MS No.: egusphere-2025-4049

A machine learning approach to driver attribution of dissolved organic matter dynamics in two contrasting freshwater systems

Daniel Mercado-Bettín[a*], Ricardo Paíz[b,c], Valerie McCarthy[b], Eleanor Jennings[c], Elvira de Eyto[d], Angeles M. Gallegos[e], Mary Dillane[d], Juan C. Garcia[e], José J. Rodríguez[e], and Rafael Marcé[a]

[a]Centre for Advanced Studies of Blanes, Spanish National Research Council, Carrer Accés Cala Sant Francesc, 14, 17300 Blanes, Spain

[b]School of History and Geography, Dublin City University, D09 YT18, Dublin 9, Ireland

[c]Centre for Freshwater and Environmental Studies, Dundalk Institute of Technology, A91 K584 Dundalk, Co. Louth, Ireland

[d]Fisheries & Ecosystem Advisory Services, Marine Institute, F28 PF65 Newport, Co. Mayo, Ireland

[e]Ens d'Abastament d'Aigua Ter-Llobregat, Ctra. Aigües, 6, 08440, Cardedeu, Spain

* Corresponding author: Daniel Mercado-Bettín

We would like to thank the reviewers for their time in reviewing this manuscript. All the suggestions and comments have been carefully considered and have been taken into account in the revised version of the manuscript. Here, we describe actions that enhanced our work, incorporating their feedback. We appreciate their time and the valuable insights provided during the revision of the manuscript.

Note:

In bold text: the comments, observations and suggestions of the reviewers.

In normal text: the responses of the authors of this manuscript.

Response to reviewer 1: Thelma Panaïotis

Disclaimer: I have strong expertise in machine learning but my primary background is in oceanography, so I am less familiar with the specific impacts and conclusions relevant to lake ecosystems. Consequently, my review emphasizes the technical and methodological aspects of the manuscript and provides fewer comments on the ecological interpretation of the results.

The study by Mercado-Bettín et al. investigates how dissolved organic matter dynamics in two contrasting freshwater lakes can be modeled using a set of machine-learning algorithms. By assembling a comprehensive set of physicochemical and meteorological predictors, the authors train and compare several regression techniques (including linear models, random forests, support vector regression and three gradient-boosting frameworks) to predict fluorescent dissolved organic matter (fDOM) concentrations. Their analysis explores which environmental drivers most strongly influence fDOM variability and evaluates model performance across the two lake systems, aiming to demonstrate that data-driven approaches can capture the temporal patterns of organic matter turnover in inland waters. Furthermore, they show that reducing the set of predictors merely decreases the prediction performance of the ML model.

We thank the reviewer for the careful reading of the manuscript and for the clear and constructive summary of our work. The focus on the technical and methodological aspects of the study is appreciated, as the expertise in machine learning aligns closely with the objectives of the manuscript. Below, we address specific comments in detail.

General comments

The manuscript is clearly written and follows a logical progression, which makes the authors' objectives easy to grasp. Nonetheless, several critical steps are lacking to support the central claim.

We thank the reviewer for this observation and we address the missing steps outlined below in order to better support the central claim of the manuscript.

1. First, the manuscript does not show whether the target variable was inspected for distributional anomalies such as skewness, outliers, or zero-inflation before model fitting. A simple histogram or density plot of fDOM (and a note on any transformation applied) would let readers judge whether the data were appropriately conditioned.

We thank the reviewer for this valuable suggestion. We included an explicit exploratory data analysis of the target variable prior to model fitting, focusing on distributional properties, skewness, outliers, and potential zero-inflation. The corresponding code of this analysis is available in the public repository (codes/1_exploratory_data_analysis.R). The results were added to the Supplementary Material of the manuscript.

Results show that fDOM at one of the study sites (Feeagh) exhibits low to moderate skewness (approximately 0.5) with relatively few outliers, whereas the other study site (Sau) displays higher skewness (greater than 1 but below 2) with some extreme values. Zero-inflation does not appear to be an issue at either site.

