2013, 2016, 2019. The black dots are observed min/max temperatures for that year. The dashed white line shows the median annual temperature.

schemes and families of prior distributions, e.g. ℓ_2-regularization in ridge regression corresponds to a multivariate Gaussian prior, and ℓ_1-regularization in LASSO corresponds to a multivariate Laplace prior (Hastie et al., 2009).

We have removed the word "smooth" from this sentence to avoid this confusion. We will also add the aforementioned references.

RC: ***I think the term "appropriate" is badly chosen. The prior should reflect the uncertainty on model parameters before considering the data set, and what is mentioned should then result in larger prior. The main challenge is maybe to consider correlation between parameters?***

AR: We have improved the wording of this sentence and added a remark about correlation:

> However, the general lack of precise error and uncertainty estimates for some model parameters (in particular, soil properties) makes the ~~selection of appropriate~~ construction of prior distributions difficult. Furthermore, it is very likely that many parameters are, in reality, highly correlated, which we are unable to account for with our current approach.

RC: ***L515-516. This is also the case for most McMC approaches.***

AR: To clarify, we meant that the noise covariance must be fixed a priori, i.e. it is not treated as a parameter. Since MCMC is a more flexible method that can work on any target density, the noise parameters are typically included in the posterior distribution (Gelman et al., 1995) and sampled to fit the data. This is not possible with EKS. We will clarify this in the revised manuscript.

RC: ***L520-523. See comment 1.***

AR: As stated above, we will include additional discussion about alternative methods in the reivsed manuscript.

RC: ***Table B5. Display prior distributions behind posterior in figures A2 to A5 to immediately grasp the reduction in parameter uncertainty?***

AR: This is a great suggestion and we have implemented it in the revised manuscript. The new plots are also shown below in Fig. B5-B8.

A. Additional comments from the authors

In the process of revisiting our simulation code to follow-up on the referee's concerns, we discovered some unrelated bugs in the configuration of the prior distributions for some of the parameters (specifically, the freeze curve parameters for the Parson's Lake site and the saturation level parameter for all sites). We fixed these issues and re-ran the simulations for all sites with the same random seed as used in the original simulations. We also increased the size of the ensemble from 256 to 512 to improve the robustness of our results as well as to provide a larger sample size for the sensitivity analysis requested by the second referee.

