Authors response

Title: Evaluation of ASCAT soil moisture retrievals and their potential to detect intraday variability
Authors: Lan Anh Dinh, Filipe Aires, and Victor Pellet
MS No.: egusphere-2026-2360
MS type: Research article

Dear referees,

We thank the two referees for their highly constructive comments and suggestions. Our responses (in blue) are listed below.

We hope that our adjustments will make the manuscript clearer and more pleasing to the referees and other readers.

Thank you & best regards,
Lan Anh & co-authors.

Referee #1:

The study presents a novel approach for retrieving ASCAT soil moisture via a convolutional neural network (CNN) that explicitly models spatial dependencies. The validation is robust, using both the ERA5 reanalysis and in situ soil moisture from the ISMN. Furthermore, the performance of the CNN-based ASCAT SM is compared with the H SAF ASCAT SM data record (derived using the change detection approach). The results are convincing, with improvements to the correlation coefficient for the CNN ASCAT SM relative to the H SAF ASCAT SM. However, the study would benefit from a clearer description of the machine learning architecture for reproducibility. Also, stratifying the ISMN validation results according to land cover and vegetation types would be informative.

Response: Thank you for your appreciation of our work. We respond to your two concerns below.

Major comments:

The description of the NN in Sections 3.1 and 3.2 should give more detail on hyperparameter tuning and overfitting prevention. Furthermore, a diagram of the steps would be beneficial for the reader.

Response: We agree that explicitly detailing these technical aspects improves the manuscript's transparency. We have revised the manuscript by including a flowchart describing the model, together with more details on hyperparameter tuning and overfitting prevention, as follows:

"We consider a localized CNN variant, incorporating a locally connected convolutional layer in which filters learn distinct weights for each spatial location. … . A detailed flowchart illustrating the proposed retrieval framework is provided in Figure 1.

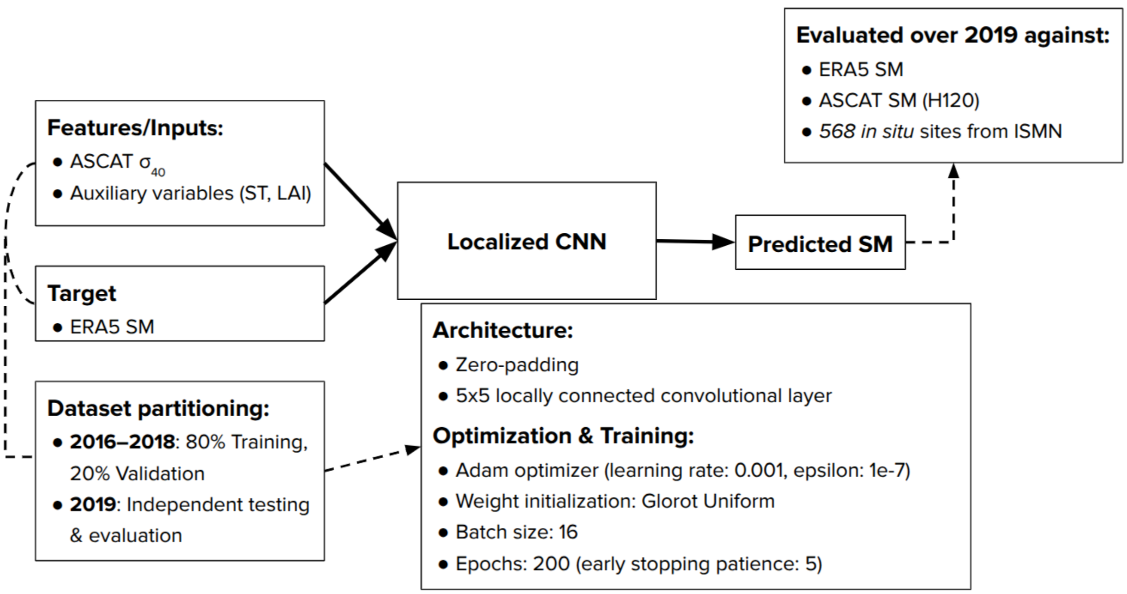

Figure 1. Workflow framework for ASCAT soil moisture (SM) retrieval, detailing the model inputs, localized CNN architecture, training parameters, and evaluation strategy.

