Response to Reviewer #2:

A deep-learning algorithm is presented that forecasts Boundary-Layer clouds based on morning-soundings and surface fluxes. The network is trained and validated with observational data of ARM's souther great plains (SGP) site. The skill of the new model is analyzed in terms of cloud triggering, their vertical structure and cloud fraction. An attribution of forecast errors is undertaken based on the network model that illustrates major factors influencing the representation of boundary-layer clouds in the context of this model.

The authors present a technically involved analysis using a broad scope of observational and modeling data which are synthesized into a deep-learning neural-network model to represent boundary-layer clouds over the ARM's Souther Great Plains (SGP) site. The analysis is novel and suits the scope of GMD well. In some aspects, the quality of the manuscript does, however, not hold up to the standards of the journal. This holds for (i) the methodological description (which in the present form would not allow to reproduce the findings), (ii) use of English Language and (iii) contextualization of the work in the scope of existing models (mixed-layer models) for the problem considered. I therefore recommend to the editors to reconsider the manuscript for review after the comments below have been addressed.

Response: We appreciate the reviewer's constructive and comprehensive feedback on our study, which are very helpful for improving the clarity and quality of the manuscript. Specifically, we have provided a clearer and more complete description of the machine-learning procedure, improved the overall readability of the manuscript, and discussed the advantages and limitations of our approach compared to mixed-layer models. All the comments and concerns raised by the referee have been carefully considered and incorporated into this revision. Our detailed responses to the reviewer's questions and comments are listed below.

Cross Review: I agree with the comments raised by anonymous Reviewer #1, in particular I want to subscribe to his first general demanding a cleaner and more complete description of the machine-learning procedure.

Response: We appreciate the reviewer's concurrence with the comments raised by anonymous Reviewer #1 and their emphasis on the need for a cleaner and more complete description of the machine-learning procedure. We have addressed this concern by providing a detailed and comprehensive description in Section 3.1, as outlined below.

Major remarks:

1. *The methodological description of the deep-learning model is rather obscure. Even after carefully reading the manuscript, it remains unclear what exactly is input to the model and what is the output? Does the model forecast an entire day or just a single instance? What is given in terms of the surface parameters – the evolution of fluxes up to the moment of forecast? Or the value at this time? How exactly are the trigger, vertical position and horizontal cloud fraction components of the model related? Shouldn't the vertical position and horizontal fraction constrain each other from the perspective of available humidity? Or does one trigger the other – if so in what direction: does non-zero cloud fraction trigger the vertical positioning or vice versa?*

Response: We are grateful for the reviewer's constructive comments on the need for a clearer methodological description. We have revised the manuscript to provide a more detailed

explanation of the inputs and outputs of the deep-learning model, as well as the relationships between its components. In general, we use hourly input and output at the same time period. Trigger serves as the input for the estimations of cloud positions and fraction, as these parameters are only meaningful under the cloudy conditions. In addition, a non-zero cloud fraction is not necessarily associated with a specific vertical position, as the cloud fraction indicates the proportion of cloud within the predicted cloud boundaries rather than at a fixed vertical grid. Specifically, we have included the following description and Table 1 of input and output lists in the revised Section 3.1, as follows:

"The model is purpose-built to consist of three distinct deep learning modules, each responsible for a critical aspect of the cloud simulation: 1) the determination of the BLC occurrence, 2) the height position of the cloud base, and 3) the cloud thickness and the normalized 10-layered shape of cloud fraction within cloud boundaries, which jointly yield the hourly-averaged vertical structures of BLCs. This modular approach ensures that the estimations are specific for each aspect of the BLCs. Combining cloud thickness and cloud fraction in one module is logical because the thickness for 10-layered clouds varies based on cloud thickness, and thickness is potentially related to the fraction, as thicker clouds are sometimes associated with larger cloud fractions. This separation of tasks enhances the overall reliability and clarity of the model in capturing the various characteristics of BLCs. Note that each of the three deep learning modules is built upon a deep neural network (DNN) with multiple hidden layers.

The occurrence module, as the first step, evaluates the likelihood of cloud formation by producing a number between 0 and 1, which we call "trigger" in the following, whose value above 0.5 indicates the presence of clouds. The target data for this module is binary (0 or 1), and the model output is a continuous value between 0 and 1. This occurrence information then feeds into the other two modules in parallel: one for locating cloud boundaries and the other for delineating the vertical shape of the cloud fraction in cloudy layers. While the cloud-base (or boundary) module and the fraction-thickness (or fraction) module are independent of each other, they collaborate to depict the vertical cloud fraction profile.

To represent the vertical structure of BLC in the fraction-thickness module, we segmented the cloud layer from the base to the top into ten levels, with each level's thickness varying according to the overall cloud thickness. These values are then interpolated to create a continuous vertical profile of cloud fraction within the BLC boundaries, offering a detailed depiction of the cloud's vertical extent. The vertical position of the layer changes based on the predicted cloud base and top to accurately represent the vertical structure of BLCs. This dynamic approach allows the fraction module to adjust and focus on the relevant portions of cloud fraction within cloudy layers. Compared to a static level approach, which requires the prediction of cloud fraction across a fixed vertical extent (e.g., multiple levels between 0-6 km), our method focuses on the shape of the fraction profile. This ensures the model is not constrained by fixed vertical levels, allowing for more efficient and robust estimations."