B. Sensitivity and density plots

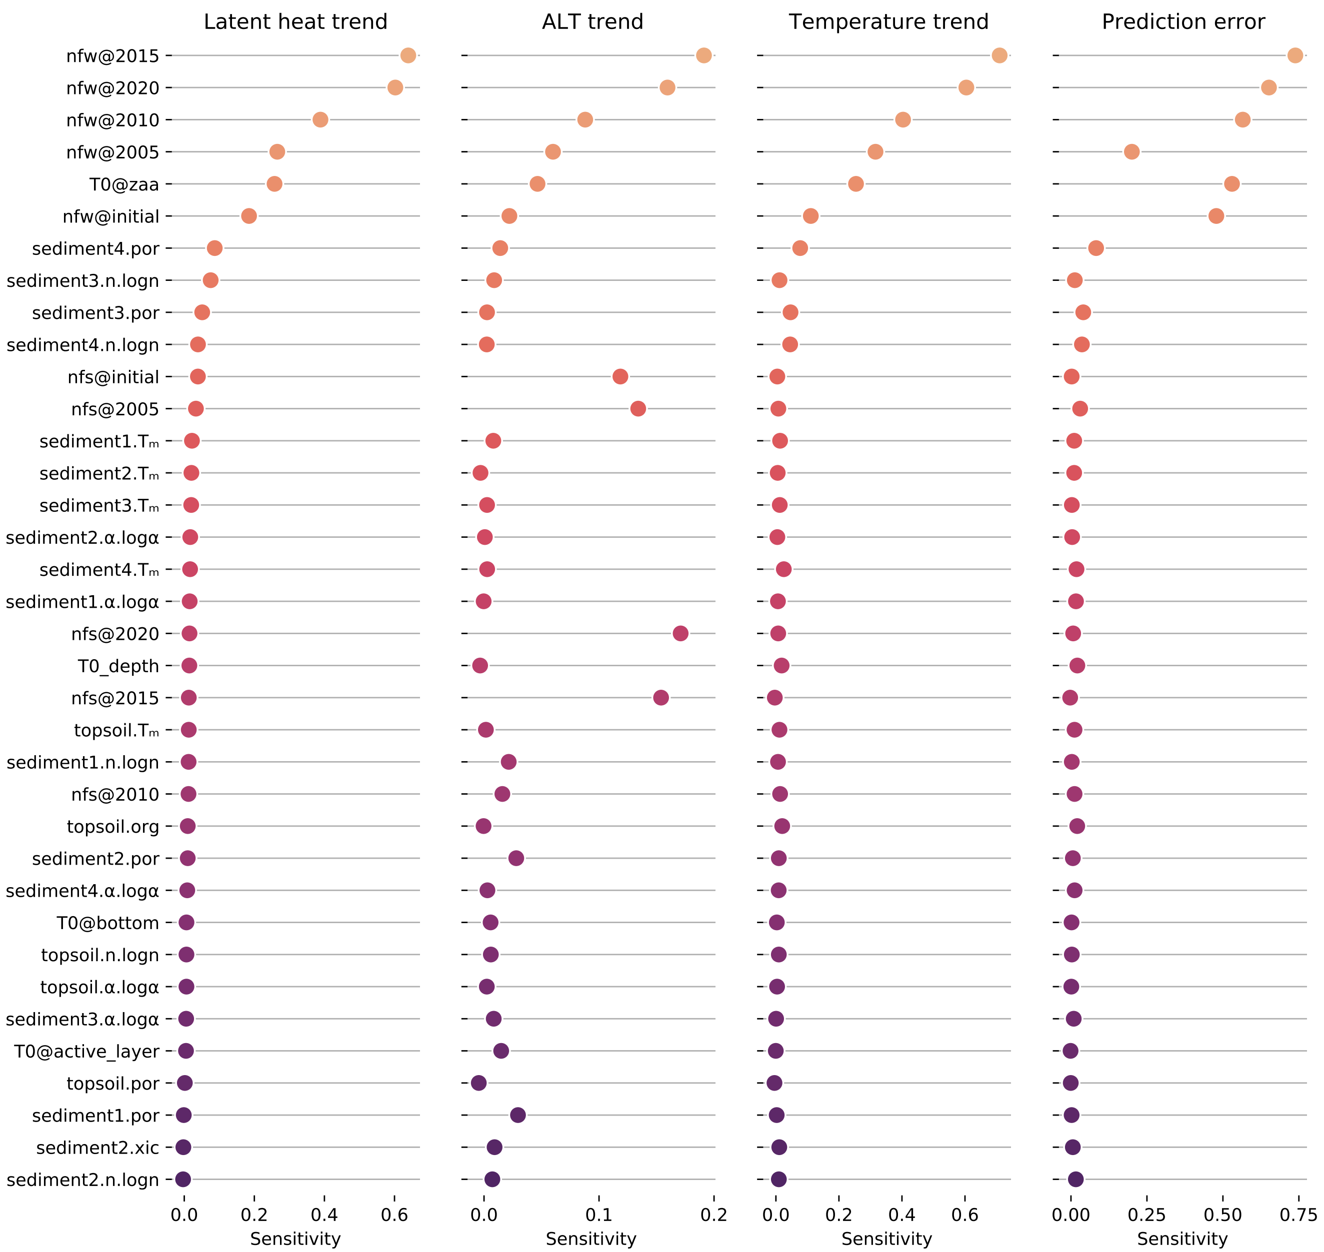

Figure B1. Samoylov Island parameter sensitivities (first order Sobol indices) w.r.t to the latent heat trend, active layer thickness (ALT) trend, temperature trend, and ground temperature prediction error estimated with the EASI algorithm