…

Because the CONUS-scale dataset rendered computationally intensive automated optimization impractical (Rabiei et al., 2025), model hyperparameters were selected using a structured manual tuning approach. Consistent with methodologies established in recent remote sensing studies (Dinh, 2026), parameters were initially anchored to standard literature benchmarks (Rabiei et al., 2025; Pellet et al., 2025) and systematically refined against validation metrics until further adjustments yielded negligible improvements. Ultimately, our model is trained using the Adam optimizer (Kingma and Ba, 2017) with an initial learning rate of 0.001 and an epsilon value of 1 e−7 for numerical stability. Weight initialization is controlled using the Glorot Uniform initialization with a fixed random seed to ensure reproducibility. A batch size of 16 is used in the training. In addition, separate models are trained for ascending and descending satellite passes.

To robustly evaluate the model and actively prevent overfitting, we employed a strict temporal data splitting strategy alongside dynamic training constraints. The dataset spanning 2016 to 2018 is partitioned into 80% for training and 20% for validation, while the 2019 data is completely held out for independent testing and performance evaluation. During the training phase, overfitting is further mitigated through the implementation of early stopping; while the model is allowed to train for a maximum of 200 epochs, the process is halted early if the validation loss demonstrates no improvement over 5 consecutive epochs, at which point the optimal weights are restored.*”*

The validation against in situ data in Section 4.2 does not give any information about the land cover and vegetation types for the different networks, and how this affects the relative performance of the SM datasets. From Figure 2 (b), it seems that the CNN soil moisture is particularly advantageous in the mountainous regions in the western US. On the other hand, H120 seems to perform slightly better in the central Great Plains, which is more agricultural.

Response: Thank you for this insightful observation. We agree with your interpretation of the regional patterns. In the revised manuscript, we included a new performance assessment stratified by land cover and vegetation type for each ISMN site, based on land-cover information from the Annual National Land Cover Database (NLCD) product (U.S. Geological Survey (USGS), 2024). This additional stratification will enable a quantitative evaluation of how land cover and vegetation types influence regional differences in model performance.

“2.4 Other datasets

Land-cover data

To stratify retrieval performance across environmental conditions, we extracted the dominant land-cover type for each ISMN site using the 2019 National Land Cover Database (NLCD; U.S. Geological Survey (USGS) (2024)). This enables a targeted evaluation across major CONUS land-cover categories.

4.2 Evaluation against in situ measurements

…

Evaluating performance across major CONUS land-cover classes (Figure 4) confirms trends seen in Table 1 and Figure 3: while H120-o achieves lower biases across all classes—an artifact of local calibration (Aires et al., 2021b)—it yields higher standard deviations. Regionally, H120-o shows marginally better correlations than CNN-lo for croplands, reflecting results over the central Great Plains in Figure 3(b). Conversely, CNN-lo excels over complex topography (e.g., the mountainous western US), outperforming H120-o in correlation over evergreen and deciduous forests due to its ability to leverage spatial context (Pellet et al., 2025). Despite being trained on ERA5, CNN-lo outperforms its training target over crops and grasslands in terms of correlation.

