
Response to Reviewers

Geoscientific Model Development

Manuscript: egusphere-2026-2288

Title: An ESMValTool-based framework for sanity checks, physical consistency and climate fidelity during model development – ICONEval v1.0

Authors: Axel Lauer, Manuel Schlund, Lisa Bock, Birgit Hassler, Gunnar Behrens, Bettina Gier, Lukas Lindenlaub, Stephan Lorenz, Jan-Hendrik Malles, Wolfgang A. Müller, Trang van Pham, Katja Weigel, Guang Zeng, and Veronika Eyring

Date: 13 August 2026

We thank Jaydeep Singh for his helpful comments. In this document, we answer each point raised by the reviewer. The original reviewer’s comments are given in **black**, our answers in **blue**. A "track changes" version of the revised manuscript highlighting all changes is available.

Reviewer #1

This manuscript by Lauer et al., entitled "An ESMValTool-based framework for sanity checks, physical consistency and climate fidelity during model development – ICONEval v1.0," presents ICONEval, an open-source wrapper around ESMValTool that can perform three sets of checks of model-development: sanity checks, physical-consistency checks, and climate-fidelity diagnostics. The capabilities of the developed tool, are demonstrated on a single ICON-XPP historical simulation. The work addresses a real and timely need of efficient and reproducible evaluation tool for model development. Also, this tool is designed to support emerging applications such as ML-enhanced ESMs and CMIP7. In its current form, however, the manuscript has a few weaknesses that should be resolved before publication. For these reasons, I recommend major revision, and the following are the comments.

We thank Reviewer #1 for providing helpful comments to improve the manuscript.

Major Comments

1. The Introduction provides a general motivation for model evaluation, but it does not sufficiently explain the specific need for ICONEval or its focus on ICON/ICON-XPP. The manuscript moves quickly from a broad discussion of ESM development and ESMValTool to the presentation of ICONEval, without clearly identifying the gap in existing evaluation workflows that this framework is intended to address. The authors should better explain why ICON-XPP requires a dedicated evaluation framework, for example, in relation to its role in CMIP7, its coupled configuration, its development workflow, or the practical challenges of rapidly assessing new model versions and parameterizations. In addition, the three central categories, sanity checks, physical consistency checks, and climate-fidelity diagnostics, can be defined more clearly in the Introduction. Some of the descriptive material currently in Sections 2.1 and 2.2 could be drawn forward to motivate the framework, while the detailed descriptions remain in Section 2.

We clarified our motivation to focus on ICON-XPP and to introduce ICONEval by describing the gap in existing workflows by adding the following paragraph to section 2.3:

"Initially, we focus on ICON-XPP, a model that will contribute to CMIP7 and is used in operational climate prediction (Müller et al., 2025), with an active development history. For example, several data driven approaches for ICON-XPP aiming at improving selected parameterizations with machine-learning based methods are being developed such as radiation (Hafner et al., 2026), convection (Heuer et al., 2026), cloud cover (Grundner et al., 2025), gravity waves (Haslauer et al., 2026) or cloud microphysics (Sarauer et al., 2025). Such model development activities require continuous performance monitoring and basic evaluation including sanity checks and physical consistency checks

already during the development phase to identify potential problems early on. Particularly during the often-needed retuning of the model when replacing a core parameterization, an efficient and easy to use tool is needed that also allows for checking a larger number of model simulations. This gap is filled by ICONEval, which allows to run a whole suite of diagnostics with a single command line enabling automatic testing e.g. right after a simulation is finished. In order to maximize the user-friendliness, the results are then also visualized. ICONEval provides for this easy-to-read overview pages linking to the individual results accessible via a simple web browser.”

Following the suggestion of the reviewer, we also added a more concise definition of each of the three diagnostic categories to introduction:

”The sanity checks can be used to assess the global representation of selected variables in an ESM simulation to ensure that the values fall within the expected range such as observed minima and maxima. The physical consistency checks can be applied to verify that a model’s equations, parameterizations, and numerical methods do not violate any known fundamental laws of physics. The climate fidelity diagnostics are used to verify that the observed basic characteristics of the atmosphere, the ocean and the land surface such as geographical distributions, seasonal and diurnal cycles or mean states are represented correctly in a model simulation.”