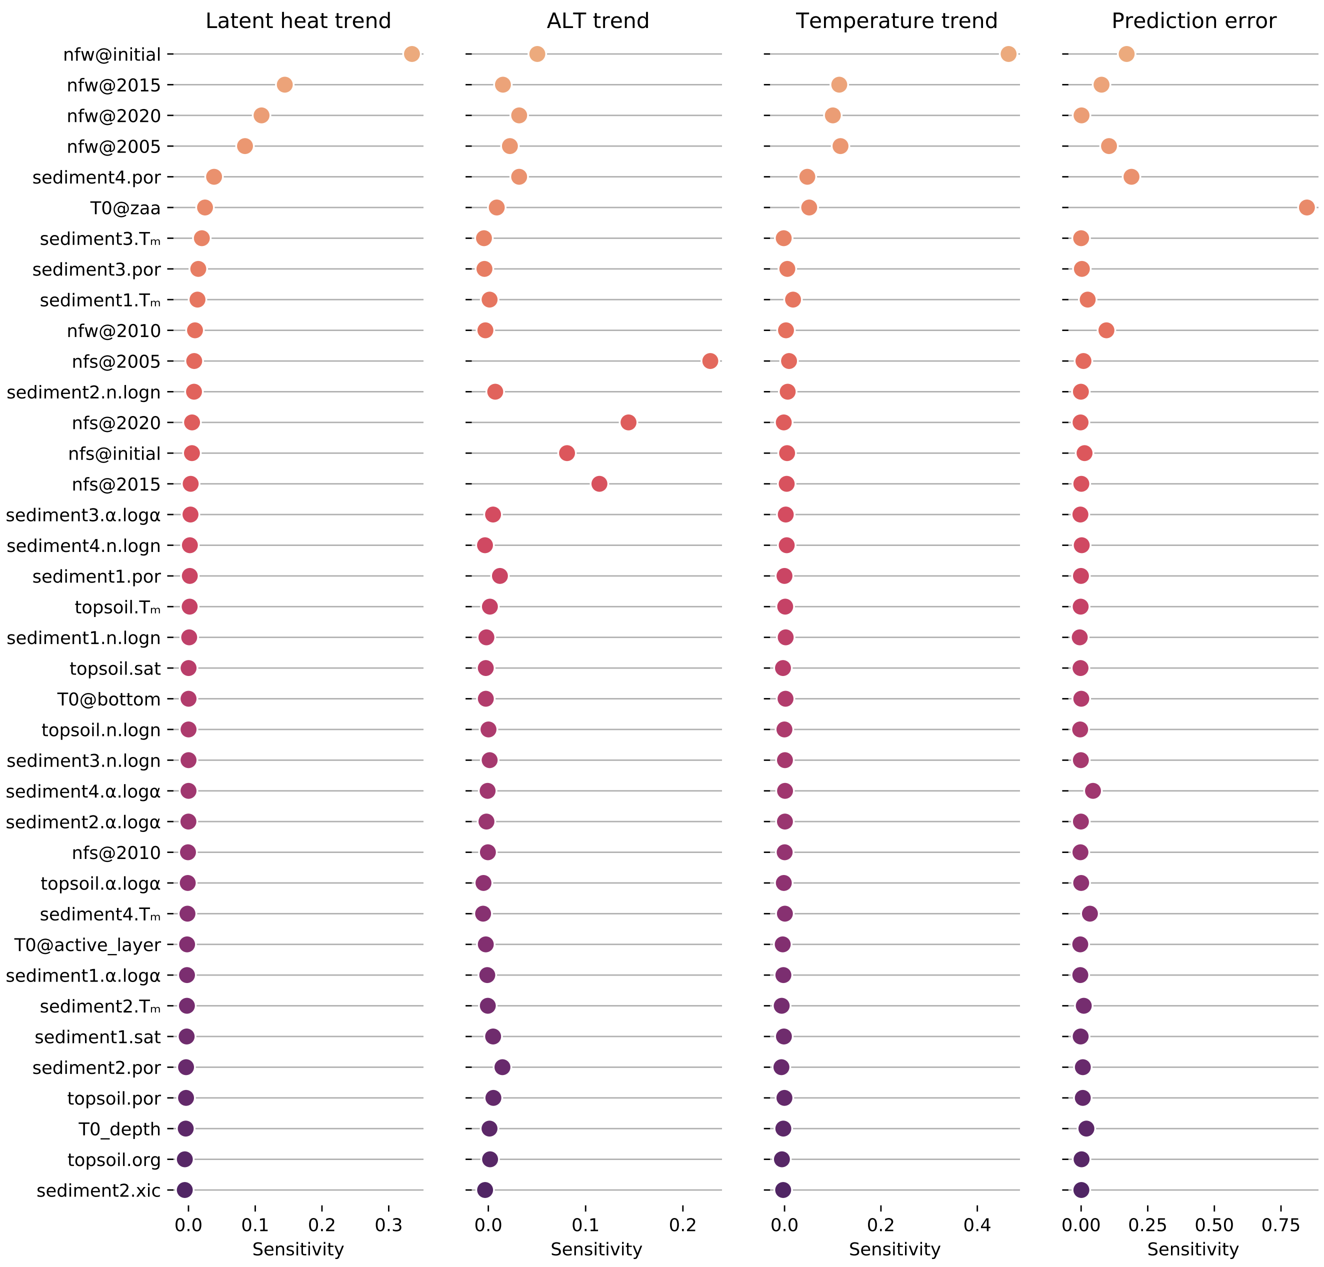

Figure B2. Barrow parameter sensitivities (first order Sobol indices) w.r.t to the latent heat trend, active layer thickness (ALT) trend, temperature trend, and ground temperature prediction error estimated with the EASI algorithm

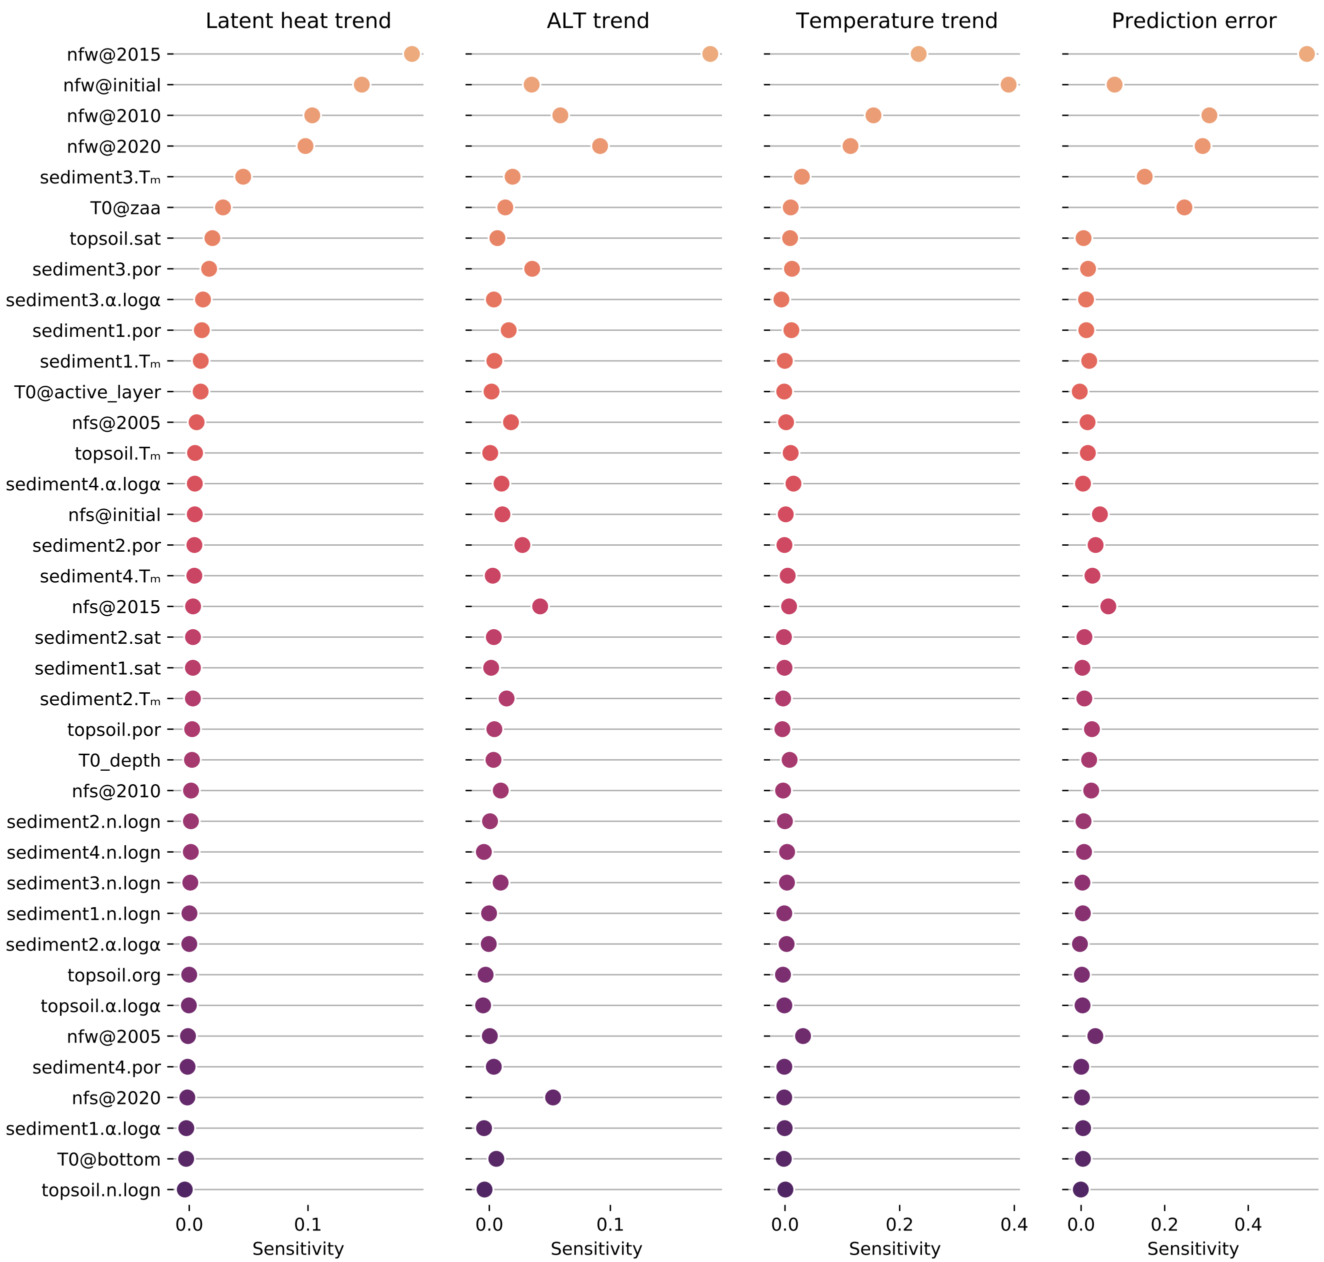

Figure B3. Bayelva parameter sensitivities (first order Sobol indices) w.r.t to the latent heat trend, active layer thickness (ALT) trend, temperature trend, and ground temperature prediction error estimated with the EASI algorithm