Feeagh study site:

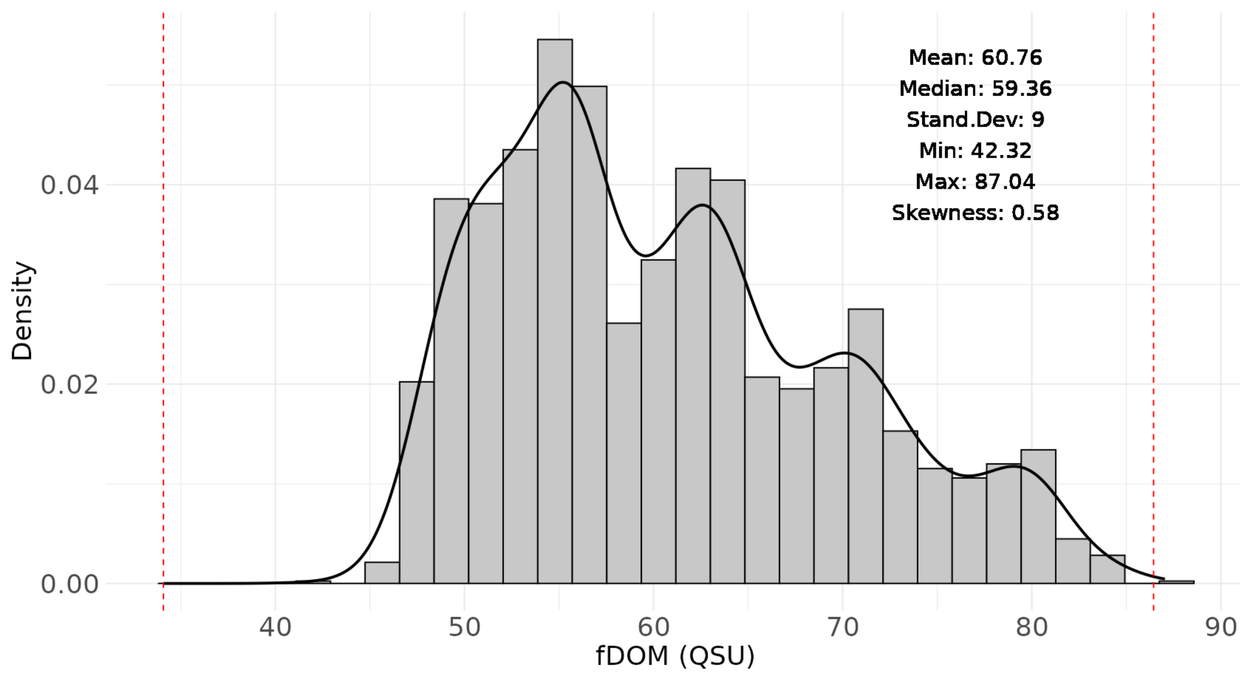

Sau study site:

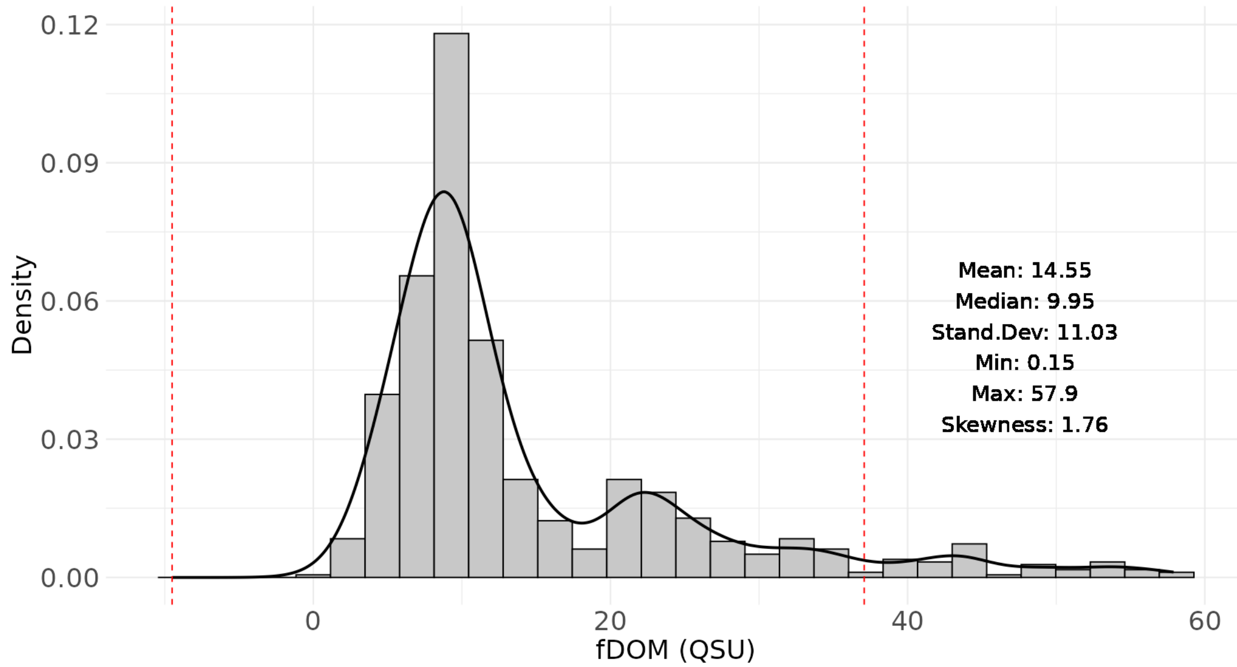

No transformation was applied to the fDOM data prior to model fitting in the main analysis. We agree with the reviewer that transformations can be beneficial for distance-based models such as k-Nearest Neighbors (KNN), kernel-based models such as Support Vector Regression (SVR), and linear models, particularly when strong skewness is present. In the case of Feeagh, the low skewness suggests that this is not critical, while for Sau, it may be more relevant, although skewness remains moderate rather than extreme. By contrast, the models that showed the best performance in this study, Random Forest (RF), XGBoost, LightGBM (LGB), and CatBoost (CTB), are generally robust to skewed distributions and can be used without data transformation (see, e.g., the "Decision Trees – Strengths" section in Molnar, 2025). IN any case, we added results, in supplementary information, for Sau with log10 transformation for comparison.

2. Second, the choice of six machine-learning models, including three gradient-boosting implementations, is left unexplained. Because these three models are largely interchangeable, the authors should either justify retaining each (for example, to compare computational efficiency or regularisation strategies) or reduce the set to a smaller, well-motivated collection, explicitly outlining the strengths and weaknesses of each algorithm and stating whether a broad model comparison is a declared aim of the study.

We thank the reviewer for this important point. We acknowledge that the rationale for including multiple machine-learning models, particularly several gradient-boosting implementations, was not sufficiently explained in the original manuscript.

Our primary objective was not to conduct an exhaustive benchmark comparison of machine-learning algorithms. Rather, given the complexity of fDOM dynamics and the relatively high number of heterogeneous input drivers, we adopted a pragmatic modelling strategy that explored supervised ML approaches with a diversity of potential drivers. Prior to the analysis, it was not clear which modelling family would be most suitable for capturing the nonlinear and site-specific behaviour of fDOM, particularly across two contrasting lake systems. We agree that the boosting methods used could be interchangeable, hence, in the new version of the main manuscript we only show one boosting method and moved the rest to supplementary, in addition to the rewriting of the whole text where we refer to boosting models instead of a single particular method.

In the manuscript, we now (i) explicitly state that a broad model benchmark is not the primary aim of the study, (ii) provide a single boosting method in the main manuscript (the one that performed best), and (iii) clarify that the workflow is intended to offer practical guidance for similar applications in lake modelling.

3. Third, hyper‑parameter tuning is mentioned but not described. The manuscript should specify which parameters were tuned for each model, the search space explored, the optimisation strategy (grid, random, Bayesian, etc.) and the validation split used. Detailing this process is essential for assessing model robustness and guarding against over‑fitting. An early subsection that summarises key performance metrics of the different models would also ground the subsequent discussion of variable importance in a known predictive skill.

We thank the reviewer for this very important point. We agree that a clearer description of the hyperparameter tuning procedure will improve transparency, reproducibility, and assessment of model robustness.

To address this, we have expanded the Methods section to explicitly describe, for each model, the hyperparameters that were tuned, the search space explored, the optimisation strategy applied, and the validation split used during tuning.

Regarding model performance metrics, we note that the metrics currently used in the manuscript were selected because they are commonly employed in lake modelling and water quality studies. However, we agree with the reviewer that an earlier and clearer contextualisation of these metrics across models would strengthen the manuscript. We added a subsection summarising key performance metrics for all models prior to the driver-attribution analysis, in order to ground the subsequent discussion of variable importance in terms of predictive skill.

4. All figures suffer from overly small text; increasing the font size for axes, legends, and annotations is essential for readability. In Figures 3 and 4 the legends contain interpretative statements, which blurs the line between description and analysis : legends should merely describe the visual content, leaving interpretation to the main text.