2. The paper’s specific contribution would be clearer if manuscript distinguish more clearly between functionality newly provided by ICONEval and functionality inherited from, or already available within, the underlying ESMValTool framework. At present, the two are sometimes discussed interchangeably, and it should also be made clearer which limitations of current ESMValTool usage are addressed by ICONEval. In addition, the manuscript highlights potential applications to hybrid machine-learning-enhanced (MLE) Earth system models and to community benchmarking standards such as ClimateBench. However, the examples presented are limited to a physics-based ICON-XPP simulation, so these capabilities are not demonstrated. The authors should either present them more explicitly as intended or future applications, including corresponding revisions to the abstract, or add a brief illustrative example showing how ICONEval is used in an MLE or benchmarking context. Finally, because the implementation and demonstration are currently ICON specific, this scope should be stated more explicitly in the title and abstract.

We agree with the reviewer that the new functionalities of ICONEval and old and new functionalities of ESMValTool need to be more clearly distinguished. We clarified this in the revised version of the manuscript by adding the following paragraph to the section 2.3:

”While all diagnostics presented in this paper can in principle also be done with the latest ESMValTool version, the task of developing a model requires an efficient and easy-to-use approach to quickly assess also a large number of simulations. For this, ICONEval provides a ready-made suite of diagnostics that we found useful during model development including model tuning. ICONEval is specifically tailored towards basic automated testing and performance assessment of ICON simulations including a user-friendly visualization of the results.”

The reviewer brings up the point that while application of ICONEval to hybrid (MLe) models is mentioned, only examples from the physics-based ICON are shown. We would like to note that the output of ICON-MLe versions including e.g. ML-based convection, cloud microphysics, cloud cover or radiation, is identical to the physics-based setup. In our opinion, showing more than one model version in the examples would therefore bring no additional insights but only lengthen the then required description of the other ICON versions and model experiments. Several different of these simulations would be needed as none of these simulations covers all domains addressed in ICONEval (atmosphere, ocean, land). We therefore think a single model run to illustrate the examples is sufficient and prefer to keep this single, well documented physics-based model simulation as an example as it provides output for all domains covered here (atmosphere, ocean, land).

The reviewer is right that the current implementation and demonstration of ICONEval is ICON specific. The workflow and the features presented can also be used for different models as long as the model output is either CMORized or the model is supported directly by ESMValTool. We clarified this in the abstract by adding the following paragraph:

”While ICONEval is focused on the ICON model, the underlying ESMValTool diagnostics can also be used for different models as long as the model output is following the CMOR (Climate Model Output Rewriter) standard or the model is supported directly by ESMValTool.”

3. Section 4.2 / Fig. 3 /Line 345: This is a useful physical-consistency check, but its interpretation should be clarified. Clausius-Clapeyron describes saturation vapor pressure; a 6-7% K-1 scaling of column water vapor additionally assumes approximately constant relative humidity. Therefore, describing the prw-T2m relation as ”close to the Clausius Clapeyron equation” is imprecise. Please describe the regression as an empirical annual mean temperature-moisture relationship over the selected period, rather than a direct warming-response scaling. Finally, since the result depends on the analysis period, ERA5 and ICON-XPP should be compared over the same period (1993-2014); the longer ERA5 record can be shown separately for context.

As suggested, we harmonized the time periods shown in figure 3 and rephrased the description of the regression as ”an empirical annual mean temperature-moisture relationship over the selected period”.

Minor comments

- Line 32: reanalysis → reanalysis.

Corrected, thanks for spotting this.

- Line 41: ’of essential climate variables are helpful...’ add few variables variables?

We added ”selected” before ”essential climate variables”.

- Line 45: It is suggested to add the description of ClimateBench.

We added the following brief description of ClimateBench to the introduction: "Climate Bench is an open-source machine learning dataset and benchmarking framework to emulate complex Earth system models for different emission scenarios."

- Line 50: "In the following, the tools, methods and datasets used are briefly described", add the paragraph at the end of the Introduction section regarding the flow of the manuscript. This line here is redundant.

As suggested, we moved this line to the end of the introduction.

- Line 52: ESMValTool has been already defined in earlier sections.

We removed "ESMValTool" from the parentheses.

- Lines 83-84: Since the details of ICON-XPP are already given in the published Müller et al. (2025), the statement "More detail about the final ICON-XPP configuration for CMIP7 can be found in Pham et al. (in preparation)" should be modified or removed, as an "in preparation" reference is not retrievable by readers.