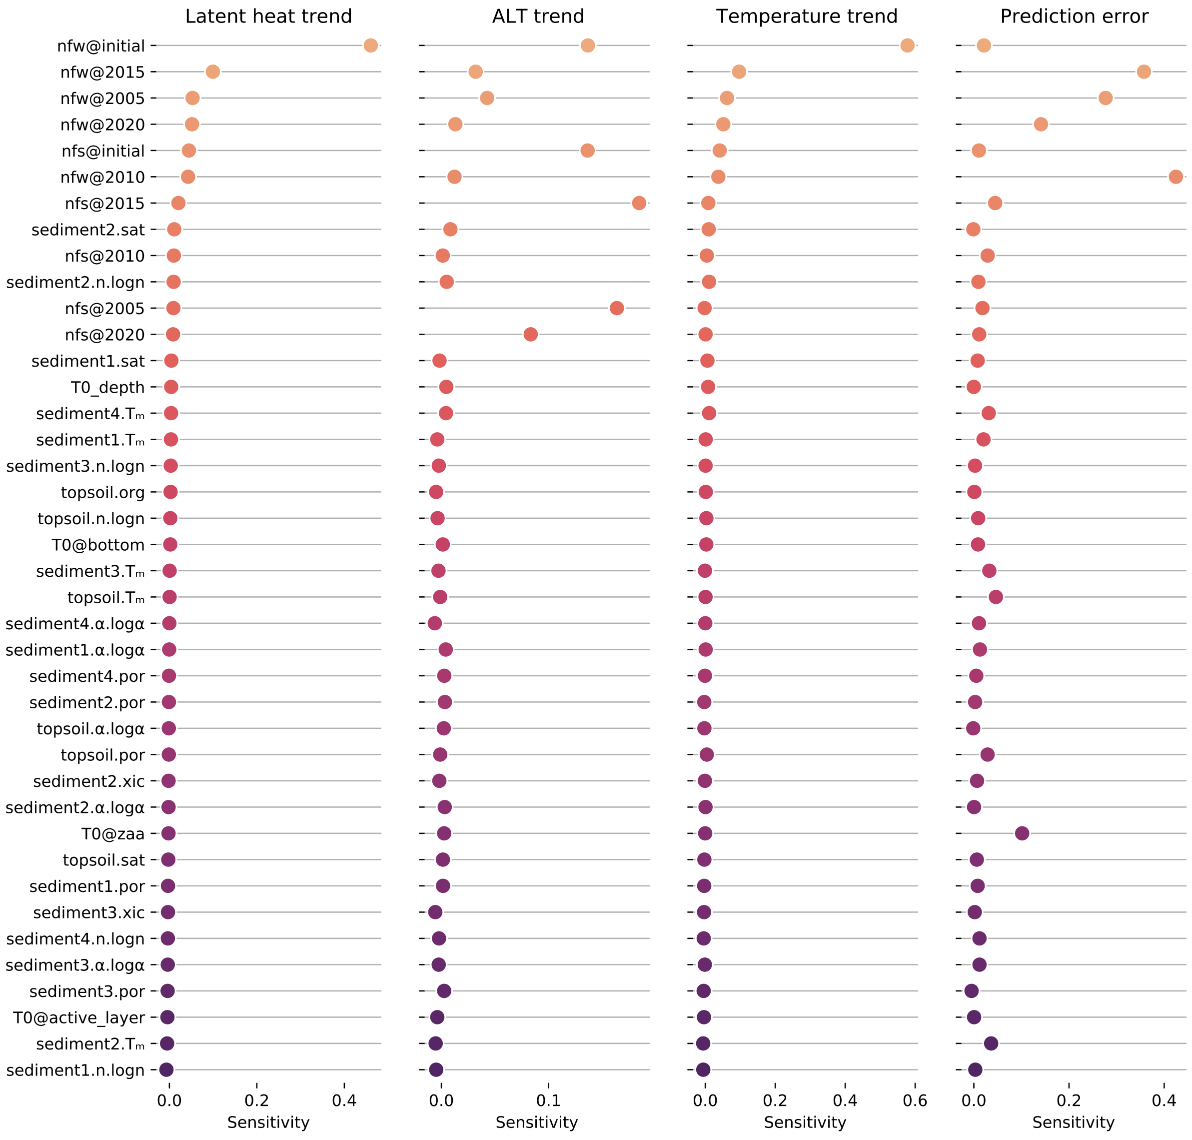

Figure B4. Parson's Lake parameter sensitivities (first order Sobol indices) w.r.t to the latent heat trend, active layer thickness (ALT) trend, temperature trend, and ground temperature prediction error estimated with the EASI algorithm