We thank the reviewer for this observation. Figure readability has been improved, increasing the font size of axes, legends, and annotations across all figures. In addition, we have revised the legends of all figures to remove interpretative statements, ensuring that legends are strictly descriptive and that all interpretation is confined to the main text. Please note that some figures are displayed at reduced size in the journal-formatted revision version. Full-resolution figures have been submitted separately and will be provided again for publication.

5. Finally, the code is shared in a reasonably structured way, but two key improvements would greatly enhance its usability. First, assigning a clear execution order, either by numbering the scripts or by providing a master driver script, would allow anyone to run the workflow sequentially without guesswork. Second, replacing absolute paths such as "~/Documents/intoDBP/driver_attribution_fdom/" with relative paths or a configurable settings file would make the repository portable across different machines and operating systems. Implementing these changes would substantially increase the rigor and reproducibility of the study.

We thank the reviewer for raising this important point, which is aligned with the FAIR principles (Findable, Accessible, Interoperable, and Reusable) and good practices for reproducible research. In response, we have added a clear execution order by numbering all scripts in the repository and replaced absolute paths with relative paths to improve portability across different machines and operating systems.

Specific comments

Figure 1: the caption should define DOC; otherwise the figure nicely illustrates the workflow

Thank you for noticing this. We think you refer to Figure 2 in the manuscript which describes the modelling workflow. The definition of DOC (dissolved organic carbon) has been incorporated into the figure.

L21-23: please indicate the direction of variation for each process (increase or decrease)

The direction of variation depends on multiple factors (land cover, climate, and topography) at the same time, hence they regulate the carbon inputs. We added a phase to help clarify this point.

L59: insert the word "respectively"

The word respectively has been added: *Lough Feeagh and Sau Reservoir are located in western Ireland (53° 56' N, 9° 35' W) and northeastern Spain (41° 58' N, 2°60 23' E), respectively (Fig. 1)*.

L93-96: you mention a 2 minutes measurement resolution but also give a number of points that roughly matches the number of days between the dates. Clarify whether the data were averaged daily before analysis.

We agree that this requires clarification. The high-frequency (2-minute) fDOM data were averaged to daily values before analysis, and this has been now explicitly stated in the revised manuscript: *Daily surface fDOM values were obtained from high-frequency data for both sites (2-minute-resolution data averaged to obtain daily values before analysis). For Feeagh, data were measured at 0.9 m depth, and for Sau, an average value was calculated between the measurements for depths 0-5 m.*

Figure 2: define the acronyms “GWLF” and “GLM” (they are only explained later, at lines 114-116). Also define “NSE”; you introduce NSE as a model-evaluation metric here but it does not appear elsewhere in the paper.

We have defined the acronyms GWLF (Generalised Watershed Loading Functions model) and GLM (General Lake Model) directly in the Figure 2 caption for clarity. In addition, NSE (Nash–Sutcliffe Efficiency) has been removed from the figure, as it is not used as a performance metric in this study.

L126-127: see also Molnar (2025) for interpretable machine learning.

We thank the reviewer for this suggestion. We have included a citation to Molnar (2025) at this point in the manuscript to enhance the statement on interpretability.

L128-135: the description of how the data were split into training and test sets, and how the models were evaluated, is unclear. Align this description with the "Prediction Workflow" paragraph and distinguish clearly between validation (hyperparameter optimisation) and testing (final performance).

This section has been revised and is now aligned with the Prediction workflow section, to clearly distinguish between training (used for hyperparameter tuning) and testing (used for final performance assessment), as also outlined in our response to General Comment 3.

L136: the term “statistical” is vague (all these models belong to supervised machine learning), and the choice of models needs justification. Explain why these particular algorithms were selected, why e.g. a neural network was not considered, and what the relative strengths and weaknesses of each method are. If three gradient-boosting frameworks (XGBoost, LightGBM, CatBoost) are used, state the reason for testing all three rather than picking one.

We thank the reviewer for this comment. We agree that the term “statistical” was not appropriate in this context and have replaced it throughout with “supervised machine-learning models.” The revised Methods section now provides a clearer justification for the selection of modelling approaches, including a brief discussion of the relative strengths and limitations of each method and the rationale for testing multiple gradient-boosting frameworks (XGBoost, LightGBM, and CatBoost). In addition, we now explicitly clarify why neural networks were not included: neural networks were not included in this study because they typically require more complex architecture design and calibration, which can increase methodological complexity and may reduce reproducibility. In contrast, the selected models allow more reproducible implementation due to their comparatively simpler hyperparameter tuning. These revisions are consistent with our response to General Comment 3.