We removed the sentence referring to "Pham et al. (in preparation)".

- In Table 1, the variable 'surface downwelling longwave radiation (rlus)' has to be checked for MERRA2 and ERA5. Isn't it rlds? Please check the surface pressure for ERA5, psl used for sea level pressure not for surface pressure.

Thanks for spotting this. "rlus" is the "surface upwelling longwave radiation", "psl" in the table should be "ps". This has been corrected.

- Check the range of Global mean TOA net downward radiation (166.3 - 231.8 W m⁻²). This values is very inconsistent.

This is the range for the surface (and not TOA as wrongly indicated in table 2) net downward radiation. We corrected this in the table.

- Is the CLARA-A2.1 in Table 1 and CLARA-AVHRR in Table 2 is similar?

Yes, this is the same dataset. We changed "CLARA-A2.1" to "CLARA-AVHRR" in table 1 to be consistent throughout the paper.

- Global net moisture flux into the atmosphere units in Table 2 should be mass per unit time.

Thanks for spotting this. We corrected the unit from "kg" to "kg s⁻¹".

- Line 111: 'listed in 2' → listed in Table 2?

We added the missing "Tab." Before "2".

- Lines 103-105: The opening of Section 3 contains a useful conceptual definition of sanity checks ("... the global representation of selected variables... the temporal evolution of global means, minima, and maxima can be assessed"). A concise definition of each of the three categories should be introduced where they are first named in the Introduction. The methodological detail on how the bounds are derived (lines 104-105) can remain in Section 3.

Following the suggestion of the reviewer, we added a paragraph to the introduction defining each of the three diagnostic categories more precisely (see text changes in reply to major comment #1).

- Figure 3/Line 185: The axis labels ("Delta 2m temperature [K]", "Delta water vapor [%]") use different notation from the text, which defines the regression variables as ΔT_{2m} (in K) and Δprw (in %). Please make the figure and text consistent and define ΔT_{2m} and Δprw in the figure caption for clarity.

We updated the labels and the caption of figure 3 as suggested.

- Line 235: biogeogemistry $\rightarrow$ biogeochemistry.

Corrected.

- Lines 244-248: The sentence beginning "These diagnostics were selected since they illustrate..." and ending "... temperature distributions (e.g. Bock et al., 2020)" is very long and please rephrase it as two or three shorter sentences.

We split the sentence into two.

- Line 273: the the El Niño... $\rightarrow$ the duplicated

Thanks for spotting. Corrected.

- Line 309: ICONEVal? Is it ICONEVal? The parenthesis can also be removed.

Corrected ICONEVal to ICONEval (line 76), parentheses removed (line 309).

- Section 5.3: Please replace "lai" with "LAI" (or set it in italics as "lai") and apply this consistently throughout the section.

We capitalized LAI throughout the manuscript.

- Line 356: new ESMValTool $\rightarrow$ ICONEval

Changed as suggested.

- Line 386: The example website provided to showcase ICONEval's output (the swift.dkrz.de link in the Code and Data Availability section) is not publicly accessible (returning an HTTP 401 (Unauthorized)) response. Since the reporting and publishing capability is central to the contribution but is currently described only in text via this external link, please ensure the example site is made publicly accessible without credentials.

We double-checked the link and the site should be publicly accessible via the link provided. We suspect that a recent system downtime due to security issues might have led to the access error described by the reviewer.

General / language

In several parts of the manuscript, information is unnecessarily placed in parentheses, which interrupts the flow and makes some sentences harder to read. Please revise these passages to integrate the parenthetical content into the main sentence where appropriate, and reduce the use of parentheses overall.

We went through the manuscript and removed parentheses where possible, i.e. in lines 19, 30, 64, 93, 97, 120, 121, 156, 157, 210, 211, 232, 241, 253, 257, 295, 303, and the captions of figures 1, 2, 15, 17 and of table 1.

