Second response to reviews for *"Parameter estimation for land-surface models using Neural Physics"*, Huang et al. (egusphere-2025-6015)

We would like to thank the reviewers for taking the time to carefully review our revised paper. We are happy that two of the reviewers have no further comments and will now address third reviewer's next set of comments. We have further revised our manuscript as a result of these comments and respond below to each of the points raised. We supply a marked-up copy of our revisions which include some improvements to the writing as well as changes made in response to the reviewer's comments (see especially lines 20–25, 80–132 and 344–371).

1. I fully agree with the authors that differentiable land models are an exciting research subject and that they could hopefully develop into a new generation of land models, however I still have some doubts about the concrete work presented here.

It's also really good to see that the authors fundamentally revised the introduction and cite related work. However, with a model and method as simple as the authors present here, it's unfortunately not clear to me what the novelty is here compared to existing differentiable land models such as JAX-CanVeg, Terrarium.jl or SINDBAD or also hybrid models such DifferLand and NeuralCrop that far exceed the model complexity presented here.

The central contribution of this work is to investigate parameter estimation, identifiability and equifinality in a fully differentiable, physics-based land-surface model implemented in a convolutional framework (i.e. Neural Physics). We are not claiming that Neural Physics is the only route to differentiable LSMs. Instead, we claim that it provides a mature framework in which numerical operators are represented using convolutions and which already supports differentiable PDE solvers, GPU execution and multigrid methods. Our contribution is to demonstrate parameter estimation for an LSM within this framework. This differs from existing differential land-surface models such as Jax-CanVeg, Terrarium.jl, DifferLand and NeuralCrop which are based on conventional numerical formulations and automatic differentiation, with support for hybrid physics-based/data-driven representations. SINDBAD does not use a gradient-based approach for optimisation, instead, it runs the model repeatedly with different parameter sets and evaluates the objective function using derivative-free calibration strategies. By contrast, our approach employs the Neural Physics framework, in which the forward model remains entirely physics-based and training-free. Our goal is therefore not to compete on model complexity, but to evaluate the suitability of the Neural Physics paradigm for parameter estimation of a simple land-surface model, and to investigate identifiability, confounding and equifinality.

2. The model of the authors is effectively just a 1D heat equation without any further parameterisations. For example they also prescribe a set Bowen ratio / latent and sensible heat flux. This is a really simple model. Parameter estimation via gradient descent of 1D PDEs like this has been demonstrated countless times across many domains.

We intentionally chose the simplest model that still includes a surface energy balance and a time-dependent PDE, in order to isolate parameter identifiability, optimisation behaviour and equifinality. Whilst gradient-based parameter estimation for simple PDEs such as the 1D heat equation has been widely studied, our objective is not to contribute a new optimisation method for this class of problems. Instead, we use this simple LSM as a minimal but physically meaningful benchmark to analyse how parameter estimation depends on model structure and observational constraints. This provides a suitable setting in which to study identifiability with a differentiable land-surface modelling framework.

3. In their reply on RC2 the authors write about that "but the demonstration that such models can be calibrated efficiently without adjoint derivation, and that the framework provides direct insight into when parameter estimation is well-posed and how this depends on the observational configuration." in their reply to RC1 they similarly write "The key advance is not the use of gradient-based optimisation per se, but that the full parameter estimation problem for a time-dependent LSM can be solved end-to-end in a differentiable framework without adjoint derivation or bespoke optimisation infrastructure. "

However, I would argue that the model the authors present here is so simple that no real conclusion can be made for full land surface models that feature many more processes and parameterisations.

We agree that the model presented in this study is simple compared to full land-surface models with many interacting processes and parameterisations. Our aim is to demonstrate the use of a differentiable, physics-based LSM, implemented within a convolutional framework, for parameter estimation, and to provide a foundation for future extensions to more complex land-surface models. Instead, the purpose of this work is to analyse parameter estimation, identifiability and equifinality, resulting from applying a fully differentiable framework to a simplified land-surface model. The conclusions presented in the paper are therefore restricted to this model and the observational configurations considered. We have revised the manuscript to ensure that all the conclusions are based on the model and results presented in this study.

4. I see two main challenges for differentiable land models:
1) The software engineering and effort needed to (re-)write such complex models in AD-compatible frameworks
2) Discontinuities in their mathematical formulation that make the model mathematically not differentiable. Land models feature a lot of discontinuities in their parameterisation, e.g. with look up tables. While this kind of control flow is not an issue for most AD libraries per se, it does still make gradient-based optimisations challenging because the gradients might just be mathematically wrong. As a remedy e.g. smoothing strategies have been proposed for example by Kaminski et 2013 for the BETHY/JSBACH model that features a handwritten adjoint. Other strategies might include learning these relationships with NNs.

From my point of view the model of the authors addresses neither of these challenges.

These two points are not the focus of our paper.

On the first point, we agree that the software engineering and effort associated with translating complex land-surface models into AD-compatible frameworks is an important practical challenge. However, addressing code modernisation strategies is beyond the scope of the present study. We briefly acknowledged this issue in the penultimate paragraph of the conclusion and noted that recent developments in automated code translation tools may substantially reduce the associated effort (Zhou et al., 2024).