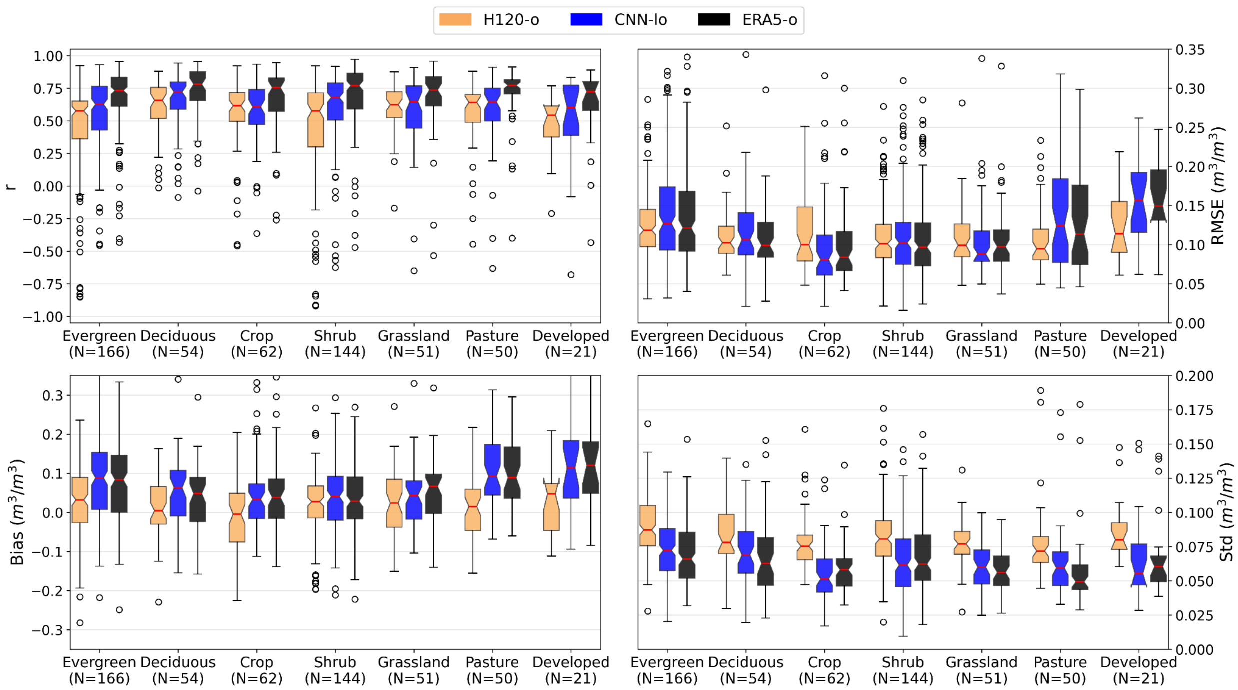

Figure. Boxplots of soil moisture (SM) retrieval performance (ASCAT H120-o, CNN-lo, and ERA5-o) compared to ISMN measurements. Results are grouped by major land-cover categories, including forests (evergreen, deciduous), crops, shrubs, grasslands, pasture, and developed areas. The number of sites within each class is shown in parentheses; underrepresented land-cover types (e.g., mixed forest, wetlands, barren, open water) are omitted.”

Minor comments:

There are many minor grammatical mistakes, and the paper would benefit from a thorough proofread.

Response: We will carefully proofread the entire manuscript and correct the grammatical errors to improve readability.

The authors should discuss the degree of dependence on ERA5 (training target vs validation independence)

Response: We thank you for highlighting this important point. Regarding the training dependence, our model is indeed heavily dependent on ERA5, as we used ERA5 soil moisture (SM) as the target. This choice is deliberate and scientifically justified. ERA5 currently provides one of the most robust large-scale SM datasets, resulting from the joint assimilation of multiple satellite observations, atmospheric forcing, and in situ measurements into a physically based land surface model, alongside bias-correction procedures (Soci et al., 2024). Using such a comprehensive reanalysis as a reference for supervised learning is a well-established practice in remote sensing (Aires et al., 2005; Rodríguez-Fernández et al., 2015; Pellet et al., 2025).

To ensure an independent evaluation, we validated the model against *in situ* observations from the ISMN, which were not used during training. Although the model was trained to reproduce the ERA5 SM target, validation against independent *in situ* measurements provides a robust assessment of its ability to retrieve physically meaningful SM information beyond the ERA5 training space. The fact that the retrieval shows comparable, and in some locations improved, agreement with *in situ* observations relative to ERA5 indicates that the model can extract additional information from the ASCAT observations beyond simply reproducing the ERA5 background.

In the revised manuscript, we have added a dedicated paragraph at the end of Section 2.2 and updated Section 2.3 to explicitly clarify this dependence and the critical role of our independent validation.

"**2.2 ERA5 database**

…

The use of ERA5-derived ST and LAI as auxiliary inputs alongside ASCAT backscatter to predict SM introduces some dependence on the ERA5 framework. However, this design is both physically and methodologically justified. Physically, ST and LAI provide crucial boundary conditions that help the model separate SM effects from vegetation attenuation and thermodynamic variability in the ASCAT signal. Methodologically, they ensure the retrieval remains consistent with the ERA5 background space, a standard requirement for data assimilation (Rodríguez-Fernández et al., 2019). While the product inherits ERA5's large-scale climatology, its sub-daily and event-scale variations are driven by the independent ASCAT observations, capturing dynamic information absent from the ancillary variables alone.