References

- Grundner, A., Beucler, T., Savre, J., Lauer, A., Schlund, M., and Eyring, V.: Reduced cloud cover errors in a hybrid AI-climate model through equation discovery and automatic tuning, *Scientific Reports*, 15, <https://doi.org/https://doi.org/10.1038/s41598-025-29155-3>, 2025.
- Hafner, K., Shamekh, S., Bertoli, G., Lauer, A., Pincus, R., Savre, J., and Eyring, V.: Representing subgrid-scale cloud effects in a radiation parameterization using machine learning: MLe-radiation v1.0, *Geoscientific Model Development*, 19, 3875-3891, <https://doi.org/10.5194/gmd-19-3875-2026>, 2026.
- Haslauer, E., Schwabe, M., Dörnbrack, A., Gerber, E., Rapp, M., Ålagar, N., and Eyring, V.: Interpretable neural networks to predict momentum fluxes of orographic gravity waves, *Machine Learning: Earth*, 2, <https://doi.org/10.1088/3049-4753/ae7df2>, 2026.
- Heuer, H., Beucler, T., Schwabe, M., Savre, J., Schlund, M., and Eyring, V.: Beyond the Training Data: Confidence-Guided Mixing of Parameterizations in a Hybrid AI-Climate Model, *Journal of Advances in Modeling Earth Systems*, 18, e2025MS005, <https://doi.org/https://doi.org/10.1029/2025MS005544>, e2025MS005544 2025MS005544, 2026.
- Müller, W. A., Lorenz, S., Pham, T. V., Schneidereit, A., Brokopf, R., Brovkin, V., Brüggemann, N., Chegini, F., Dommenges, D., Fröhlich, K., Früh, B., Gayler, V., Haak, H., Hagemann, S., Hanke, M., Ilyina, T., Jungclaus, J., Köhler, M., Korn, P., Kornbluh, L., Kroll, C. A., Krüger, J., Castro-Morales, K., Niemeier, U., Pohlmann, H., Polkova, I., Potthast, R., Riddick, T., Schlund, M., Stacke, T., Wirth, R., Yu, D., and Marotzke, J.: The ICON-based Earth System Model for climate predictions and projections (ICON XPP v1.0), *Geoscientific Model Development*, 18, 9385-9415, <https://doi.org/10.5194/gmd-18-9385-2025>, 2025.
- Sarauer, E., Schwabe, M., Weiss, P., Lauer, A., Stier, P., and Eyring, V.: A physics-informed machine learning parameterization for cloud microphysics in ICON, *Environmental Data Science*, 4, e40, <https://doi.org/10.1017/eds.2025.10016>, 2025.

We thank the reviewer for his/her helpful comments. In this document, we answer each point raised by the reviewer. The original reviewer’s comments are given in **black**, our answers in **blue**. A ”track changes” version of the revised manuscript highlighting all changes is available.

Reviewer #2

In the paper, the authors introduce ICONEval, an open-source evaluation framework built on top of the Earth System Model Evaluation Tool (ESMValTool) designed for rapid, reproducible, and automated performance monitoring during the development of Earth System Models (ESMs). ICONEval introduces parallel processing and HTML summary reports categorized into three diagnostic pillars: sanity checks against observational bounds, fundamental physical consistency checks (e.g., mass conservation), and multi-domain climate fidelity diagnostics. The tool is demonstrated using the ICON-XPP model and is positioned as a highly efficient benchmarking utility for comparing model configurations, testing updated parameterizations, and assessing emerging hybrid machine learning-enhanced climate models.

This manuscript is well-suited for the journal’s scope, as it introduces an open-source evaluation framework (ICONEval) that directly advances our ability to benchmark, monitor, and improve ESMs, including emerging hybrid machine learning configurations. The paper is well-prepared, presenting a highly structured and logical three-pillar approach—spanning sanity checks, physical consistency, and climate fidelity—that is clearly articulated and thoroughly demonstrated using the ICON-XPP model. Given the increasing complexity of modern climate modeling, the development of reproducible, automated, and customized tools for specific models is highly relevant to the journal’s audience and represents a timely, high-quality contribution to the community.

I suggest accepting the manuscript subject to the authors addressing the following a few specific comments.

We also thank Reviewer #2 for helping us to improve the manuscript.

Specific comments

Line 145 and Figure 2: Still, it is not very clear to me that how the ”reasonable” upper and lower limits (red lines in the figure) were defined, in particular with using multiple reference datasets. For example, for minimum, was the minimum monthly value across all years of global spatial average taken from each reference dataset, then minimum value among them was taken as the ”reasonable minimum lower limit?” I suggest further clarifying it in the manuscript.