L140-141: consider using permutation-importance as an alternative that works across all models

We thank the reviewer for this suggestion. We decide to use the model that perform the best to show the variable importance and moved the rest of model to supplementary for comparison.

L165: I would like to see a short "Model performance" subsection before the analysis of important drivers, so readers first see how well the models actually performed.

We thank the reviewer for this suggestion. We revised the Results section to improve the flow and consistency with the figures and overall narrative of the manuscript. Specifically, we first highlight the most influential drivers using the best-performing model (CatBoost, CTB), which is then described explicitly in the subsequent text. We adopted this structure because model performance is closely linked to the most influential drivers identified, and presenting results in this order we consider improves flow and coherence. The text has therefore been reorganised accordingly.

Figure 3: I'm confused by what you mean regarding the "four ML models that directly provide feature importance: Random Forest (RF), eXtreme Gradient Boosting (XGBoost), Light Gradient Boosting (LGB), CatBoost (CTB)". These models do not provide direct feature importance, you have to compute it (that you did with node purity and gain contribution. On the opposite, a linear model directly provides feature importance, through its parameters (see Molnar (2025)).

We thank the reviewer for this important clarification. We agree that the wording "directly provide feature importance" is inaccurate and has been removed. In the revised manuscript, we explicitly stated that feature importance is computed for these models using metrics such as node purity (Random Forest) and gain contribution (XGBoost, LightGBM, and CatBoost).

Table 1: RMSE should be reported with its units, since it shares the same unit as the predicted variable. The fact that XGBoost yields a very poor performance (R^2 = 11%) compared to others (including a linear model with R^2 = 45%) suggests that the training

could be largely improved. Which hyperparameters did you choose for the XGBoost? I dug in the code and found that only a few hundred trees (100-300) were used, which may be insufficient (but this also depends on the learning rate). Typically, boosting procedures require thousands of trees (see Hastie et al., 2009; Chapter 10). Moreover, the reported R^2 of 99 % on the training data in Table A1 strongly suggests overfitting.

We agree that RMSE units should be reported and we have added them in the revised manuscript. The reviewer is correct that the low number of trees used during XGBoost hyperparameter tuning led to poor performance, particularly in Feeagh. We have rerun the analysis using thousands of trees, and XGBoost performance is now comparable to the other boosting methods. The issue of overfitting, as indicated by the high training performance, is acknowledged as a limitation in the current discussion section; we now mentioned that this is a limitation in the revised manuscript.

Figure 4: mixing a partial‑dependence plot (PDP) for the Random Forest with SHAP values for CatBoost makes the interpretation confusing. Choose one model and interpret it thoroughly. Do not embed interpretation inside the legend. For the SHAP plot, indicate how the variables are ordered (e.g. by mean absolute SHAP value). Also, the cosine‑scaled Julian day axis should be accompanied by a note translating the cosine values back to calendar seasons, because "cos(julian day) = ‑0.5" is not intuitive.

We thank the reviewer for this comment. We agree that combining PDPs (Random Forest) and SHAP values (CatBoost) in the same main figure can make interpretation less clear. In the revised manuscript, Figure 4 now focuses on the best-performing model (CatBoost, CTB) and presents SHAP-based results only; the PDP analysis has been moved to the Supplementary Material.

We have also revised the caption to remain strictly descriptive and removed interpretative statements, in line with our response to General Comment 4. Finally, we have added a note translating the cosine-transformed Julian day values to approximate calendar seasons.

L198 - 208 & Figure 4: the text mentions an experiment using "the most influential drivers" and another using "a reduced subset of reanalysis‑based and easily accessible drivers." In the figure the purple line is labelled "Testing with all drivers," which appears to correspond to the reduced set described in the text. Either correct the label or clarify the distinction between the two experimental setups.

We thank the reviewer for pointing out this inconsistency. We have clarified the distinction between the two experimental setups in both the text and figure labels. The violet colour corresponds to simulations using only the most influential drivers, while the green colour corresponds to the reduced subset of reanalysis-based plus Julian day.

L249: discuss the implications of predicting fDOM instead of total DOC. Is the former a simpler target, and does that affect the ecological relevance of the results?