2.3 In situ data from the International Soil Moisture Network

We used in situ SM data from the ISMN … . This ISMN dataset was completely withheld from training, which enables a robust, independent evaluation of whether the retrieval effectively leverages ASCAT observations to generate physically meaningful SM estimates beyond the ERA5 background.

"

In Section 1, the authors should define what the surface soil moisture represents (i.e., the top few centimetres of soil)

Response: We agree with the reviewer. We have added this definition into the manuscript: "Surface soil moisture (i.e., water in the top few centimeters of soil, hereafter referred to as SM) represents a relatively small component of the hydrological cycle (Trenberth et al., 2007) … "

Line 139: It seems this sentence should start with a bullet point like the previous two bullet points; Line 168: The second ypred should be yref; Line 233: The red point is shown in Figure 4 (a), not Figure 4 (b); Line 252: section 55.1 typo

Response: Thank you; these errors have been corrected in the revised manuscript.

Line 269: ERA5-24h should be defined before the abbreviation is first used

Response: Yes, this has been explained earlier in the original manuscript, on line 229: "For reference, the full 24-hour SM amplitude of ERA5 (ERA5-24h)."

References:

- Rabiei S, Babaeian E, Grunwald S (2025). Deep learning-based short- and mid-term surface and subsurface soil moisture projections from remote sensing and digital soil maps. Remote Sens. 17, 3219. https://doi.org/10.3390/rs17183219
- Pellet V, Aires F, Boucher E, Volden E (2025). Enhancing soil moisture statistical retrieval from SMOS using partial convolutions and localization strategies. J. Appl. Meteorol. Climatol. 64, 1951–1965. https://doi.org/10.1175/JAMC-D-25-0041.1
- Dinh LA (2026). ASCAT soil moisture retrieval using deep learning: A focus on localization strategy. Front. Remote Sens. 6, 1718353. https://doi.org/10.3389/frsen.2025.1718353
- U.S. Geological Survey (2024). Annual NLCD Collection 1 Science Products. https://doi.org/10.5066/P94UXNTS
- Soci C, Hersbach H, Simmons A, Poli P, Bell B, Berrisford P, et al. (2024). The ERA5 global reanalysis from 1940 to 2022. Q. J. R. Meteorol. Soc. 150, 4014–4048. https://doi.org/10.1002/qj.4803
- Aires F, Prigent C, Rossow WB (2005). Sensitivity of satellite microwave and infrared observations to soil moisture at a global scale: 2. Global statistical relationships. J. Geophys. Res. Atmos. 110. https://doi.org/10.1029/2004JD005094
- Rodriguez-Fernandez N, Aires F, Richaume P, Kerr Y H, Prigent C, Kolassa J, et al. (2015). Soil moisture retrieval using neural networks: Application to SMOS. IEEE Transactions on Geoscience and Remote Sensing 53, 5991–6007. https://doi.org/10.1109/TGRS.2015.2430845
- Rodriguez-Fernandez N, de Rosnay P, Albergel C, Richaume P, Aires F, Prigent C, et al. (2019). SMOS neural network soil moisture data assimilation in a land surface model and atmospheric impact. Remote Sens. 11. https://doi.org/10.3390/rs11111334

Referee #2:

General comments:

This is an interesting paper demonstrating the potential of machine learning to enhance existing soil moisture products derived from EUMETSAT's ASCAT instruments on MetOp satellites. A neural network (NN) is trained using ERA5 simulations as a reference. The resulting soil moisture retrievals are less noisy than the original ASCAT soil moisture product. In addition, the authors provide some evidence suggesting that sub-daily soil moisture variability could be observed through NN retrievals. While the paper is well written, the addition of a diagram presenting the NN training and retrievals with inputs and outputs would be useful.

Response: Thank you for your appreciation of our work. As also mentioned in the response to Referee #1, we have expanded the model description and added a detailed diagram describing the CNN architecture and processing steps in the revised manuscript.