On the second point, we agree that many operational land-surface models include discontinuities arising from lookup tables, threshold-based parameterisations and conditional logic, which can complicate gradient-based optimisation in practice. However, the model considered in this study does not include such discontinuities, as it is based on smooth, explicitly discretised governing equations. The purpose of this work is therefore not to address non-smooth parameterisations, but to isolate and study parameter estimation in a fully differentiable setting. Extending our approach to discontinuous or hybrid formulations is an important direction for future work. More generally, whilst modern automatic differentiation frameworks can in some cases handle piecewise-smooth operations, the treatment of strongly discontinuous functions remains an active area of research. We have included this point in the future challenges discussed in the conclusion.

5. About Neural Physics: Why do the authors use the approach they refer to as "Neural Physics"? With this method they compute the finite difference operator as a convolution. The authors stress this approach several times throughout the paper to be important to make their model differentiable and name it to be the "core" of their paper in the reply on EC1. However, finite difference operators are not hard to differentiate through without it as well. When you write out the finite difference operator in sparse or banded matrix form, the differentiability is also easily achieved with basically every AD library as the operation can be simply computed as a matrix multiplication. What makes the Neural Physics special? It might be that the convolution of PyTorch is better optimised and computationally more efficient than a banded or sparse matrix multiplication, but the authors didn’t investigate this.

We are not claiming that Neural Physics is the only route to differentiable LSMs. Instead, we claim that it provides a mature framework in which numerical operators are represented using machine-learning primitives (in this case, convolutions) and which already supports differentiable PDE solvers, GPU execution and multigrid methods. We agree also that differentiability does not depend on the use of convolutions. The present problem could equally be implemented using sparse matrix operators and differentiated with modern automatic differentiation tools. We have revised the manuscript to clarify that the motivation for using Neural Physics is not purely differentiability, but rather the use of an existing framework in which numerical discretisations are expressed through convolutional kernels. More generally, the Neural Physics framework can support a broader class of numerical methods, including higher-order discretisations, multigrid solvers and GPU execution, which we hope to exploit for a range of inverse problems in the future.

There may be advantages in computational efficiency of the Neural Physics approach over conventional numerical approaches. For example, in contrast to explicit sparse matrix assembly and multiplication, convolution-based implementations avoid explicit index management and can also exploit mixed precision arithmetic. However, we did not compare these methods in this study, as our focus is on parameter estimation within a fully differentiable framework arising from the use of convolutional operators.

6. To summarise: I share the author's excitement about the possibilities of differentiable models. However, unfortunately I don't really see much novelty in this paper. As ways forward I would suggest to either: significantly increase the complexity of the model they investigate, e.g. also including parameterisations with discontinuities, or to focus much more prominently on their application on the dataset they investigated and make this the main focus of the paper. Another idea would be to include gradient-free optimisations techniques as benchmarks into their study to show that the gradient-based optimisation really works better.

We thank the reviewer for this summary and for the constructive suggestions regarding possible extensions of the work. We agree that the present study does not aim to increase model complexity, introduce discontinuous parameterisations or provide an exhaustive benchmarking study against gradient-free optimisation methods. Instead, the central contribution of this work is more narrowly defined: we investigate the behaviour of parameter estimation, identifiability and equifinality in a fully differentiable, physics-based land-surface model realised through a convolution-based framework.

In this sense, the value of the study lies not in model complexity, but in isolating and analysing the structure of the parameter estimation under controlled but physically meaningful dynamics. This allows us to examine how observational configurations affect identifiability and optimisation behaviour in a differentiable setting. We agree that extending the model to include additional process complexity, discontinuous parameterisations or benchmarking against gradient-free optimisation methods would be valuable directions for future work. However, these extensions would constitute separate studies, as they introduce additional confounding factors that would obscure the specific identifiability and equifinality questions addressed here.

We have revised the manuscript to ensure that this positioning is made explicit and that the claims are restricted to the scope of the model and experiments presented.

Response to Editor:

Dear Editor,

We would like to thank you and the reviewers for the careful and constructive assessment of our manuscript. We have substantially revised the paper in response to the comments received during the review process.

In particular, we expanded the literature review on differentiable and hybrid land-surface models, clarified the distinction between fully physics-based differentiable formulations and hybrid machine-learning approaches, strengthened the discussion of parameter identifiability and equifinality, added out-of-sample evaluation for the Phoenix dataset, revised the discussion of Neural Physics to avoid overstating the role of convolutional formulations, and expanded the discussion of limitations and future directions.

We recognise that Reviewer 2 remains unconvinced regarding the level of novelty and the intentionally simplified nature of the model. We respectfully believe that this reflects a difference in perspective regarding the intended scope of the paper rather than a fundamental flaw in the methodology or analysis. The purpose of the manuscript is not to present the most comprehensive differentiable land-surface model, but rather to use a deliberately controlled, fully differentiable, physics-based setting to investigate parameter estimation, identifiability, equifinality and observational design.

We believe the revised manuscript now positions this contribution much more clearly and avoids implying that the work is intended as a replacement for more complex operational differentiable LSM frameworks. At the same time, we believe the results provide a useful and rigorous contribution to the growing literature on differentiable Earth-system modelling, particularly in relation to inverse problems and parameter identifiability.

We therefore hope that the revised manuscript will now be considered suitable for publication in GMD.

Kind regards,

Maarten van Reeuwijk

(on behalf of all authors)