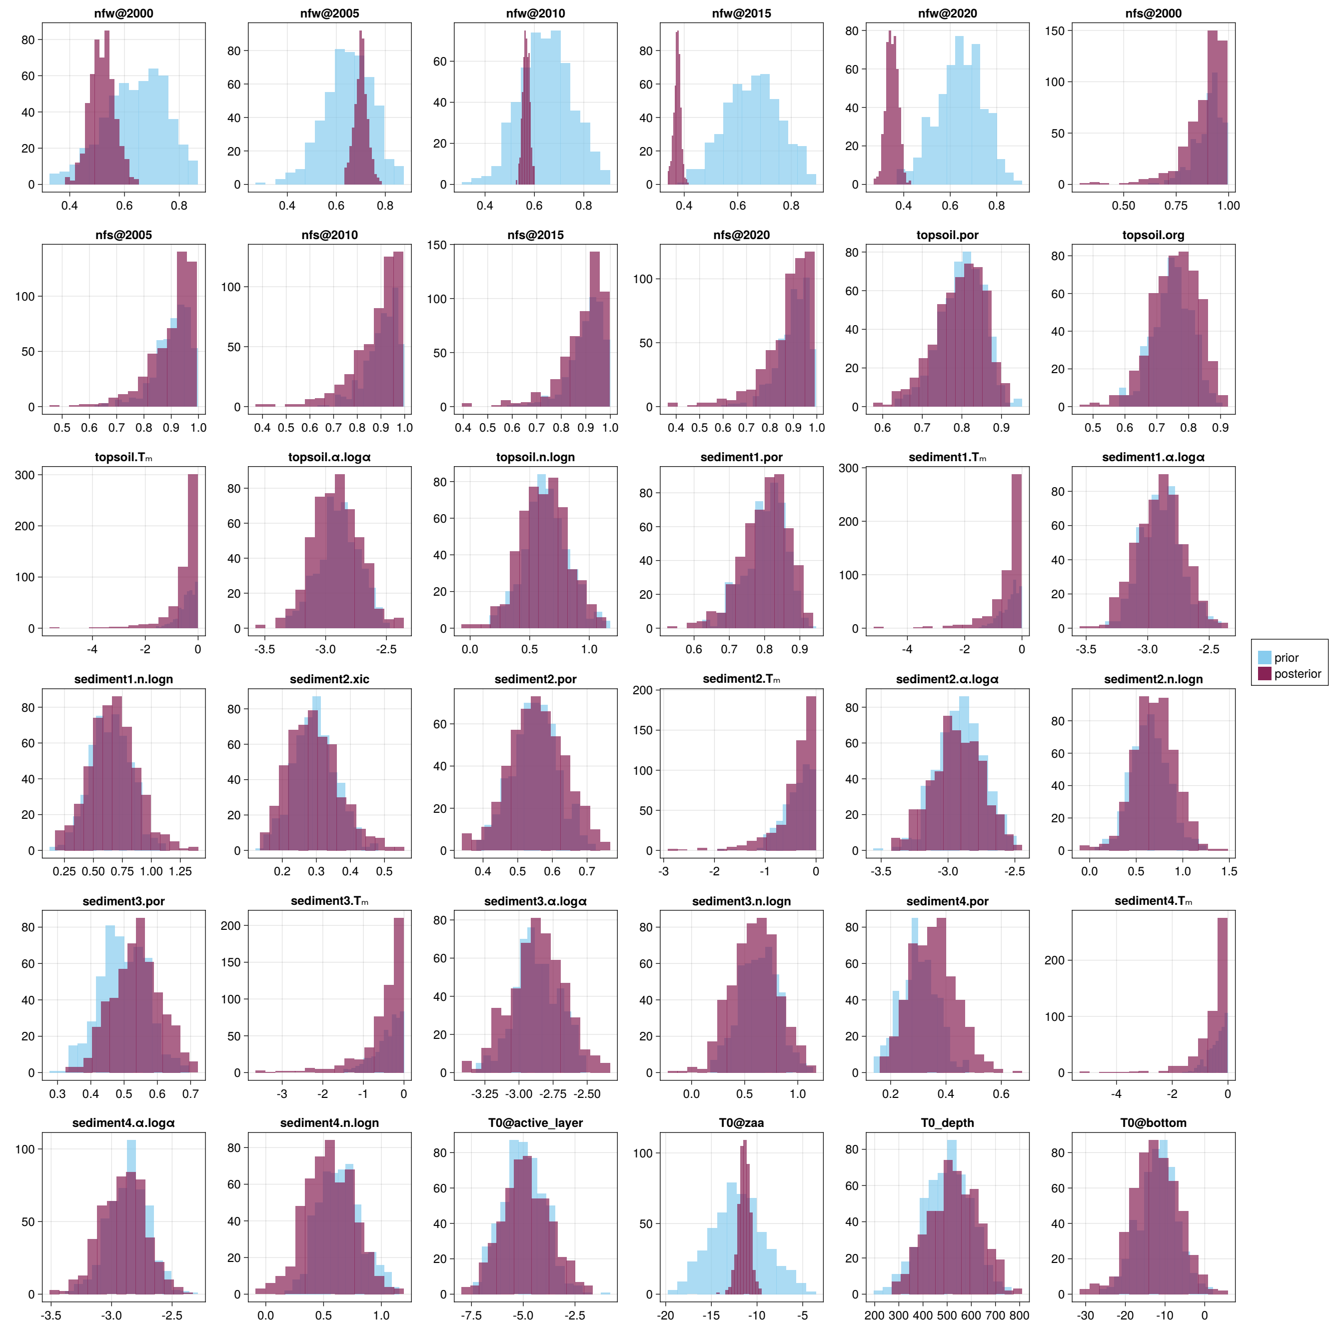

Figure B5. Histograms of prior vs posterior samples for the Samoylov Island site

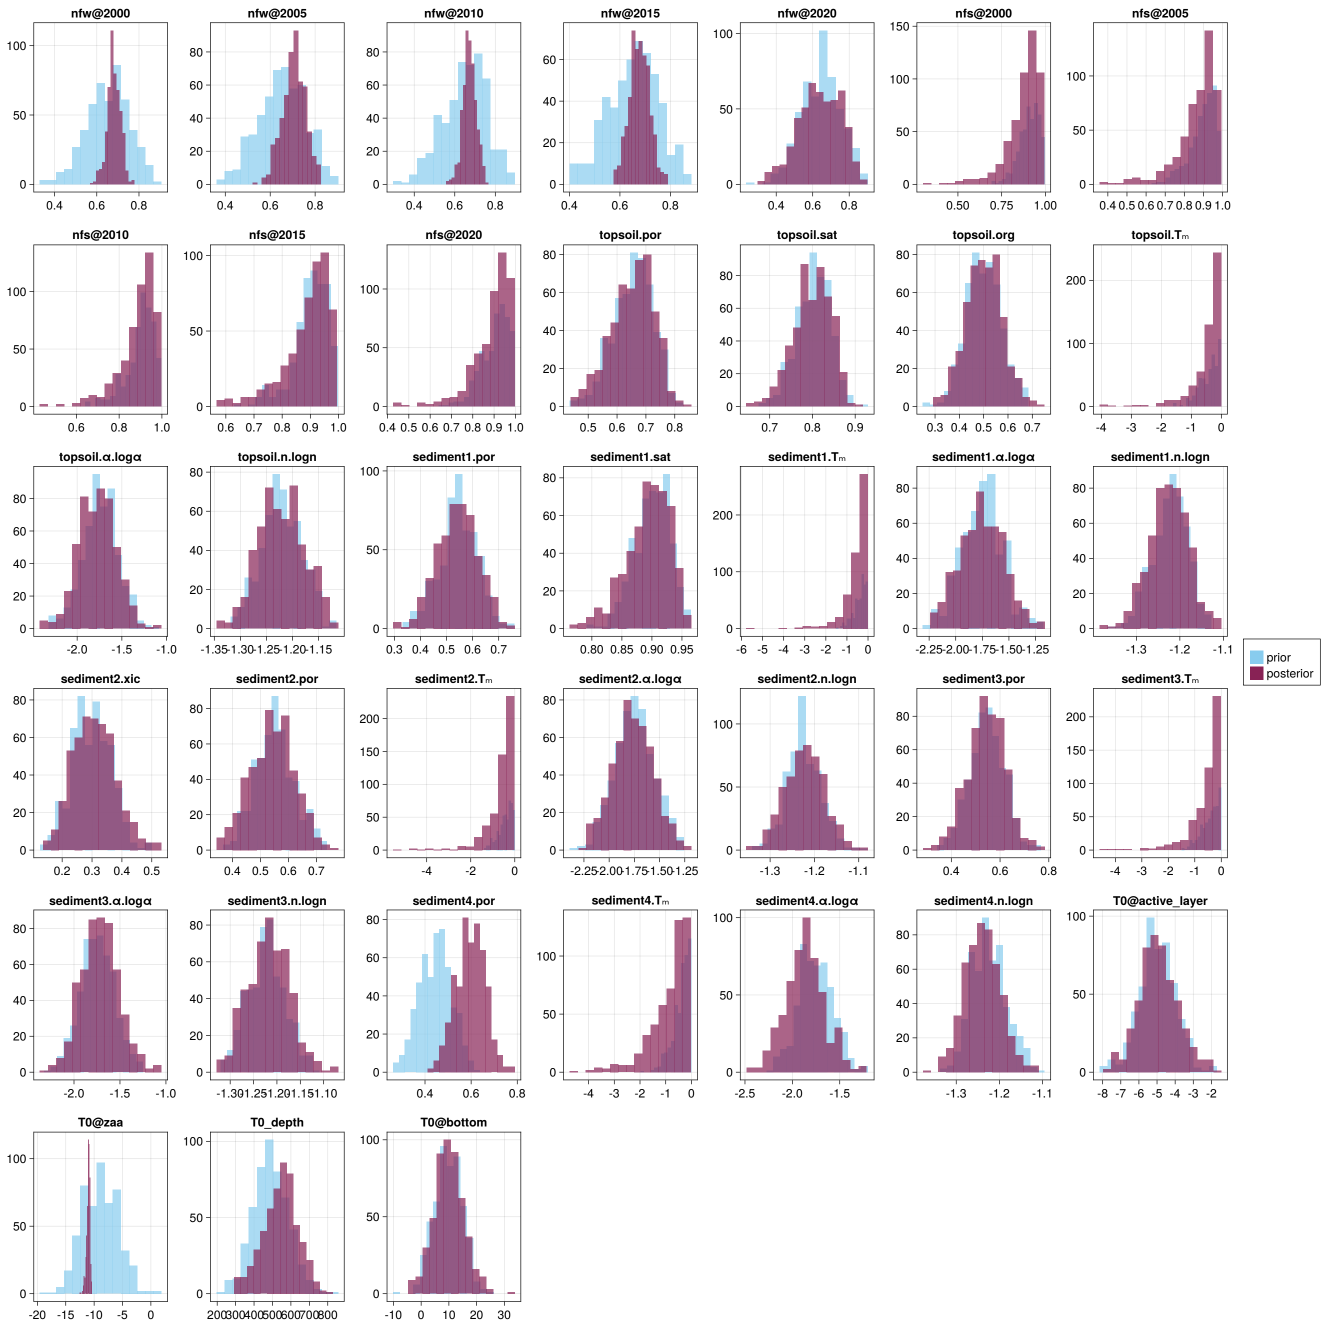

Figure B6. Histograms of prior vs posterior samples for the Barrow site

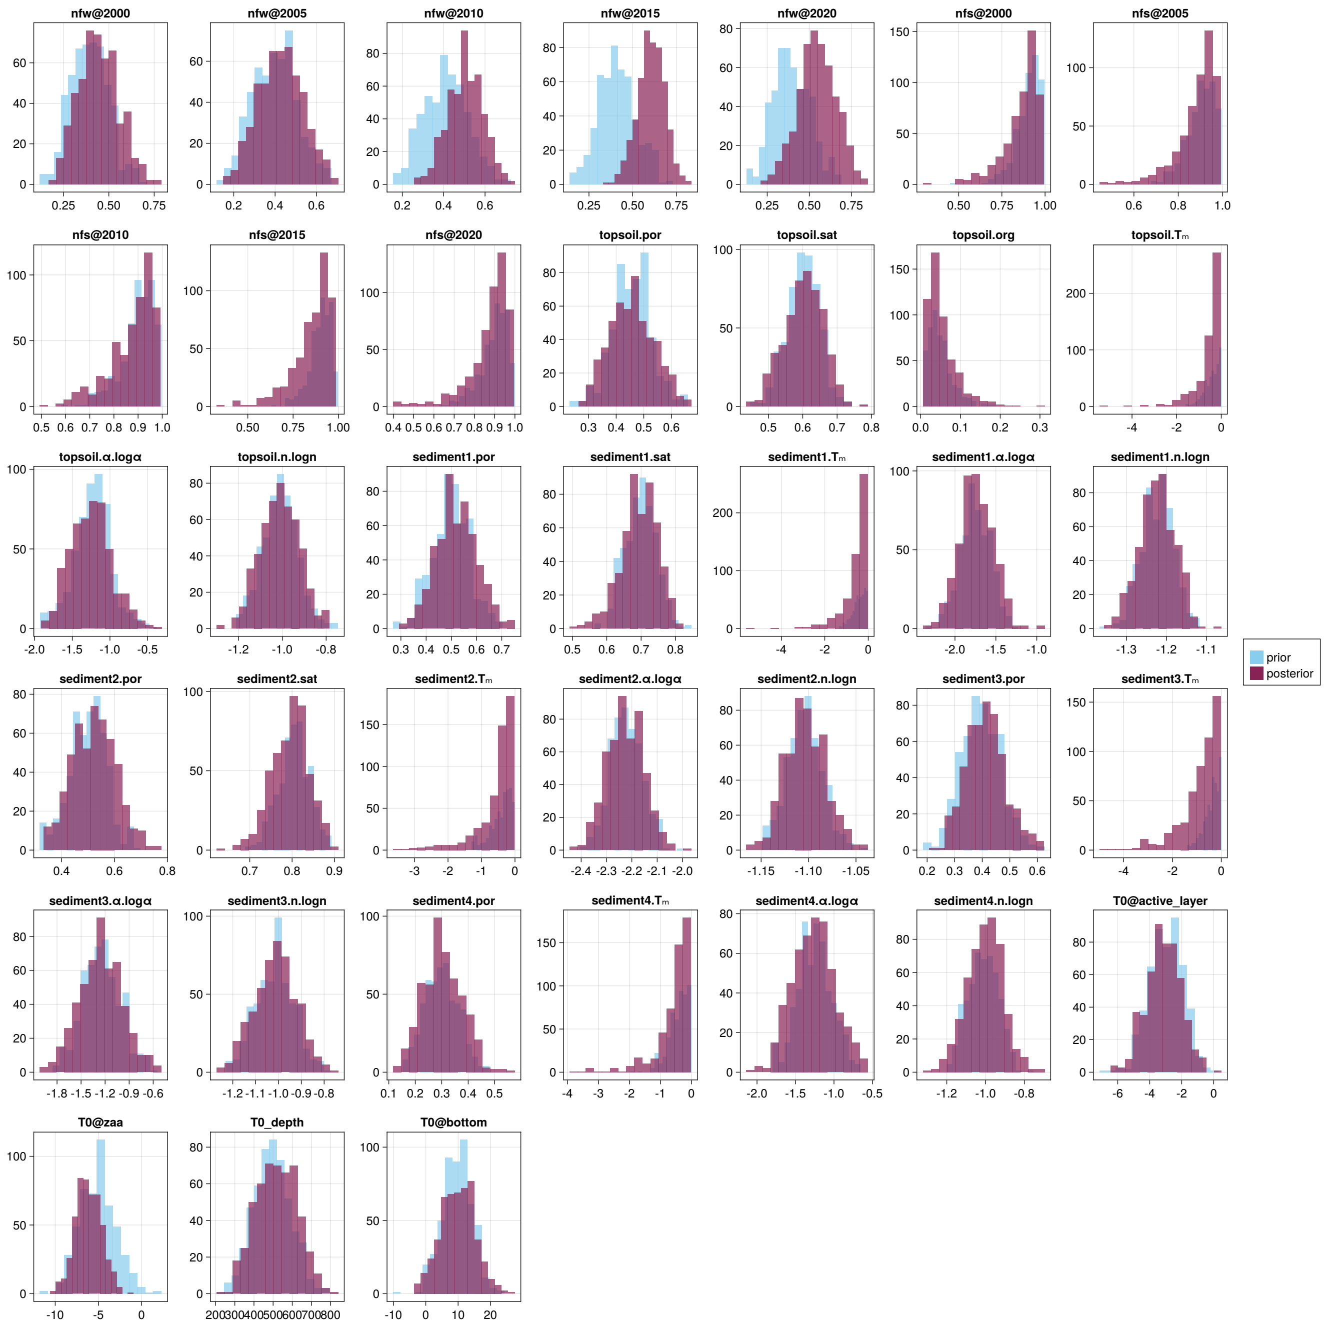

Figure B7. Histograms of prior vs posterior samples for the Bayelva site

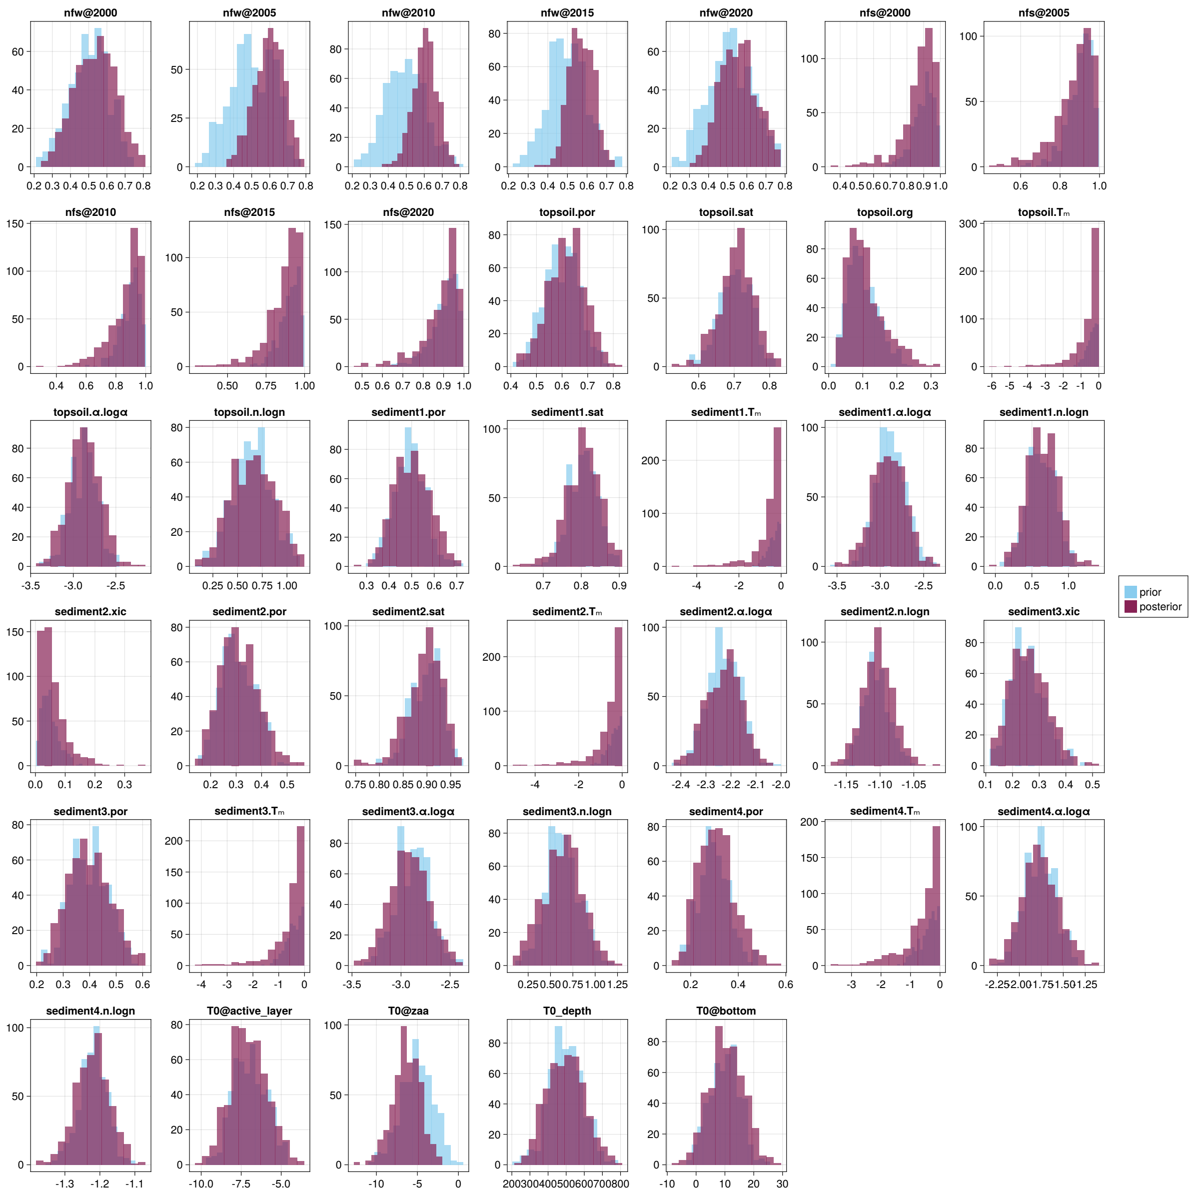

Figure B8. Histograms of prior vs posterior samples for the Parson’s Lake site

References

Calvetti, D. and Somersalo, E.: Inverse Problems: From Regularization to Bayesian Inference, WIREs Computational Statistics, 10, https://doi.org/10.1002/wics.1427, 2018.

Dixit, V. K. and Rackauckas, C.: GlobalSensitivity. Jl: Performant and Parallel Global Sensitivity Analysis with Julia, Journal of Open Source Software, 7, 4561, 2022.

Doucet, A., de Freitas, N., and Gordon, N.: An Introduction to Sequential Monte Carlo Methods, in: Sequential Monte Carlo Methods in Practice, edited by Doucet, A., de Freitas, N., and Gordon, N., Statistics for Engineering and Information Science, pp. 3–14, Springer, New York, NY, https://doi.org/10.1007/978-1-4757-3437-9_1, 2001.