Table 1: Detailed descriptions of input and output variables used in the deep learning models for predicting boundary layer clouds (BLCs). The table includes the variable names, descriptions, and data sources. For the input parameters, surface meteorology and fluxes are taken from the current and previous hours, while morning profiles comprises 46 values spanning from 0-8 km at 06 LT. Note that the output data is derived from ARSCL (Active Remote Sensing of Clouds). The three

outputs correspond to the trigger module, cloud-base module, and fraction-thickness module, respectively.

Variable	Description	Data Source
Input		
Month	Range from 1-12	Time Record
LT	Local Time	Time Record
PS	Pressure at surface level (2m)	Surface Meteorology Station
RH	Relative Humidity at 2m	Surface Meteorology Station
U	Zonal wind at 2m	Surface Meteorology Station
V	Meridional wind at 2m	Surface Meteorology Station
T	Temperature at 2m	Surface Meteorology Station
LCL	Lifted Condensation Level	Derived from T, RH, PS
SH	Sensible Heat	Energy Balance Bowen Ratio
LH	Latent Heat	Energy Balance Bowen Ratio
RH Profile	Morning RH profiles	Radiosonde
U Profile	Morning U wind profiles	Radiosonde
V Profile	Morning V wind profiles	Radiosonde
θ Profile	Morning potential temperature profiles	Radiosonde
BLH_{SH}	PBLH derived from sensible heat	Derived from θ Profile and SH
BLH_{Parcel}	PBLH derived from parcel method	Derived from θ Profile and T
Output		
Trigger	Cloud occurrence	ARSCL
Position	Cloud-base height	ARSCL
Fraction Profiles	Cloud fraction and thickness	ARSCL

2. *Data Representativity should be assessed and discussed. The entire analysis focuses on a single site. This comes with two major restrictions that need careful consideration: (1) the current findings are constrained to the surface configuration of the SGP site; as such the insight to cloud—surface coupling may be substantially limited. (2) When comparing to reanalysis data, it needs to be taken into account that the reanalysis is representative for a large area – in comparison to the local nature of the DNN model and observation data. What is the local heterogeneity of cloud fields in the area covered by an ERA5 or MERRA grid cell? In other words – to what extent is the current method to be understood us a downscaling rather than an improved parameterization. To a lesser, but possibly non-negligible (?), extent, this also applies to time-locality.*

Response: We appreciate the reviewer's comments on the need to address data representativity and the focus on a single site. We have added a discussion in the manuscript acknowledging that our findings are specific to the SGP site and the need for further studies in different surface environments and synoptic patterns to generalize the findings. Meanwhile, the limitations of the study have been discussed, and we have outlined future work as follows:

" Moving forward, future work is warranted to test and extend this diagnostic tool to different synoptic patterns over a large region, which can be integrated into multiple-scale

models or reanalysis data. However, several challenges need to be addressed to achieve this. One significant limitation is the lack of high-quality, detailed observations of clouds and radiosonde profiles globally. This scarcity of data can hinder the model's ability to generalize effectively across different regions. To overcome this, there are several potential strategies. First, using transfer learning techniques can help adapt the model trained in one region to other regions with limited data. Integrating data from global observational networks (i.e., ARM) can also create a more diverse and representative training dataset, capturing a wider range of atmospheric conditions and cloud characteristics. Meanwhile, leveraging satellite data can provide broader coverage and enhance the robustness of the model. We plan to explore these approaches in future work to enhance the model's performance and applicability on a global scale."

Meanwhile, we acknowledge the local heterogeneity of cloud fields in the area covered by an ERA5 or MERRA grid cell in the revised Section 2.3. This inherent discrepancy between the model and observations arises from the difference between point-based and grid-based measurements. However, we believe it is not a significant issue for this study. The point observation at the SGP site, as a plain region, is frequently used as a benchmark to evaluate cloud coverage and fraction for reanalysis data or climate models, which typically have grid sizes ranging from tens of kilometers to over 100 km (e.g., Song et al., 2014; Zhao et al., 2017; Zheng et al., 2023; Zhang et al., 2017). Thus, we believe it is safe to compare our model to reanalysis data, including ERA-5 with a 0.25-degree grid (~30 km) and MERRA-2 with a 2/3° × 1/2° degree grid. For example, Figure R1 demonstrates the local heterogeneity of cloud fields within the grid cells. The four sites around the central site are located at distances of approximately 30-50 km. Their total cloud fractions are very similar, with differences ranging from 0-7% in terms of climatology. In summary, we believe it is suitable to use the point-based observations to represent the grid-based reanalysis data over this region, similar to other numerous studies. It is also important to note that for meteorological sites in plain regions, the diagnostic of clouds from models may largely differ from observations, which is an inherent limitation that applies to most models.