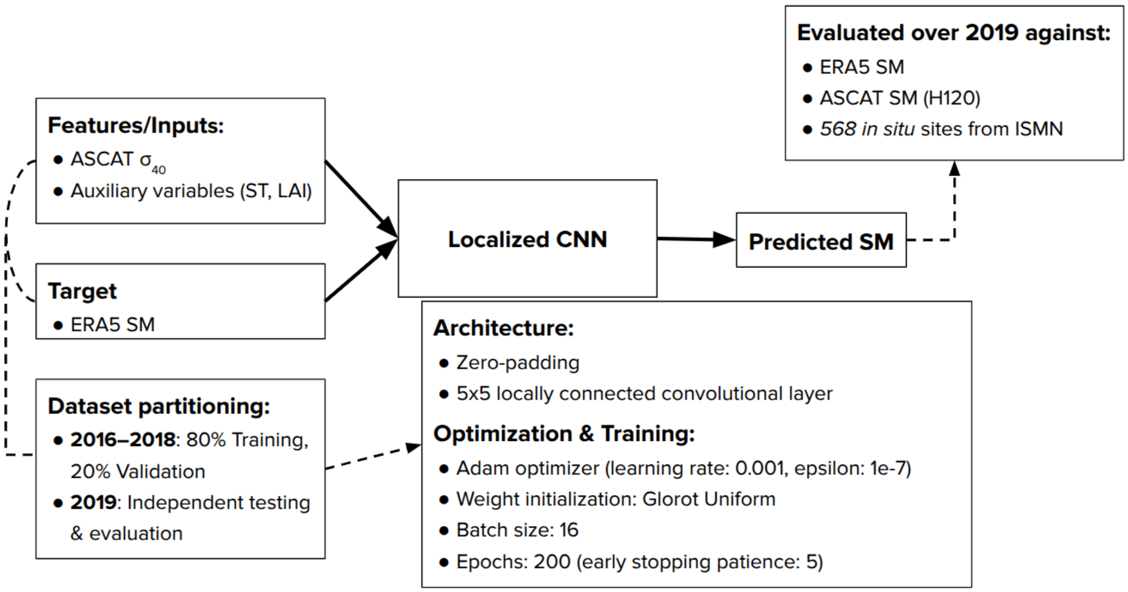

Figure 1. Workflow framework for ASCAT soil moisture (SM) retrieval, detailing the model inputs, localized CNN architecture, training parameters, and evaluation strategy.

Also missing is a table summarising mean score values per experiment with respect to in situ observations. In addition to scores covering the entire four-year period, the table should also indicate seasonal scores (e.g. DJF, MAM, JJA and SON across years).

Response: Thank you for this helpful suggestion. We agree that summarizing the performance metrics in a table greatly improves the clarity of our results, and we have added a new table (Table 1) to the revised manuscript.

As a brief clarification regarding the evaluation timeframe: our experimental setup partitions the four-year dataset by using three years (2016, 2017, 2018) for model training and validation, and reserving one independent year (2019) exclusively for testing and evaluation. Consequently, the performance metrics are computed only over this single year (2019) rather than the full four-year period. Given this one-year window, we focused the table on the annual overall scores rather than splitting this single year into seasonal subsets.

The corresponding text and Table 1 have been added to the manuscript as follows:
"Table 1. Performance metrics (median values across 568 sites) for soil moisture (SM) retrievals (ASCAT H120-o, CNN-lo, and ERA5-o) evaluated against ISMN *in situ* data for 2019. Metrics include temporal correlation (r), bias, and standard deviation (Std).

SM estimates		r	Bias ($m^3\ m^{-3}$)	Std ($m^3\ m^{-3}$)
ASCAT retrievals	H120-o	0.59	0.024	0.080
	CNN-lo	0.65	0.059	0.064
ERA5-o		0.75	0.056	0.061

"

Finally, the approach is somewhat circular since ERA5-derived variables (ST and LAI) are used to predict soil moisture, as well as ASCAT backscatters. This shortcoming should be clearly acknowledged, and the reasons for it should be made clear.

Response:
We agree that the use of ERA5 variables for both the target (SM) and inputs (ST and LAI) must be clearly acknowledged and justified in the text.