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B.: Bayesian Data Analysis, Chapman and Hall/CRC, 1995.

Geman, S. and Geman, D.: Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images, IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-6, 721–741, https://doi.org/10.1109/TPAMI.1984.4767596, 1984.

Harp, D. R., Atchley, A. L., Painter, S. L., Coon, E. T., Wilson, C. J., Romanovsky, V. E., and Rowland, J. C.: Effect of Soil Property Uncertainties on Permafrost Thaw Projections: A Calibration-Constrained Analysis, The Cryosphere, 10, 341–358, https://doi.org/10.5194/tc-10-341-2016, 2016.

Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H.: The Elements of Statistical Learning: Data Mining, Inference, and Prediction, vol. 2, Springer, 2009.

Hermans, T., Nguyen, F., Klepikova, M., Dassargues, A., and Caers, J.: Uncertainty Quantification of Medium-Term Heat Storage From Short-Term Geophysical Experiments Using Bayesian Evidential Learning, Water Resources Research, 54, 2931–2948, https://doi.org/10.1002/2017WR022135, 2018.

Hermoso de Mendoza, I., Beltrami, H., MacDougall, A. H., and Mareschal, J.-C.: Lower Boundary Conditions in Land Surface Models – Effects on the Permafrost and the Carbon Pools: A Case Study with CLM4.5, Geoscientific Model Development, 13, 1663–1683, https://doi.org/10.5194/gmd-13-1663-2020, 2020.

Kantas, N., Beskos, A., and Jasra, A.: Sequential Monte Carlo Methods for High-Dimensional Inverse Problems: A Case Study for the Navier–Stokes Equations, SIAM/ASA Journal on Uncertainty Quantification, 2, 464–489, https://doi.org/10.1137/130930364, 2014.

Moradkhani, H., Hsu, K.-L., Gupta, H., and Sorooshian, S.: Uncertainty Assessment of Hydrologic Model States and Parameters: Sequential Data Assimilation Using the Particle Filter, Water Resources Research, 41, https://doi.org/10.1029/2004WR003604, 2005.

Neal, R. M.: MCMC Using Hamiltonian Dynamics, in: Handbook of Markov Chain Monte Carlo, pp. 139–188, Chapman and Hall/CRC, 2011.

