Major changes

We thank Reviewer for the detailed and constructive review. The comments led to a major revision of the manuscript. Beyond correcting technical issues, we substantially redefined the scientific scope of the study.

The manuscript has been transformed from a station-specific Bayesian prediction study into a Bayesian nonlinear hierarchical analysis of atmospheric processes. The revised study jointly models aerosol loading, cloud cover and clearness index, emphasizes conditional associations rather than causal inference, compares tropical and temperate climates (Côte d'Ivoire and Germany), and greatly expands the physical interpretation and discussion.

General comments

This paper investigates the dependence of Global Horizontal Irradiance (GHI) on weather conditions at the surface (temperature, humidity, wind speed), cloud cover and aerosol concentration, for 12 stations located in Côte d'Ivoire. To this end it relies on a sophisticated Bayesian model to derive the posterior distributions of the parameters relating the predictors to the predicted variable. The paper is overall well written, in particular the introduction that gives a nice overview of the related works, with a special focusing on the uncertainty estimation for the prediction models. The paper is very technical, particularly hard to follow for non-statisticians, and lacks of physical interpretations. In many places reordering would help, and some paragraphs are not specifically useful and could be removed or at least condensed. The section dedicated to the analysis of the results is very general and does not demonstrate the interest of the work in terms of processes understanding or improved prediction capability, which somehow deserves the primary objective of the paper. Likewise the discussion is very limited, while obviously there would be much to say. Unless the paper is augmented in terms of original content and its added value to the field is better emphasized, I doubt it can be of great interest for the readers of GMD, except maybe for the pure methodological approach that could be used by others.

The originality is now clearly stated. The contributions include: (1) quantify the nonlinear conditional associations between temperature, relative humidity, wind speed, aerosols, cloud cover and surface solar irradiance; (2) characterize the interaction between aerosols and cloud cover and evaluate how their combined influence on solar irradiance varies across atmospheric conditions; (3) identify climate-dependent differences in the atmospheric regulation of solar irradiance by comparing tropical (Côte d'Ivoire) and temperate (Germany) environments; (4) account for spatial heterogeneity through hierarchical modelling of station-level variability while fully propagating posterior uncertainty; and (5) provide a physically interpretable probabilistic approach for improving the understanding of atmospheric controls on surface solar irradiance and their implications for solar resource assessment.

The objective has been clarified and is now explicitly stated as process understanding, rather than giving the impression that the work is focused solely on prediction. The Methodology section was reorganized and simplified, the Results section rewritten around physical interpretation, numerous new figures added (posterior predictive checks, nonlinear response curves, interaction surfaces, seasonal effects, quantitative GHI contrasts, station effects).

Specific comments

1) There has been a lot of recent work on the forecasting of solar energy. It should be better emphasized what is the originality of the present study. In terms of methodology, efficiency (then it should be compared to state-of-the-art alternative models), location, etc. These elements should be motivated in the introduction, and highlighted in the conclusion. Likewise, in Section 3.3 some dependences between GHI and meteorological variables are highlighted, that could probably be mentioned in a dedicated paragraph of the introduction, to point out what is already known in terms of such correlations.

We sincerely thank the reviewer for this valuable suggestion. We have revised the Introduction to better emphasize the originality and scientific contribution of this study. Specifically, we now clearly distinguish our work from previous studies by highlighting that the proposed Bayesian Nonlinear Hierarchical Structural Model (BNHSM) jointly represents nonlinear atmospheric relationships, hierarchical spatial variability and uncertainty propagation within a unified probabilistic framework. We also expanded the motivation for the comparative analysis between the tropical climate of Côte d'Ivoire and the temperate climate of Germany to demonstrate the broader applicability of the proposed framework beyond a single climatic region. Furthermore, we added dedicated paragraphs reviewing the current understanding of the relationships between solar irradiance and key meteorological variables, including temperature, relative humidity, wind speed, aerosols and cloud cover. This additional background provides the physical context for the hypotheses investigated later in the manuscript. Finally, the Conclusion has been revised to more clearly highlight the novelty and methodological contributions and practical implications of the proposed framework.