The reviewer is right. When multiple reference datasets are available, the minimum and maximum global monthly mean values are obtained as the minimum and maximum values across all global values. We clarified this in the manuscript as follows:

””Reasonable” upper and lower limits for global means are defined using the global minimum and maximum monthly values across all available years of multiple observational and reanalysis datasets. This is equivalent to the lowest and highest monthly global average value observed in any of the reference datasets.”

Figure 8a: While I understand there could be some artifact due to difference in the data frequency interval (i.e., ERA5 1 hourly while ICON-XPP 6 hourly), the sudden magnitude change over the pacific seems odd, which could have been resulted from incorrect local time handling. Could you please double check on this?

We agree with the reviewer that particularly the stripe over the Pacific Ocean looks strange. We double-checked the calculation and believe everything is correct. When converting UTC to local solar time, the temporal frequency of the input data is preserved. For example, hourly input data will be transformed into hourly output data. For this, a location's exact LST will be put into corresponding bins defined by the bounds of the input time coordinate (in this example the bin size is 1 hour). This means that the lower the time resolution of the input data, the larger the resulting errors. The errors can reach up to time frequency / 2 explaining the bands visible in figure 8. For this reason, a time resolution of 1 hour or better is recommended for this diagnostic. We added the following clarification to the text:

"When converting from UTC to local solar time (LST), the temporal frequency of the input data is kept. A location's exact LST will be put into corresponding bins defined by the bounds of the input time coordinate. This means that the lower the time resolution of the input data, the larger the resulting errors that can reach up to time frequency / 2. It is therefore recommended to use this diagnostic with hourly or higher temporal resolution."

Figure 9. In addition to showing the difference field only, I suggest showing the full field from each data (ICON-XPP and ERA5) as well, so the pattern could be easily compared, like in Figure 10.

Following the suggestion of the reviewer we extended figure 9. The revised version now also contains the full fields from ICON-XPP and ERA5 in addition to the differences between the two.

Figure 11. In the Historical experiment, year-to-year comparison between model and observation does not say much, as their matching is not even expected. Although the authors have described this in the manuscript (line 275-276), I wonder if it would be beneficial to re-emphasize this point in the figure caption so advise readers what to focus on the figure. One additional option could be comparing the standard deviation or variance values obtained from each time series as an indicator for the intensity of the variability, or trend line. The same comment applies to Figs. 12 and 14 as well. For Fig. 11, I am not sure how much useful information could be derived from the CMIP6 results in the background.

We added the following reminder to the caption of figure 11:

"Historical ESM simulations are not expected to replicate the timing of distinct observed ENSO events, but should have a comparable amplitude and frequency of the interannual variability."

Following the suggestion of the reviewer, we added the temporal standard deviations of the examples model (ICON-XPP) and the reference dataset (ESACCI-SST) to the figure as a measure of the variability.

We would prefer to keep the CMIP6 models in the figure to provide an envelope for the amplitude range that similar models show. This helps to quickly assess with this consistency check whether potential deviations from the observed amplitude are within the typical range of other ESMs. We clarified this in the manuscript by adding the following paragraph:

”The ensemble of CMIP6 models included in this figure provide an envelope for possible anomaly amplitudes that can be expected from current ESMs. As a measure for the strength of the internal variability, the temporal standard deviations of the example model and the reference dataset are given in the top of the panel.”

Similarly, we also added the temporal standard deviation to figures 12 and 14 and a reminder that historical ESM simulations are not expected to reproduce distinct events to the figure captions.

We think this does not apply to the mean seasonal cycles shown in figure 13, which therefore remains unchanged.

The original comments of GMD's executive editor Astrid Kerkweg are given in **black**, our answer in **blue**.

Executive Editor of GMD Astrid Kerkweg

Dear Axel and Co-authors,

please note that GitLab links are not enough to refer to the code used. Thus, please cite the DOI of the respective ICON release. (I guess doi:10.35089/WDCC/IconRelease2025.10) in the code availability section. If this is not the correct DOI and the respective ICON (sub)version is not accessible via DOI please store the version in a permanent archive yourself or get the ICON C5 to provide also a doi for this subversion, and cite that DOI.

Best regards, Astrid Kerkweg (executive Editor of GMD)

Dear Astrid,

Thank you for pointing this out. The doi for the ICON version used in our paper is as you already guessed 10.35089/WDCC/IconRelease2025.10. We added this information to the code availability section.

Kind regards,

Axel Lauer