Reviewer Report for "Global natural wetland methane emissions (2000-2025)" by Li et al.

General Comments

The manuscript presents a valuable machine-learning framework (XGBoost) to extend the Global Methane Budget (GMB, 2000-2020) through 2025. The low-latency tracking goal is highly relevant to the community, and the model demonstrates good skill across testing periods covering the years prior to and after the training window. Furthermore, the 2021-2025 extrapolation appears reasonable at the regional level. Overall, the manuscript is well-written and the core results are clear. However, some methodological details — particularly regarding feature selection and data processing — and certain visualization choices require further clarification and justification before publication. Below are my specific comments:

Specific Comments

Line 85 (Abstract): “...the most recent Global Methane Budget (GMB; Saunois et al., 2025)” Please explicitly specify here that this budget covers the 2000-2020 baseline period to establish the timeframe early for the reader.

Lines 220-235: The inclusion of 1-month and 2-month lagged versions of all 11 climate variables raises concerns regarding multicollinearity. The authors already acknowledged that the original 11 climate predictors are inherently multicollinear; introducing repeating variables with shifted time axes likely exacerbates this issue. Furthermore, the specific choice of 1- and 2-month lags is not adequately justified. Why were these chosen over temporal smoothing methods, such as seasonal means or running averages, which might better mitigate repeated data usage? Please provide a brief rationale for this specific lag selection and, if available (not requested but I think the authors might have already tested it), including a brief discussion on the statistical consequences of introducing these extra lagged covariates into the model.

Line 310: What exactly does "detrended monthly global emissions from GMB" entail? It is unclear how the emissions from the GMB were detrended and whether this implies the data was smoothed. Please explicitly describe the detrending and/or smoothing methodology applied here.

Line 335 (Section 3.1): The performance evaluation section transitions directly into grid-level details. Before diving into the spatial granularity, I recommend including an initial comparison of the modeled vs. GMB-estimated global (and/or broad latitudinal zone) annual emissions across both the training and testing periods. This macroscopic overview is essential for readers to assess the XGBoost model's performance at the aggregate scales most relevant to global budgeting analyses. While the authors do provide regional-level figures to support the 26-year trend discussion (Section 3.3), those figures focus primarily on temporal anomalies rather than explicit model performance evaluation against the GMB baseline.

Line 430 (Figure 4): The use of a logarithmic scale on the y-axis for regional emissions visually compresses the data, making it very difficult to distinguish meaningful differences between the GMB estimates and the model predictions. The authors may consider using a linear y-axis to allow for a clearer comparison of the bar heights.

Line 540 (Data Availability) & Conclusion: A question regarding the future of this work: Do the authors plan to use this specifically trained XGBoost model to provide low-latency natural methane emissions on a regular, operational basis, or is this paper primarily a demonstration of the emulator framework methodology? The conclusion currently reads like the latter. In either case, it would be highly beneficial for the community if the authors provided the trained, reusable XGBoost model. Additionally, briefly clarifying the plan for future operational updates (if any) in the conclusion would be nice.