2) It is not clear why it is useful to derive such prediction models, instead of simply measuring the GHI. Maybe the point would be to use forecasts of meteorological parameters and to predict GHI from that. In this case, would the estimation be better than the GHI forecasts itself (from satellite products of weather models?). A comparison could help convince the reader of the utility of the new model. Otherwise, the physical understanding of the links can in itself be valuable, but it is very weak *in this paper, so it does not bring much novel physical insight to the community. The applications could be mentioned in the discussion, or at least in the opening of the conclusion.*

We thank the reviewer for this important comment. We agree that the original version of the manuscript did not sufficiently clarify the primary objective of the study. In the revised manuscript, we have repositioned the study. Rather than developing a forecasting model intended to replace direct GHI measurements or existing satellite-based products, the proposed Bayesian Nonlinear Hierarchical Structural Model (BNHSM) is designed to improve the understanding of the nonlinear atmospheric dynamics regulating surface solar irradiance. Specifically, the revised model explicitly represents the conditional relationships among meteorological variables, aerosols, cloud cover and the clearness index within a unified probabilistic model. We have also clarified that the objective is not to outperform operational satellite products or numerical weather prediction systems in forecasting accuracy, but rather to provide a physically interpretable framework for investigating atmospheric interactions and their influence on solar irradiance. Following the reviewer's suggestion, the practical applications of the proposed framework have been expanded in both the Discussion and the Conclusion.

3) The paper is very focused on the statistical model, in a way that is hard to follow for non-experts. To increase the readability it would be great to better explain the physics behind the parameters, scores, tools etc. This probably implies a partial rewriting of the Method section.

We thank the reviewer for this valuable suggestion. The Methodology section has been extensively revised with new figures to improve its readability for non-specialists. The presentation of the model has been simplified by reducing the number of equations and introducing a conceptual description of the model before the mathematical formulation. A new conceptual diagram illustrating the hypothesized atmospheric pathways has been included to facilitate understanding of the model structure. The physical mechanisms linking variables are now described in greater detail.

4) As said at point 2) the physical interpretation of the results is extremely limited. While the method only puts forward correlations, the authors try to find some causality, and for each explanation we could argue for the opposite causal chain. Also the fact that covariance between variables is not really discussed is such that apparent correlations between a variable and GHI can be directly related to the correlation of one variable with the other. The opposite may also happen when the authors do not see a strong impact of PM2.5, the inclusion of which sounded however as an originality of the paper. In practice, the authors mention many physical processes that can explain the link between meteorological variables and GHI, but they struggle to convince the reader about what really drives the correlations. Section 3.3 hence appears as a sum of physical explanations and *suggestions without clear, ordered conclusions. This is critical for a method that puts forward the interpretability.*

We sincerely thank the reviewer for this insightful and constructive comment. We agree that the previous version did not sufficiently distinguish statistical associations from causal relationships and that the physical interpretation of the results required improvement. In the revised manuscript, the interpretation of the model has been fundamentally clarified. The proposed Bayesian Nonlinear Hierarchical Structural Model (BNHSM) is now explicitly presented as a framework for estimating conditional associations along hypothesized atmospheric structural pathways, rather than for causal inference. We clearly state that the directed pathways reflect current physical understanding of the atmospheric system but do not imply causal relationships, since observational data alone are insufficient to establish causality. To better account for the interdependence among atmospheric variables, the revised model explicitly represents the conditional relationships linking meteorological conditions, aerosols, cloud cover and the clearness index. Instead of analysing each predictor independently, the model estimates their associations while accounting for their hypothesized structural dependencies. This reduces the risk of misleading interpretations arising from covariation among predictors. In addition, the Results and Discussion sections have been extensively rewritten and reorganized. The interpretation is now structured around the major atmospheric processes, including meteorological regulation, aerosol effects, aerosol–cloud interactions, seasonal variability, and climate-dependent differences between Côte d'Ivoire and Germany.

5) The paper is somehow too short, with the largest share dedicated to the statistical model. It would benefit from additional figures (ex: an illustration of the performances of the final model in terms of predictions vs observations). The discussion is definitely too short as well, and very weak as is.

We sincerely thank the reviewer for this valuable suggestion. We agree that the previous version placed excessive emphasis on the statistical methodology, while the interpretation and discussion of the results were insufficiently developed. In the revised manuscript, nous avons porte plus d’attention a la presentation des resulats et a leur interpretation. Several new figures have been added to evaluate model performance and to facilitate the interpretation of the results. These include posterior predictive checks, Bayesian R^2, residual analyses, PSIS-LOO and Pareto-k diagnostics, nonlinear response curves, aerosol–cloud interaction surfaces, seasonal patterns and physically interpretable contrasts expressed as changes in GHI. The Results and Discussion sections have also been extensively expanded and reorganized. The revised manuscript now presents quantitative findings with their physical interpretation and discusses the differences observed between the tropical climate of Côte d'Ivoire and the temperate climate of Germany.