Physically, microwave backscatter is highly sensitive to vegetation attenuation and the thermodynamic state of the land surface. Integrating ERA5 ST and LAI provides the necessary boundary conditions that help the model separate SM effects from other sources of variability in the ASCAT signal. Methodologically, using ERA5 ancillary variables ensures the retrieved ASCAT product remains fully consistent with the ERA5 background space, which is a standard requirement for data assimilation preparatory frameworks (e.g., Rodríguez-Fernández et al., 2019).

We acknowledge that the retrieved product inherently adopts the large-scale climatology and background state of ERA5. However, the sub-daily and event-scale temporal variations are primarily constrained by the independent ASCAT observations, providing dynamic information not contained in the ERA5 ancillary variables alone. As also noted in our previous response to Referee 1, this is evidenced by the model's performance against independent ISMN stations.

We have revised Section 2.2 of the manuscript to explicitly acknowledge this design choice and to clarify the physical and methodological rationale for the chosen framework.

"**2.2 ERA5 database**

…

The use of ERA5-derived ST and LAI as auxiliary inputs alongside ASCAT backscatter to predict SM introduces some dependence on the ERA5 framework. However, this design is both physically and methodologically justified. Physically, ST and LAI provide crucial boundary conditions that help the model separate SM effects from vegetation attenuation and thermodynamic variability in the ASCAT signal. Methodologically, they ensure the retrieval remains consistent with the ERA5 background space, a standard requirement for data assimilation (Rodríguez-Fernández et al., 2019). While the product inherits ERA5's large-scale statistics, its sub-daily and event-scale variations are driven by the independent ASCAT observations, capturing dynamic information absent from the ancillary variables alone."

Particular comments:
- L. 41: "SM estimates at sub-daily resolution" is not clear. Do you mean that several ASCAT observations are available for a given location on the same day? Or does the time of observation vary from one day to the next?

Response: We are referring to the availability of multiple ASCAT observations for a given location on the same day (corresponding to native overpass times), rather than daily aggregated soil moisture. This distinction is now added in the revised manuscript:

"Notably, the Metop ASCAT CDR (H SAF, 2021a), which we use later in this study for comparison, already provides SM estimates at sub-daily resolution. Specifically, due to the orbital geometry of the ascending and descending passes, combined with the use of multiple ASCAT instruments (e.g., Metop-A and -B), a given location can be observed multiple times at different local overpass times on the same day."

- L. 44: ASCAT observations are global. Why focus on CONUS? Any reason for that?

Response: Here we selected the CONUS as it offers a highly dense network of in situ SM measurements, which are essential for robust validation. Also, CONUS encompasses pronounced spatial and climatic variability, making it an ideal testbed (Bernhardt et al., 2018).

This is now clarified in the revised manuscript:
"This study aims to explore the potential of deep learning to improve sub-daily SM from ASCAT observations over the contiguous United States (CONUS). Although ASCAT provides global coverage, CONUS serves as an ideal study region because it features pronounced spatial and climatic variability (Bernhardt et al., 2018) and hosts a dense network of *in situ* measurements essential for robust validation. "

- L. 102: Does "the ERA5 LAI product" exist? As far as I am aware, ERA5 neither simulates dynamic LAI nor integrates LAI observations. This sentence needs to be clarified.

Response: Thank you for pointing this out. We fully agree; ERA5 does not simulate dynamic LAI nor assimilate LAI observations, but rather uses a prescribed monthly climatology (Duveiller et al., 2023). Our original phrasing ("the ERA5 LAI product") was inadvertently misleading, as it implied a dynamically varying dataset.
However, it is important to note here that our rationale for using this specific climatological background, rather than dynamic external satellite datasets, was to ensure strict physical

consistency with our training target. Feeding the model dynamic vegetation anomalies that the ERA5 soil moisture target inherently does not possess would introduce an input-target mismatch. We have made this point clearer in the revised manuscript:

"To maintain consistency with the ECMWF assimilation framework, we use the LAI forcing from ERA5 rather than externally sourced satellite datasets (Myneni et al., 2015; Yan et al., 2024). While ERA5 does not simulate dynamic LAI or assimilate LAI observations—relying instead on a prescribed monthly climatology (Duveiller et al., 2023)—using this specific background was necessary to ensure strict physical consistency with our training target."

