We thank Anonymous Referee #1 (AR1) for their helpful feedback on our manuscript. AR1's comments are copied below in black with our responses to each in red.

Deep savanna roots:

The authors focus on deep roots for evergreen broadleaf forests, while the savanna roots are assumed to be shallow. In Singh et al. (2020) the presence of deep roots in savanna regions is highlighted. It is recommended to discuss how your analyses relate to this opposing result.

Singh, Chandrakant, et al. "Rootzone storage capacity reveals drought coping strategies along rainforest-savanna transitions." Environmental Research Letters 15.12 (2020): 124021.

Thank you for bringing up this important point. The findings of Singh et al. (2020) and similar works should be considered. The following has been added to the manuscript (L264-270) , Results section):

Note that deep-rooted woody vegetation has been shown to exist in savanna ecosystems (Canadell et al., 1996; Oliveira et al., 2005a; Singh et al., 2020). Most of the areas classified as savanna in our domain are a result of deforestation, with forest giving way to grass-dominant pastures with a shallower rooting depth (Gash and Nobre, 1997; Roberts et al., 2002; Von Randow et al., 2004; Piontekowski et al., 2019; Honey, 2023). Additionally, Noah-MP does not capture heterogeneity in growth form within a given vegetation class, making it impossible for us to account for the proportion of potentially deep-rooted woody vegetation in savanna. A single canopy height range is assumed for every grid point classified as savanna (minimum 0.1 m, maximum 10 m) and forest (minimum 8 m, maximum 20 m). This canopy height is considered in calculation of e_j.

Please let us know if you have any follow-up comments on this issue.

Comparison to GLEAM:

The authors compare their model results with the GLEAM evaporation product. It should be noted that the 'algorithms applied to satellite observations' (L233) include assumptions on roots and root water uptake. Given this, referring to GLEAM as 'observational estimates' (L372) could be misleading.

We agree with AR1 that we need to be careful about how we refer to the GLEAM product. 'Observational estimates' in L391 has been changed to 'observations-based estimates' and other places in the text with similar language were also changed.

Additionally, a related point referring to the limitations of GLEAM was made by AR2 and has been addressed (see the document with our responses to AR2).

The results presented in Sect. 3.2 and Fig. 9 could benefit from difference maps and/or a quantification in terms of for example correlation, because the visual comparison of the maps is not entirely intuitive. Moreover, a short sentence on how to interpret the Theil-Sen slope could help the reader in understanding Fig. 9; what does a + and – slope mean in terms of water dynamics?

Thank you for these suggestions. We agree that quantification of the differences between plots in Fig. 9 would strengthen our results. We chose to calculate the correlation coefficient between plots, which requires regridding of the Noah-MP model results and the GLEAM data to a common spatial grid. Unfortunately, due to computation issues that arose in the last few days, we were unable to obtain the correlation coefficients before the response deadline. However, I am happy to provide them as soon as the issue is resolved and include them in the final version of the manuscript.

L243 – L248 in Methods section 2.4 have been updated to include more detail on the meaning of the Theil-Sen slope:

We apply the Mann–Kendall test for monotonic trend (Mann, 1945; Kendall, 1948; Gilbert, 1987) to the mean time series of transpiration and ET between days 150 and 250 of the year (the height of the dry season) for each point in the domain. We then calculate the slope of the Mann–Kendall trend (known as the Theil-Sen estimator) for the mean time series for each point (Theil, 1950; Sen, 1968). The sign of the Theil-Sen estimator indicates the direction of the trend in the mean time series; a positive (negative) value is associated with generally increasing (decreasing) transpiration or ET. We compare spatial patterns of the Theil–Sen estimator between the SOIL and ROOT Noah-MP model runs and the gridded GLEAM evaporation estimates.

Conclusion:

The conclusion provides more a summary of the findings than a strong conclusion. It is recommended that the authors rename the section, or reframe the section into a stronger ending. Furthermore, the numbers mentioned in L451 cannot be found in the previous sections, indicating that this does not fit in the conclusion section. It is also advised to stronger emphasize the relevance and the potential of this approach for the climate model community, and the accuracy of climate predictions.

The numbers mentioned in L451 (now L458-459) are found earlier in section 3.1. However, we agree with AR1 that the conclusion should include less details of the study findings and more emphasis on the relevance of the study to the larger scientific community. We have reduced the detail included in the conclusion section considerably. Additionally, to improve the organization of the conclusion section, we have added subsection headings (‘Summary of findings’) and (‘Significance and future work’).

Elaboration on the significance of the study is done via addition of the following paragraph to the conclusion:

Noah-MP with DynaRoot enabled can be used to investigate a number of different science questions with wide-ranging implications. Given the role of plant trait diversity in resilience of the Amazon as studied by Sakschewski et al. (2016) and the identification of deep-rooting as a drought resilience strategy by Chen et al. (2024), DynaRoot could be used to study changes in forest resilience under deforestation scenarios. Moreover, given that moisture varies slowly in subsurface soils (Amenu et al., 2005), DynaRoot makes it possible to characterize the role of deep soil moisture memory in influencing surface moisture via transpiration in a coupled land-atmosphere framework. Such research has been alluded to in Niu et al. (2020) and Zanin (2021), and could have implications for predictability of atmospheric moisture on longer timescales. Dominguez et al. (2024) discuss two multidecadal convection-permitting simulations that were completed for the entire South American continent. In their analysis of these runs, Zilli et al. (2024) identified land-atmosphere coupling in CPMs as an outstanding area of investigation. This motivates potential future work that focuses on the role of fine-scale land surface characteristics—such as water table depth and vegetation traits (including rooting depth)—in simulating convection. DynaRoot would be applicable in such work, particularly in global convection-permitting simulations that have become a priority in the climate modeling community (Satoh et al., 2019; Caldwell et al., 2021; Feng et al., 2023).

*[***Specific comments:**

- L100: reference brackets
  - Reference brackets have been added.
- 1c,d: the colorbars of the soil texture and land cover types are lacking information. It is suggested that the authors either show only the classes that are present in the regions including labels, or all classes with all labels.
  - Figure 1 has been updated to include only the classes that are present in the region for both soil texture and land cover.
- L242: ‘uptake above 1m’ is a bit confusing, could be solved by mentioning that we talk about uptake from soil layers shallower than 1m.

  - ‘uptake above 1 m’ has been changed to ‘uptake shallower than 1 m’. This change has also been made to the caption and label of Fig. 2.
- L260: ‘water table depth is 2m’ would be more suitable
  - In the GW case (unmodified Noah-MP with the MMF groundwater scheme activated), direct interactions between resolved soil layers and the water table can occur when the water table is shallower than the cumulative depth of all soil layers (in this case, 2 m). Thus, we will leave this part of L260 (now L276) as is.
- 3g: the units of the axes are missing and it is not explicit that this is for the ROOT experiment or for another experiment.
  - Fig. 3g shows the root activity function (Eq. 9), which is unitless. This is now specified in the figure. The root activity function is part of the DynaRoot scheme, which is only active in the ROOT case. Thus, root activity output can only be from the ROOT case. The figure caption now specifies that the data in 3g comes from the ROOT case.
- L338: reference brackets
  - Reference brackets have been added.
- L427: remove the ‘is’
  - Removed.