Technical comments

l.5: "this relationship" is poorly defined as so far the predictors are not explicited. Do they correspond to individual weather variables such as temperature, humidity, cloud cover, liquid water path? In which case it would be obvious that solar radiation is not linearly related to some of those variables.

We thank the reviewer for this helpful comment. We agree that the phrase *"this relationship"* was insufficiently defined in the original manuscript. In the revised version, we explicitly specify that the predictors include the main atmospheric drivers of surface solar irradiance, namely air temperature, relative humidity, wind speed, aerosol loading and cloud cover, with their interactions. We also clarify that the nonlinear relationships refer not only to the individual associations between these atmospheric variables and solar irradiance, but also to the interaction effects among predictors (e.g., aerosol–cloud interactions) that jointly influence surface solar irradiance. Although such nonlinear behaviour is well recognized in atmospheric processes, many conventional statistical approaches still approximate these relationships using linear formulations, which may not adequately capture threshold effects or saturation for example. This clarification has been incorporated into the Introduction to better motivate the proposed nonlinear Bayesian architecture.

l.8: "prior knowledge" of what?

We thank the reviewer for this helpful comment. We agree that the expression "incorporates prior knowledge" was ambiguous in the original manuscript. Our intention was not to suggest that previous knowledge of the atmospheric relationships was imposed on the model. Rather, the Bayesian framework employs weakly informative prior distributions to regularize parameter estimation, improve numerical stability and constrain parameters to physically plausible ranges while allowing the observational data to dominate the posterior inference. We have therefore revised the manuscript to explicitly describe the role of the prior distributions and avoid the potentially misleading expression "prior knowledge."

l.12: if the graph neural networks is a perspective (is it an alternative to the Bayesian approach?), it should not appear (at least that early) in the abstract.

We thank the reviewer for this valuable comment. We agree that introducing Graph Neural Networks (GNNs) as a future research direction in the abstract was premature and could be confusing, particularly because they are not part of the proposed methodology. Our intention was to emphasize that the significant station-specific effects identified by the hierarchical model suggest the existence of spatial dependencies among observation sites that could potentially be exploited by graph-based learning approaches in future studies. However, we agree that this perspective is beyond the scope of the present work. Accordingly, the abstract has been revised and no longer includes any reference to Graph Neural Networks or future methodological extensions. The revised abstract focuses exclusively on the objectives, methodology, principal findings and contributions of the present study.

l.16: as the domain of application (space and time) has not be defined, allusion to "similar tropical climates" is unclear. The abstract clearly lacks of a mention of the data used.

We sincerely thank the reviewer for this helpful comment. We agree that referring to *"similar tropical climates"* without first defining the geographical and temporal scope of the study was unclear. In the revised abstract, we now explicitly describe the study domain by specifying the two countries investigated (Côte d'Ivoire and Germany), the observational datasets used and the study period. The reference to *"similar tropical climates"* has been removed to avoid ambiguity and to ensure that the abstract accurately reflects the scope of the present study.

l.26: it is not clear whether "predict" refers to some leadtime (in the future) or to estimate the present solar radiation from present weather characteristics.

We thank the reviewer for this helpful comment. The use of the term *"predict"* was ambiguous in the original manuscript. In the present study, the proposed model estimates the contemporaneous solar irradiance from atmospheric conditions observed at the same time, rather than forecasting solar irradiance at a future lead time. To avoid any misunderstanding, we have replaced the term *"predict"* with *"estimate"* or *"model"*, where appropriate and clarified this distinction throughout the revised manuscript.

l.46: when citing a paper its content should be a bit more detailed. Here the reported result is too general (no information on location, type of situations etc.).

We thank the reviewer for this helpful comment. We agree that the description of the cited study was too general. In the revised manuscript, we have expanded the corresponding sentence to briefly summarize the main findings of the cited work, including the study region, the atmospheric conditions investigated, and the principal conclusions that are relevant to the present study.

l.46-47: remove parentheses for citations. Beware also with the use of \cite{}

We thank the reviewer for pointing this out. The citation style has been revised throughout the manuscript. Parentheses have been removed where inappropriate and the use of LaTeX citation commands has been corrected to ensure consistent formatting across all references.

l.50: it's not clear why adding a possibly useless predictor can be detrimental to the model performance.