- L. 225: Replace "extracting a reliable diurnal signal is a true challenge!" by "extracting a reliable signal is challenging".

Response: Thank you; the change has been made.

- L. 232 (Fig. 4): ERA5 precipitation data can be subject to significant biases. Could you include in situ precipitation observations in Fig. 4b?

Response: We thank you for these valuable suggestions. The primary purpose of using ERA5 precipitation (as in Fig. 4a) was to illustrate the broad spatial pattern of a precipitation event over the CONUS. However, we completely agree that including an independent observational precipitation record would strengthen the interpretation of the case study presented in Fig. 4b.

In the revised manuscript, we have integrated the high-resolution Multi-Radar Multi-Sensor (MRMS) gauge-corrected precipitation (Zhang et al., 2016) product to provide a robust, independent observational reference. To match the spatial resolution of our SM retrievals, the MRMS data were aggregated to a 0.25° grid. Although there are differences in the absolute precipitation magnitude and the precise timing of individual peaks between the MRMS and ERA5 datasets, they exhibit consistent overall temporal trends. Importantly, the broader intra-day precipitation dynamics in both datasets correspond with the observed SM responses.

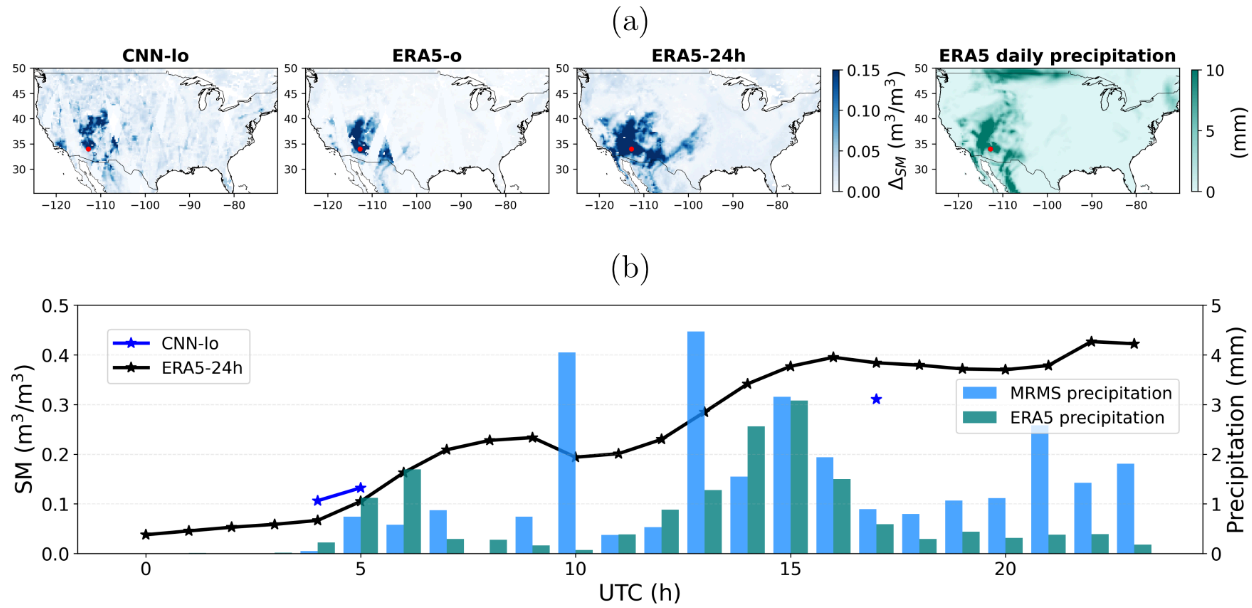

Figure. (a) Maps of daily SM amplitude (ΔSM) on November 20th, 2019, derived from CNN-lo and ERA5-o. ERA5-24h (full 24-hour amplitude) and daily accumulated precipitation are also shown for reference. (b) Hourly SM time series at [34◦N, 112.75◦W] (marked with a red dot in the maps), comparing CNN-lo retrievals to ERA5-24h. CNN-lo estimates correspond to three ASCAT overpasses: two ascending (04:00–05:00 UTC) and one descending (17:00 UTC). Precipitation data from both the independent MRMS gauge-corrected product and ERA5 are also presented.