Noh, S. J., Tachikawa, Y., Shiiba, M., and Kim, S.: Applying Sequential Monte Carlo Methods into a Distributed Hydrologic Model: Lagged Particle Filtering Approach with Regularization, Hydrology and Earth System Sciences, 15, 3237–3251, https://doi.org/10.5194/hess-15-3237-2011, 2011.

Plischke, E.: An Effective Algorithm for Computing Global Sensitivity Indices (EASI), Reliability Engineering & System Safety, 95, 354–360, https://doi.org/10.1016/j.ress.2009.11.005, 2010.

Sisson, S. A., Fan, Y., and Tanaka, M. M.: Sequential Monte Carlo without Likelihoods, Proceedings of the National Academy of Sciences, 104, 1760–1765, https://doi.org/10.1073/pnas.0607208104, 2007.

Thibaut, R., Compaire, N., Lesparre, N., Ramgraber, M., Laloy, E., and Hermans, T.: Comparing Well and Geophysical Data for Temperature Monitoring within a Bayesian Experimental Design Framework, Water Resources Research, 58, e2022WR033 045, 2022.

Wang, J. and Zabaras, N.: Hierarchical Bayesian Models for Inverse Problems in Heat Conduction, Inverse Problems, 21, 183–206, https://doi.org/10.1088/0266-5611/21/1/012, 2005.