We thank the reviewer for this helpful comment. We agree that the original statement was insufficiently justified. The corresponding sentence has been revised to clarify that the relevance of atmospheric predictors depends on the climatic and environmental conditions of the study region. Predictors that are informative in one climate may contribute little additional explanatory power in another, increasing model complexity without necessarily improving model performance.

l.51: when referring to air pollutants, it would be more rigorous to mention the actual quantities used (concentrations near the surface, total column, AOD?). Also we expect given this sentence that you'll use BC in your study.

We sincerely thank the reviewer for this helpful comment. We agree that referring to *air pollutants* without specifying the corresponding quantity could lead to ambiguity. In the revised manuscript, we now explicitly state that aerosol loading is represented by the Aerosol Optical Depth at 550 nm (AOD550). In addition, we have removed references to atmospheric variables that are not included in the present analysis, such as black carbon (BC), to ensure that the description of the predictors is fully consistent with the methodology.

l.68: "efficiency" has not been defined, and can be misinterpreted.

We sincerely thank the reviewer for this helpful comment. We agree that the term *"PV efficiency"* was ambiguous and could be misinterpreted. During the major revision of the manuscript, the Introduction was substantially rewritten, and the corresponding statement has been removed. The revised Introduction now focuses on the scientific motivation of the study and no longer refers to photovoltaic efficiency in this context.

l.70: in the mentioned "products" the variables are not really "parameterized", they're rather "estimated".

We thank the reviewer for this helpful comment. We agree that the term *"parameterizing"* was inappropriate in this context, since meteorological and aerosol variables in satellite-based and reanalysis products are estimated or retrieved rather than parameterized. As part of the major revision, the Introduction has been substantially rewritten, and the corresponding paragraph has been removed. Consequently, this issue is no longer applicable in the revised manuscript.

l.70-71: "due to limited data availability and an incomplete understanding of their regional dynamics" is unclear. You can mention their intrinsic limitations (spatial resolution, precision, etc.)

l.72: satellite products of which quantities?

l.75: why would satellite miss these low-level clouds? Small size, poor contrast with the surface? Any reference to support this statement?

l.75: is CAMS a satellite product? I mostly know the CAMS model

l.75: "the authors of the present study found" could be turned into "it was recently shown", as the citation comes in the end

l.76: please specify the considered variables.

We sincerely thank the reviewer for these helpful comments. We agree that several statements in the original Introduction lacked sufficient precision regarding the atmospheric variables considered, the satellite and reanalysis products used, and their associated limitations. As part of the major revision, the Introduction has been substantially rewritten and reorganized. The revised version now provides a clearer motivation for the study, explicitly defines the atmospheric variables investigated, accurately describes the data sources and presents a more focused review of the relevant literature. Consequently, the specific statements referred to in these comments have been removed or replaced, and the concerns raised are no longer applicable to the revised manuscript.

We agree that CAMS itself is not a satellite product but a modelling and data assimilation service. Our original wording was intended to refer specifically to the CAMS Radiation Service (CRS), whose surface solar irradiance products are derived by combining geostationary satellite observations with atmospheric information from the CAMS modelling system. Consequently, although these irradiance products are frequently referred to in the literature as *satellite-derived* or *satellite-based* products, we agree that this terminology may be imprecise.

l.84: an outline of the paper at the end of the introduction would be welcome. To detail what kind of data are used in particular.

We sincerely thank the reviewer for this helpful suggestion. We agree that an outline of the manuscript at the end of the Introduction improves readability and helps guide the reader through the paper. Accordingly, a concluding paragraph has been added to the Introduction summarizing the organization of the manuscript. The new outline explicitly indicates where the study areas, ground-based and atmospheric datasets, data preprocessing procedures, the proposed Bayesian methodology, model evaluation diagnostics, the main results, and the discussion are presented.

l.87: could the twelve stations be highlighted in Fig. 1?

l.109: the objective of this study

l.113: in details

l.113: the references should not be in parentheses.