- L. 240: This study does not actually demonstrate the intraday capability of the method. In situ observations of precipitation would provide a more convincing demonstration.

Response: We agree that the original wording in Line 240 overstated the conclusions that can be drawn from a single case study. We have revised the text to clarify that this example is intended specifically to illustrate the model’s behavior in capturing short-term SM dynamics during a heavy precipitation event, rather than to serve as a comprehensive demonstration of intraday capability. By incorporating the independent MRMS dataset, we provide an objective baseline to confirm that the timing of the retrieved SM response is physically consistent with actual observed rainfall, and not merely an artifact of the ERA5 forcing used during training.

“While a single case study cannot serve as a formal demonstration of the model's overall capability, this example illustrates the model’s ability to capture intraday SM dynamics during heavy precipitation events, despite sparse ASCAT observations. Crucially, the alignment with independent MRMS precipitation data confirms that the retrieved SM response is physically consistent with actual observed rainfall, rather than merely reflecting the ERA5 forcing. Capturing these dynamics is critical for hydrological applications, as rapid soil saturation can provoke flood risks. However, the spatial depiction of intraday variability remains strongly constrained by ASCAT’s sampling frequency. Integrating observations from multiple instruments (e.g., SMOS, SMAP) is expected to enhance both retrieval accuracy and temporal resolution in future work.”

- *L. 252: "section 55.1"?*

Response: We referred to section 5.1 here, and it has been corrected in the revised manuscript.

References:

- Marinescu P J, Abdi D, Hilburn K, Jankov I, Lin L (2024). An Evaluation of NOAA Modeled and In Situ Soil Moisture Values and Variability across the Continental United States. Weather and Forecasting, 39(3), 523-540. https://doi.org/10.1175/WAF-D-23-0136.1
- Bernhardt J, Carleton A M, LaMagna C (2018). A comparison of daily temperature-averaging methods: spatial variability and recent change for the conus. J. Clim.31, 979–996. https://doi.org/10.1175/JCLI-D-17-0089.1
- Duveiller G, Pickering M, Muñoz-Sabater J, Caporaso L, Boussetta S, Balsamo G, Cescatti A (2023). Getting the leaves right matters for estimating temperature extremes, Geosci. Model Dev., 16, 7357–7373, https://doi.org/10.5194/gmd-16-7357-2023
- Soci C, Hersbach H, Simmons A, Poli P, Bell B, Berrisford P, et al. (2024). The ERA5 global reanalysis from 1940 to 2022. Q. J. R. Meteorol. Soc. 150, 4014–4048. https://doi.org/10.1002/qj.4803
- Aires F, Prigent C, Rossow WB (2005). Sensitivity of satellite microwave and infrared observations to soil moisture at a global scale: 2. Global statistical relationships. J. Geophys. Res. Atmos. 110. https://doi.org/10.1029/2004JD005094
- Rodriguez-Fernandez N, Aires F, Richaume P, Kerr Y H, Prigent C, Kolassa J, et al. (2015). Soil moisture retrieval using neural networks: Application to SMOS. IEEE Transactions on Geoscience and Remote Sensing 53, 5991–6007. https://doi.org/10.1109/TGRS.2015.2430845
- Rodriguez-Fernandez N, de Rosnay P, Albergel C, Richaume P, Aires F, Prigent C, et al. (2019). SMOS neural network soil moisture data assimilation in a land surface model and atmospheric impact. Remote Sens. 11. https://doi.org/10.3390/rs11111334
- Pellet V, Aires F, Boucher E, Volden E (2025). Enhancing soil moisture statistical retrieval from SMOS using partial convolutions and localization strategies. J. Appl. Meteorol. Climatol. 64, 1951–1965. https://doi.org/10.1175/JAMC-D-25-0041.1
- Zhang, J., and Coauthors, 2016: Multi-Radar Multi-Sensor (MRMS) Quantitative Precipitation Estimation: Initial Operating Capabilities. Bull. Amer. Meteor. Soc., 97, 621–638, https://doi.org/10.1175/BAMS-D-14-00174.1