We thank the reviewer for this comment. We have added text to the Discussion section addressing the implications of predicting fDOM instead of DOC: *While both fDOM and DOC are widely used indicators of dissolved organic matter in lakes, they differ in measurement principles and in the fractions of organic matter they represent. Both remain, however, ecologically relevant for understanding DOM sources, transport, and transformation processes, which aligns with the context of this work*.

L273: good point, thank you for highlighting the limitation of dataset shift.

Appreciate the comment.

Figure A5: define all acronyms

The mentioned acronyms have been defined in the caption.

References:

Molnar, C. (2025). Interpretable Machine Learning: A Guide for Making Black Box Models Explainable (3rd ed.). christophm.github.io/interpretable-ml-book/

Reviewer 2: Anonymous Referee #2

This study incorporated multiple machine learning approaches to identify the key drivers among 24 environmental variables to predict the fDOM in two contrasting lentic systems. This is overall an interesting study and its methodologies may be applied for future research in similar fields. Please find my detailed comments below:

We thank the reviewer for the assessment. We appreciate the recognition of the applicability and reproducibility of our study, which are key objectives of this work. Below, we address the specific comments in detail.

Line 69: The word "also" is confusing. Do you mean both natural processes and human activities are important for Sau? If so, this sentence is not consistent with lines 67-68. If only human activity is the driver, "also" should be removed.

We have clarified in the new version of the manuscript that we intended to convey that both natural processes and human activities are important drivers at Sau, reflecting its more heavily intervened and generally more populated catchment. We have adjusted the wording accordingly and ensured consistency between the referenced lines to avoid confusion: *The dynamics of DOM in both study sites have been previously explored in Ryder et al. (2014), which identified that natural dynamics related to soil temperature, river discharge and drought were important drivers in Feeagh, and in Marcé et al. (2021), which showed human activities (e.g., wastewater effluents and agricultural runoff) were important for Sau, in addition to its environmental dynamics.*

Lines 84- 85 and Figure 2 step (4): It is unclear why the second simulation should only use the reanalysis data, since you already know that the human drivers are also important, at least for Sau. Could you please explain the study objective or hypothesis? Also, please double-check if you did include the Julian day for the second simulation. If it was included, a purple rectangle that represents "Others – Julian day" and relevant text should be added in step 4 model 2 because the last two lines in the figure caption said so.

We thank the reviewer for this comment. The rationale for using only reanalysis-based drivers (plus Julian day) in the second simulation is directly linked to one of the main objectives of the study, namely to evaluate the scalability and reproducibility of the proposed modelling workflow. Specifically, this modelling exercise was designed to test whether fDOM dynamics can be reasonably captured using only globally available data sources, without relying on site-specific or human-activity variables (e.g., in-situ measurements and site-specific modelling outputs) that are often difficult to obtain or unavailable in many regions .

We agree that human drivers can be also important at specific sites, such as Sau. However, our results show that a model trained exclusively with reanalysis-based variables can still capture a part of its DOM dynamics. This is because reanalysis variables, such as soil moisture and soil temperature, can implicitly reflect aspects of land use and human influence on catchment functioning (e.g. altered soil properties or hydrological responses in more urbanised or intensively managed catchments). As such, this second simulation is not intended to replace site-specific modelling or monitoring, but rather to demonstrate the utility of the workflow in data-limited contexts.

Regarding the reviewer's second point, we confirm that Julian day was included in the second simulation. We have revised Figure 2 and its caption accordingly by explicitly adding the "Others – Julian day" component to step (4), ensuring consistency between the figure and the text.

Lines 102-103: Was fDOM data for Lake Feeagh also corrected? Figures A2 and A3 only show the correction for Sau. Line 301: Please add the linear regression equation in Figure A2 and cite Figure A2 in A3.

We thank the reviewer for raising this point. The fDOM data for Feeagh were corrected for temperature quenching following established approaches described by Watras et al. (2011) and Ryder et al. (2012). To improve clarity, we have expanded the Supplementary Material to describe the temperature-correction procedures for both study sites, including a brief explanation of the correction approach, the equations and coefficients used, and their basis in the published literature. The relevant figures are now clearly referenced.

Lines 106-107: "All input 107 data variables, including their respective units and source, are displayed in Step 1 of Figure 2." should be moved to the beginning of this paragraph.

We agree that this sentence provides a useful starting context and have moved it to the beginning of the paragraph to improve clarity and flow.

Lines 217-218: Besides boosting biological activity, the increased soil moisture also indicated the stronger soil-lake hydrological connectivity. More terrestrial DOM can be exported into the lake during wet periods.