We sincerely thank the reviewer for these helpful comments. The revised manuscript has been carefully proofread to address the remaining typographical, grammatical and formatting issues. Specifically, the citation formatting has been revised to conform to the journal's style. Regarding the suggestion to highlight the twelve stations in Figure 1, this comment is no longer applicable because the figure has been removed from the revised manuscript.

l.114: what is cloud cover here? A fraction of cloud (%) or a cloud mask at the spatial resolution of the satellite product? In this paragraph the satellite products used and their characteristics (resolution, frequency, type (column, profile), etc.) should be detailed. Also, do not mention satellite products if a reanalysis is actually used (for aerosols).

We sincerely thank the reviewer for this helpful comment. We agree that the description of the atmospheric variables and data sources required greater precision. In the revised manuscript, cloud cover is explicitly defined as the Total Cloud Cover (TCC), expressed as the fraction of the sky covered by clouds (ranging from 0 to 1). In addition, the Data section has been substantially revised to provide a comprehensive description of all datasets, including their spatial and temporal resolutions and data sources. We have also carefully revised the terminology throughout the manuscript to clearly distinguish satellite-derived products from reanalysis products. In particular, aerosol loading is now represented by AOD550 and its source is described.

l.119: any motivation to only consider PM2.5 (not larger particles) to estimate the impact on solar radiation? Also do you mean here mass concentration again?

Eq.1: I have the feeling that if some particles grow they become larger than 2.5 micrometers, hence would not be included anymore in [PM2.5]. Does it deserve some explanation?

l.125: what is a wet/dry relative humidity? 0 and 100%? Is this naming standard? As it sounds a bit awkward

l.128: should this ratio fOM:OC appear in Eq. 1?

We sincerely thank the reviewer for these insightful comments. We agree that the previous methodology for deriving PM2.5 from individual aerosol species required additional clarification regarding the aerosol components considered, the hygroscopic growth factors and the assumptions underlying Equation (1). As part of the major revision, this methodology has been completely removed from the manuscript. In the revised study, aerosol loading is represented directly by the Aerosol Optical Depth at 550 nm (AOD550), which is physically more appropriate for investigating aerosol–radiation interactions because it directly quantifies the atmospheric aerosol burden affecting the attenuation of incoming solar radiation. Consequently, the PM2.5 formulation, the hygroscopic growth equations and the associated assumptions (including the OM/OC conversion factor) are no longer part of the revised manuscript.

l.146: "pollution" seems loosely defined. Does it correspond to all aerosols? Why not using aerosols instead?

We sincerely thank the reviewer for this helpful comment. We agree that the term *"pollution"* was too broad and could be misleading, since the MERRA-2 products represent atmospheric aerosols rather than anthropogenic pollution alone. In the revised manuscript, this terminology has been corrected. The analysis now explicitly uses Aerosol Optical Depth at 550 nm (AOD550) as a measure of atmospheric aerosol loading, and all references to *"pollution"* have been replaced by the more appropriate term *"aerosols"* or *"aerosol loading"*, where applicable.

l.153: not clear what you mean by "attenuation" here. The fact that the air mass increases as the Sun is low, or just the fact that the Sun is low hence the GHI lower?

We sincerely thank the reviewer for this helpful comment. We agree that the original wording was imprecise. The reduction of surface solar irradiance outside midday hours is primarily governed by changes in solar geometry, particularly the increasing solar zenith angle, rather than by atmospheric attenuation alone. As part of the major revision, the corresponding preprocessing procedure has been removed, and the revised manuscript instead explicitly accounts for deterministic solar-geometry effects through the use of the clearness index, and by adding seasonal and diurnal terms to the model.

l.155: the wording is surprising. Maybe just state that all values outside the range were discarded because unexpected? But then what does it mean if you discard some values (and how many?), that the measurements are not reliable? Hence they could be unreliable also for values within the range?

We sincerely thank the reviewer for this insightful comment. We agree that the justification for excluding observations outside the prescribed temperature range was not sufficiently rigorous in the original manuscript and could lead to confusion regarding data quality. As part of the major revision, the preprocessing strategy has been substantially revised, and this temperature-based filtering criterion has been removed. The revised manuscript therefore no longer excludes observations based on predefined temperature thresholds.

l.157: could you specify for which variables the outliers are removed?

l.164: I think you can use GHI here

l.166: it seems that some predictors are missing Table 2: the acronyms should be understandable

l.182: not clear in Eq. 4 what is known and what should be retrieved (the parameters)

l.192: what are the "global parameters"?

l.192: "combination" is unclear. Is it a sum, a product? How are the actual nonlinear functions chosen for each predictor?

