

Authors Response to Referee #2

We thank referee #2 for all the time they invested in reading and understanding this paper and in particular, for their constructive comments and suggestions to the manuscript we have submitted. We took all the comments into consideration for generating a revised version of the manuscript. We respond to the original comments (italic) point-by-point.

Summary & General comments:

The MESMER-M-TP v0.1.0 module is designed to emulate spatially explicit monthly mean precipitation fields using monthly mean temperature fields as predictors. It offers runtime efficiency, which facilitates the efficient exploration of the precipitation forecast and its uncertainty under different scenario or parameters. Despite for some misses in extreme events, the emulator can effectively generate spatially coherent statistics of temperature and precipitation time series in general.

Overall, I enjoyed reading this paper very much. I appreciate that the emulator is carefully designed with physical knowledge and intuition. I also value the simplicity of the emulator, making it easy to understand and interpret. In terms of writing, the methodology is described in detail, the results are well-presented with sufficient explanation of the calculations. The significance and the implications are also discussed. Although I think the paper is already in a good shape, I do have a few minor comments below that I hope can further improve its quality and readiness.

AR: Thank you very much for this assessment. We found all your comments very helpful and appreciated the precise suggestions. We find that implementing them in the revised manuscript has overall increased the quality of our paper.

Specific comments:

L6: MESMER(-M): I suggest put the full name here

L94: row-vector: isn't it a column vector?

L103: the local precipitation... -> I suggest "the emulated local precipitation"

AR: All of them were very good catches and we have adjusted the manuscript accordingly.

A general comments for notations and equations: Please make sure all the notations are used and defined in a consistent way.

E.g., (1) Equations (3)-(6); (2) Equation (8); (3) Figures 7&8; (4)L154, (5)Equation(7)

AR: We are very sorry for the confusion regarding the notation. We did change the notation multiple times during the writing of this paper (for clarity). It seems that we accidentally mixed the different notations on multiple instances. We incorporated all your suggestions in our paper and carefully checked the rest of the manuscript for inconsistencies. We hope that the new set of equations is now concise. (L89ff)

L170-175: I suggest describing the methodology in more detail, perhaps including a few equations for clarity.

AR: We are describing the Kernel Density Estimation and the subsequent sampling approach in more detail now and to this end also included two equations (L165-182). We hope that the new description is clearer now.

L213: How sensitive are the results to the number of closest grid points n ? Does using PCA to define the predictor make the results less sensitive to the choice of n ?

AR: The results are rather robust with respect to the number of closest grid points. We have tested for n between 75 and 400. Larger n become too computationally expensive, while setting $n=2652$ results in a single PCA for the entire globe which strongly varies with region. The main reason for including the PCA is the strong spatial correlation of temperature, which leads to difficulties when fitting a Linear Model. We have tried to overcome these difficulties by using Lasso/Ridge/ElasticNet regression (rather than performing a Regression on the PCA transformed temperatures). Lasso Regression selects data from only a few grid-points to build a regression model, making the precipitation estimates overly sensitive to the temperature field, while Ridge Regression proved to be very sensitive to the training dataset during cross-validation. We found the solution relying on the PCA a good compromise. We have added some of this information to the manuscript (L219-227).

L218-219: I suggest including an equation to show the actual form of the prediction model, including the interaction terms used in this study.

AR: Great suggestion. We introduced equation (10) (~L242)

L243-249: My understanding is that, given the temperature field around a grid point, the emulator predicts a probability distribution of precipitation at the point rather than a single predicted value. Is that correct? If so,

AR: Your understanding is exactly right.

1. (1) *Do you actually draw random samples from the pdf to characterize the distribution? If so, how many samples are drawn?*

AR: We approximate the empirical distribution of the residuals using a Kernel Density Estimation (KDE) and then, we draw new samples from the fitted distribution to generate additional realizations of variability. We rely on a 1-on-1 coupling meaning we sample a single time series from the KDE for each gridded temperature field we use as forcing. This implies that when we force the emulator with gridded temperatures from an ESM, we draw as many realizations as we have ESM ensemble members available. For the three models we employ for validation that would be 30, 40 and 50 realizations, respectively. When forcing the emulator with emulated temperatures, we work with 100 gridded temperature fields per model and scenario, meaning we draw 100 time series from the KDE per model and scenario. We have added some information in the manuscript in L165-182 and L195-222.

(2) *How do you verify the EMU results (which is a distribution) against ESM (which is a single value)?*

AR: We verify this in multiple steps. During the fitting procedure, we compare the empirical distribution against the distribution fitted using the KDE and, for example, tune the bandwidth parameter of the KDE relying on k-fold cross-validation (L243-246). Once we have generated the full precipitation emulations (trend + variability) we compare the ensemble we generate against the entire available ESM ensemble. For example, CanESM5 has 50 ensemble members. We train our emulator on one ensemble member across SSPs. Next, we use the temperatures from the remaining 49 ensemble members as input for our emulator and pair the trend estimate from each of the ensemble members with one variability realization each, giving us 49 emulated precipitation fields. We then compare the statistics from the ESM precipitation ensemble to that of the emulated precipitation ensemble, both of which have 49 members (L208-L216 and L245-265). We focus on verifying the emulation approach as a whole, rather than verifying the individual modeling steps as we are interested in the performance of the emulation framework. The reasoning being that the GLM already captures large fractions of the precipitation signal (up to 80% of the variance of the original signal) and our focus is on validating the total variability (contribution from GLM + contribution from KDE).

(3) *What exactly does the quantile refer to in Figure 2? I assume for the ESM, the quantiles are calculated from the ensemble members (so sample size = # of ensemble member), but what about the EMU results (what are the samples size for the quantiles)?*

AR: Excellent question. We actually had not included information on emulated ensemble sizes in the manuscript. We have added information in L208-220. In short: the sample size for the ESM quantiles is always the number of ensemble members (see Table A1). The ensemble size for emulated results equals the ensemble size of the ESM for the results for the direct emulation (i.e. forcing MESMER-M-TP with temperatures from ESMs). When we emulate precipitation from emulated temperatures (Appendix C), we use 100 realizations.

L296: Fig.3 and B1 -> should be B2?

AR: Good catch!

A general suggestion for the figures: it might be helpful to add subtitle like (a),(b),(c) for subplots in some figures. This can make the figure captions easier to describe and the references to the figure in the text clearer.

AR: We agree and have adjusted where necessary.

Figure 2 - L4: Orange/Blue lines represent precipitation estimates of a single ESM/EMU ensemble member -> should be "dashed" lines to distinguish with the solid lines

AR: Another very good catch!

L300-302: This is an assumption in the modeling framework, is that correct? If so, I suggest change "suggests" -> "assumes"

AR: Adjusted accordingly.

L314-316: It is not immediately clear to me why this suggests that the model struggles to disentangle the trend and variability?

AR: We agree that the explanation is a bit unclear. What we meant to say is that it seems that the emulated division into trend and variability is not always accurate. A systematic underestimation of the median quantile suggests that our trend estimates are too low and the subsequent overestimation of the 10th quantile and underestimation of the 90th quantile then hints at the emulated data not having enough variability. We have adjusted the comment in the manuscript (L350-354)

We would again, like to thank referee #2 for their valuable comments. We have found them very helpful and tried to address them all as best as we could in the revised version of the manuscript.