We thank the reviewer for highlighting this point. We agree that increased soil moisture can also indicate stronger soil–lake hydrological connectivity, particularly during flushing events triggered by high precipitation. We have added information in the Discussion emphasizing this: *In addition, higher soil moisture can reflect stronger soil–lake hydrological connectivity, particularly during wet periods and flushing events, facilitating the transport of terrestrially derived DOM into the lake*.

Lines 257: the indices in Figure A5 should be spelled out instead of abbreviations or acronyms.

The acronyms have been added in the new version

Figure 1: (1) For the "land cover" legend, I do not think "organic rich soil" is one of the land cover classes in CORINE, and it is merely used as a land cover metric. Please explain which classes were accounted for in it, or replace it with a commonly used category. Does it refer to peatland according to line 72? (2)"Regular seasonality" is a vague description in the fourth line of the figure caption. What does "regular" mean? Does it have distinguished seasonality every year or have no seasonal difference every year? Both patterns can be "regular".

We thank the reviewer for the careful reading of Figure 1. (1) Yes, that is correct - organic-rich soil is not a standard land-cover class in the CORINE classification. In both catchments, the overlap between the catchment boundaries and the CORINE Land Cover 2018 dataset resulted in a relatively large number of individual land-cover classes. To improve readability and facilitate comparison between sites, we grouped the original CORINE classes into a smaller number of categories (the five categories displayed in Figure 1) representing dominant catchment characteristics, as follows:

Lough Feeagh			
Code	Colour	CODE_18	CORINE Land Cover Classification
312		312	Coniferous forest
321		321	Natural grassland
322		322	Moors and heathland
324		324	Transitional woodland/shrubs
333		333	Sparsely vegetated areas
412		412	Peatbogs
511		511	Water courses
512		512	Water bodies

Land cover group	% of total area
Organic-rich soil LC	45%
Agriculture	1%
Grassland	27%
Urban	0%
Forest	25%
Watercourse	2%
	100%

Sau Reservoir			
Code	Colour	CODE_18	CORINE Land Cover Classification
111		111	Continuous urban fabric
112		112	Discontinuous urban fabric
121		121	Industrial or commercial units, public services and military installations
131		131	Mineral extraction sites
142		142	Green urban areas
211		211	Non-irrigated arable land
212		212	Permanently irrigated arable land *
231		231	Pastures
242		242	Complex cultivation patterns
243		243	Land principally occupied by agriculture
311		311	Broad-leaved forest
312		312	Coniferous forest
313		313	Mixed forest
321		321	Natural grassland
322		322	Moors and heathland
323		323	Sclerophyllous vegetation*
324		324	Transitional woodland/shrubs
332		332	Bare rock
333		333	Sparsely vegetated areas
511		511	Water courses
512		512	Water bodies

Land cover group	% of total area
Organic-rich soil LC	6%
Agriculture	18%
Grassland	12%
Urban	4%
Forest	59%
Watercourse	0%
	100%