We sincerely thank the reviewer for these detailed and constructive comments. We agree that several aspects of the original preprocessing and model specification sections required clarification, including the preprocessing strategy, variable definitions, model notation, parameter interpretation and the mathematical specification of the nonlinear functions.

As part of the major revision, these sections have been completely rewritten. The previous Bayesian Nonlinear Mixed-Effects (BNLME) formulation has been replaced by a new Bayesian Nonlinear Hierarchical Structural Model (BNHSM) specifically developed to investigate the conditional relationships among meteorological variables, aerosols, cloud cover and the clearness index. The preprocessing strategy has also been substantially simplified and the previous filtering procedures, generic nonlinear mixed-effects formulation and associated notation have been removed.

The revised methodology now defines all variables and datasets, describes the proposed model in terms of its hierarchical Bayesian structure, specifies the role of each parameter and details the nonlinear components and interaction terms.

l.194: even after normalization there remains an impact of the diurnal cycle due to the air mass.

We sincerely thank the reviewer for this important observation. We agree that normalizing GHI using the clearness index substantially reduces, but does not completely eliminate, the influence of solar geometry. In particular, residual diurnal variability remains because the atmospheric optical path length (air mass) still varies with the solar zenith angle and affects atmospheric extinction, even after normalization. To account for these remaining deterministic temporal effects, the revised model explicitly includes diurnal and seasonal smooth terms. Doing so, the model is able to capture residual temporal variability that is not fully explained by the clearness index alone. The manuscript has been revised accordingly to clarify this point.

l.203: for non-experts, maybe explain why values larger than 1 can be observed, otherwise it suggests that the measurements could be erroneous.

We sincerely thank the reviewer for this valuable comment. We agree that the original manuscript did not sufficiently explain why clearness index values greater than one may occasionally occur. Although the clearness index is theoretically bounded by unity under ideal conditions, values slightly exceeding one can arise in practice due to a combination of physical and observational factors. These include cloud-edge enhancement effects, whereby scattering from nearby clouds temporarily increases the direct and diffuse irradiance reaching the surface, as well as uncertainties in irradiance measurements and clear-sky model estimates. We have revised the manuscript to clarify that such values do not necessarily indicate erroneous measurements. Because the revised model employs a Beta likelihood, whose support is restricted to the open interval (0,1), the observed values are transformed to lie strictly within this interval before model fitting.

l.205: should this Beta distribution be explicited? At least I'm not familiar with it.

We sincerely thank the reviewer for this helpful suggestion. We agree that the original manuscript introduced the Beta distribution too briefly. In the revised manuscript, the Beta likelihood is explicitly motivated by the fact that the response variable, the clearness index, is continuous and naturally bounded within the open interval (0,1). We also clarify that the Beta distribution is parameterized by its mean and precision, making it particularly suitable for modeling bounded continuous variables while accommodating varying levels of dispersion.

l.207: not clear where the precision parameter comes from, and what value it takes.

We thank the reviewer for pointing out this lack of clarity. In the revised manuscript, we explicitly describe the role of the precision parameter ϕ in the Beta distribution. Rather than being fixed, ϕ is treated as an unknown model parameter and is estimated jointly with the regression coefficients within the Bayesian framework. We further specify its prior distribution in the model specification and explain that larger values of ϕ correspond to lower variability of the response around its conditional mean, whereas smaller values allow greater dispersion.

l.225: maybe explicit that the sigmoid fit was always the best one, except for Pit. Also not clear how you compare R values for different sites. Do you choose the function that has the maximum mean R value?

We sincerely thank the reviewer for this insightful comment. We agree that the original manuscript did not clearly describe the procedure used to select the nonlinear functional forms. As part of the major revision, this preliminary model-selection procedure has been completely removed. Rather than selecting predefined functional forms (e.g., sigmoidal or quadratic relationships) based on goodness-of-fit statistics, the revised manuscript employs Bayesian penalized splines within the hierarchical model to estimate nonlinear relationships directly from the data. This approach allows the shape of each predictor–response relationship to be inferred flexibly during model fitting.

l.228: should it be H = 3? Also define dt

Figure2: the caption contains important information about the ranges of observations used, that may deserve to be in the main text. Also, what does it mean to obtain fits for different ranges, sometimes inconsistent? In terms of general applicability of the final regression? This point probably deserves more explanations.

l.231: again it is not clear what the BNLME model does or finds once you've found the regression functions for each variable. What are the free parameters to be optimized?

l.233: not clear what the "raw" covariates are

l.235: what is the motivation of using so many models, while one would expect only the more detailed one to perform best (according to Eq. 8)?

We thank the reviewer for these constructive comments. We agree that the original methodology was not sufficiently clear regarding the selection of nonlinear functional forms, the interpretation of the preliminary regression analyses, the definition of model parameters and the rationale for comparing multiple intermediate model variants.

As part of the major revision, this entire model-building strategy has been fundamentally redesigned. The previous approach, which relied on selecting predefined nonlinear functions (e.g., sigmoidal or quadratic relationships) from exploratory analyses and progressively constructing several BNLME model variants, has been removed. In the revised manuscript, nonlinear effects are represented directly using Bayesian penalized splines within the hierarchical Bayesian framework. Consequently, the shapes of the predictor–response relationships are estimated during model fitting rather than being determined in advance through separate regression analyses.

The revised methodology also introduces a single unified Bayesian hierarchical structural model instead of a sequence of increasingly complex candidate models. All model parameters, including spline coefficients, interaction effects, random effects and hyperparameters, are estimated simultaneously from the data within a coherent probabilistic framework. As a result, the questions concerning the selection of regression functions, the interpretation of the exploratory figures, the definition of "raw" covariates, the number of seasonal harmonics, and the rationale for multiple intermediate models are no longer applicable.

l.254: the references do not appear correctly here (should be in parentheses).

We thank the reviewer for pointing out this formatting issue. The citation formatting has been corrected in the revised manuscript, and the references now appear in parentheses in accordance with the journal's citation style.

l.260: it is not clear so far why you need to define such sky conditions. Maybe add a sentence to tell why you use such a classification.

We sincerely thank the reviewer for this helpful comment. We agree that the motivation for introducing the sky-condition classification was not sufficiently justified in the original manuscript, and this classification was poorly used in the remainder of the manuscript. As part of the major revision, this classification has been completely removed. The revised methodology models the continuous effects of cloud cover and aerosol loading directly within the Bayesian hierarchical framework, eliminating the need for predefined sky-condition categories. Consequently, this comment is no longer applicable.

Section 3.1. Not clear what is the added value of this very descriptive paragraph, given that Fig. 2 contains already much information about the variables. In the text details are provided for some arbitrary stations, not all, with no specific motivation for the choice made. There does not seem to be critical results highlighted in that paragraph, I believe Table 4 could be self-sufficient. Saying that the minimum wind speed is 0 m s-1 and the minimum insolation nearly 0 W m-2 as well is not very useful.

We sincerely thank the reviewer for this valuable observation. We agree that the original descriptive statistics section was overly detailed and did not sufficiently emphasize the key characteristics relevant to the subsequent modelling. In the revised manuscript, rather than providing station-by-station descriptions, we now focus on the principal spatial and climatic patterns observed across the two study regions and highlight only those features that are directly relevant for interpreting the Bayesian hierarchical model. Redundant descriptions of obvious quantities (e.g., minimum wind speed or near-zero irradiance values) have been removed. The Results and Discussion section now focuses on scientifically meaningful findings.

l.269: Highest variability for this variable

We thank the reviewer for identifying this typographical error. This comment is no longer applicable to the revised manuscript.

l.291: I confess that elpd.loo is not very meaningful to me. I wonder if it would require more low level explanation. Likewise, what is the model "weight"?

Figure 4: what are ELPD and ELPD difference? And how is it meant to be interpreted?

l.301: MCMC has not been defined, neither the corresponding diagnostics.

l.304: I'm not sure I understand what parameter is the intercept.

Fig. 5: to what station correspond each color in the first row? Also, there are different lines. To what do they correspond?

l.309: does bi correspond to the first row (beta)? In that case this is misleading, as beta appears for each variable as well in Eq. 8. Also the variable b appears twice in Eq. 8 (at the beginning and for the temporal variations).

l.324: I'm not sure the prior distributions of all parameters have been introduced.

We sincerely thank the reviewer for these comments. We agree that the original validation section did not provide sufficient explanations of several Bayesian concepts and diagnostics, including the interpretation of the expected log predictive density (ELPD), model weights, Markov chain Monte Carlo (MCMC) diagnostics, posterior parameters and prior distributions. We also acknowledge that some figures lacked sufficient annotation to facilitate interpretation by readers unfamiliar with Bayesian inference.