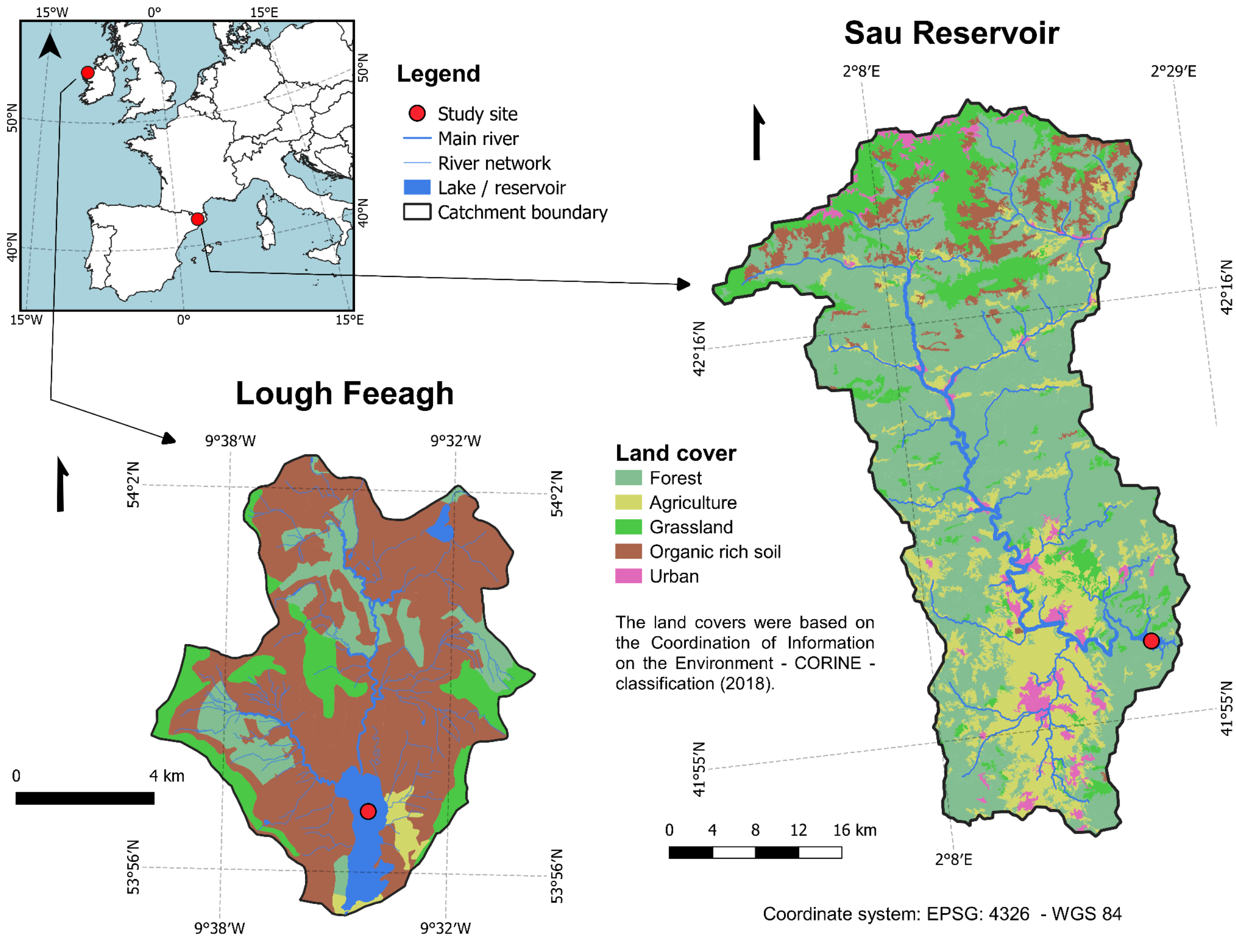

In this context, the category labelled organic-rich soils aggregates CORINE classes associated with high organic matter content, including peatbogs, moors and heathlands, and related classes. Similarly, CORINE classes such as Broad-leaved forest, Coniferous forest, and Mixed forest were grouped into a single category (Forest), and analogous groupings were applied to the remaining classes. This aggregation was used solely for visualisation and interpretative clarity in Figure 1 and is not intended as a replacement for the original CORINE classification.

To avoid ambiguity, we have added this information to the Supplementary Material explicitly listing the original CORINE land-cover classes identified in each catchment and showing how they were aggregated into the grouped categories used in the figure. We have also clarified this grouping and explicitly referred to it in the caption of the figure.

(2) We agree that the term regular seasonality was vague in this context. We have removed this wording and reduced the text in the caption to provide a more precise description of what is only included in the figure.

Figure 2: The abbreviations in step 1 should be explained in the figure caption, such as the ERA, GLM, and GWLF.

We thank the reviewer for this comment. All abbreviations used in Step 1 of Figure 2, including ERA5 RA 5 (the fifth-generation atmospheric reanalysis product produced by ECMWF), GLM (The General Lake Model), and GWLF (The Generalised Watershed Loading Functions Model), have been defined in the figure caption in the revised manuscript.

References:

Watras, C.J., Hanson, P.C., Stacy, T.L., Morrison, K.M., Mather, J., Hu, Y.H. and Milewski, P., 2011. A temperature compensation method for CDOM fluorescence sensors in freshwater. Limnology and Oceanography: Methods, 9(7), pp.296-301.

Ryder, E., Jennings, E., de Eyto, E., Dillane, M., NicAonghusa, C., Pierson, D.C., Moore, K., Rouen, M. and Poole, R., 2012. Temperature quenching of CDOM fluorescence sensors: temporal and spatial variability in the temperature response and a recommended temperature correction equation. Limnology and Oceanography: Methods, 10(12), pp.1004-1010.