As part of the major revision, the model validation section has been completely rewritten and substantially expanded. The previous comparison among several BNLME model variants has been replaced by a comprehensive Bayesian model evaluation framework. The revised manuscript now explicitly describes the purpose and interpretation of posterior predictive checks, Bayesian R^2, PSIS-LOO, WAIC, Pareto-k diagnostics, R, effective sample sizes, and Hamiltonian Monte Carlo convergence diagnostics. The role of each model parameter, including the intercept, random effects, spline coefficients, interaction terms and hyperparameters, is also clearly defined, and the prior distributions assigned to all estimated parameters are explicitly presented in the methodology. Figures have been redesigned with clearer annotations and captions to improve their interpretation.

Table 7: It is not clear to me whether there should be a set of "parameters distributions" for the whole country, or one for each station. In the latter what is the mean in Table 7? The mean of all the posterior distributions? It's hard to understand how a set of posterior distributions is treated and why it would be meaningful to give these statistics (compared to showing all the posterior distributions).

We sincerely thank the reviewer for this important observation. We agree that the presentation of the posterior summaries in the original manuscript could be confusing. As part of the major revision, the model formulation and the presentation of posterior results have been revised. The reported summaries (posterior means, credible intervals, effective sample sizes, and convergence diagnostics) refer to the posterior distributions of the estimated model parameters. The accompanying text has been expanded to clarify the interpretation of these posterior summaries and their role within the hierarchical Bayesian framework.

l.349-350: this point could have been clarified earlier

l.354-357: I'm not convinced by this link between surface temperature and cloudiness, and by the dissipation argument. Could it be reversed? Less clouds means more GHI hence higher temperatures? Likewise the impact of wind is very questionable. Wind could suspend dust and decrease GHI, and often wind is associated with cloudy conditions. You're looking at correlations hence it is very difficult to find causal effects, which could be one way or the other.

l.365: what do you mean by "attenuation effect of relative humidity"? The actual absorption by water vapor (to be analyzed in clear-sky conditions only) or the fact that the dependence is stronger wrt humidity (through the intermediate formation of clouds maybe)?

l.375: again, their might be strong covariance that could mask the first-order effect of pollution, for instance if pollution is correlated with wind.

We thank the reviewer for these comments. We agree that the original discussion occasionally implied causal interpretations that are not justified from observational data alone. Throughout the revised manuscript, the estimated coefficients are now consistently described as conditional associations within the hypothesized Bayesian hierarchical structural model rather than direct causal effects.

The revised discussion also explicitly acknowledges that several atmospheric predictors (e.g., temperature, humidity, cloud cover, aerosols and wind speed) are strongly interrelated. Consequently, the estimated effects should be interpreted as the association of each predictor with the clearness index after accounting for the other variables included in the model, rather than as isolated physical mechanisms.

l.384: it is the only place where sky conditions are used, and this is marginal. Would it be useful to show a Figure to illustrate this particular case?

We thank the reviewer for this helpful suggestion. In the revised manuscript, the sky-condition classification has been removed following a substantial revision of the study design. The proposed Bayesian framework is fitted on all-sky conditions and now directly models the continuous influence of cloud cover and aerosols on surface solar irradiance.

l391: indeed AOD would be more relevant, and I see no particular reason not to use such a product, for instance from AERUS-GEO (https://www.icare.univ-lille.fr/projects/user-driven projects/aerus-geo/).

We sincerely thank the reviewer for this valuable suggestion. We agree that Aerosol Optical Depth (AOD) is a more appropriate variable for investigating the influence of atmospheric aerosols on surface solar irradiance because it directly characterizes the aerosol burden in the atmospheric column responsible for the attenuation of incoming solar radiation. Following the reviewer's recommendation, the revised manuscript now uses AOD at 550 nm (AOD550) as the primary aerosol indicator. This variable replaces PM2.5, whose estimation in our previous analysis relied on derived products and was therefore subject to additional sources of uncertainty. We believe that this modification improves the physical consistency of the analysis and the interpretation of aerosol–radiation interactions.

The authors would like to express their sincere appreciation to the reviewer Quentin Libois for the time and effort he devoted to evaluating this manuscript. His constructive comments and valuable suggestions were instrumental in improving the quality of this work.