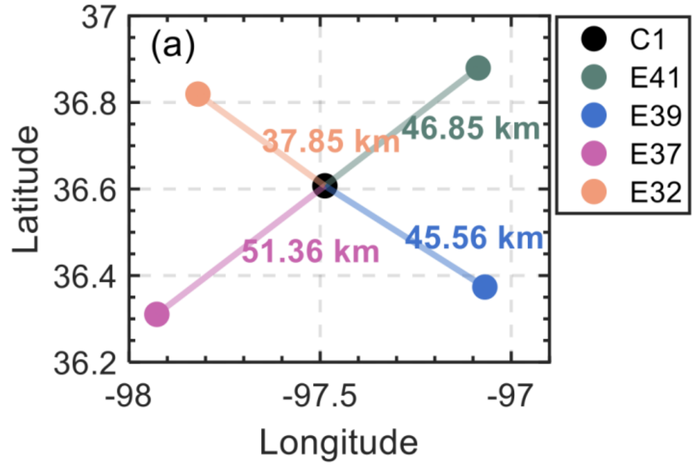

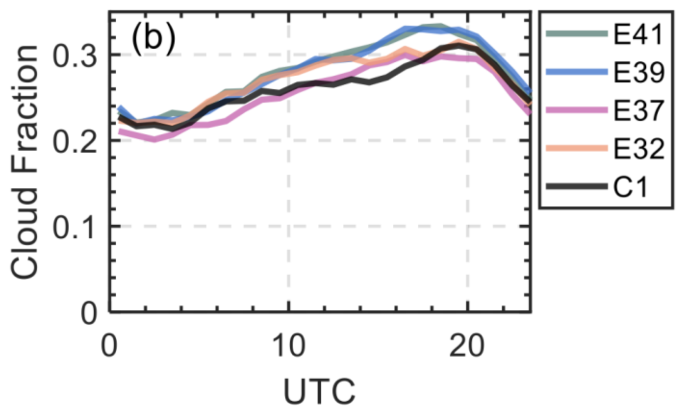

Figure R1. (a) Locations of the five sites, including the central site (C1) and four surrounding sites (E32, E37, E39, E41), with distances between C1 and the other sites indicated. (b) Total cloud fraction for these five sites derived from Doppler lidar, showing variations over time in UTC. The cloud fraction is averaged for the period where all five sites have available data (2016-2023).

Reference:

Song, H., Lin, W., Lin, Y., Wolf, A. B., Donner, L. J., Del Genio, A. D., ... & Liu, Y. (2014). Evaluation of cloud fraction simulated by seven SCMs against the ARM observations at the SGP site. Journal of climate, 27(17), 6698-6719.

Zhao, W., Marchand, R., & Fu, Q. (2017). The diurnal cycle of clouds and precipitation at the ARM SGP site: An atmospheric state-based analysis and error decomposition of a multiscale modeling framework simulation. Journal of Geophysical Research: Atmospheres, 122(24), 13-387.

Zheng, X., Tao, C., Zhang, C., Xie, S., Zhang, Y., Xi, B. and Dong, X., 2023. Assessment of CMIP5 and CMIP6 AMIP simulated clouds and surface shortwave radiation using ARM observations over different climate regions. Journal of Climate, 36(24), pp.8475-8495.

Zhang, L., Dong, X., Kennedy, A., Xi, B. and Li, Z., 2017. Evaluation of NASA GISS post-CMIP5 single column model simulated clouds and precipitation using ARM Southern Great Plains observations. Advances in Atmospheric Sciences, 34, pp.306-320.

3. *Embedment in large-scale models would break the physical consistency of the reanalyzed state. While I agree that there is merit in the (local) DNN cloud representation vs. the large-scale*

reanalysis, I doubt that the "representation" of clouds can be improved by simply using the DNN output in the context of the reanalyses. We should not forget that the reanalysis produces a heavily constrained, physically consistent approximation with the observations. Simply changing the cloud representation would thus, most likely deteriorate many other parameters as it would break the consistency.

Response: We appreciate the reviewer's insightful comment regarding the potential challenges of embedding the DNN cloud representation within reanalysis data. Models or reanalysis data are integral and complex systems. Any modification to a specific parameter may improve that parameter itself but could potentially worsen the performance of others. This is a common issue, not limited to deep learning, but applicable to any notable modifications to models.

In our study, we propose using the DNN as a diagnostic tool for BLCs from both observations and reanalysis data. Regarding the scale issue. As demonstrated in Figure R1, we believe that the local-scale observations in the SGP are representative of reanalysis grids. Thus, the DNN model can improve cloud representation over this site. However, further tests are warranted for its wide application across different sites and regions.

To address the specific concern about breaking the physical consistency of the reanalyzed state, we have revised the manuscript to clarify that our DNN model is intended to complement rather than directly replace existing reanalysis cloud representations. Specifically, we suggest that the DNN outputs can be used in a diagnostic capacity to identify BLCs and help to understand deficiencies in representing BLCs of reanalysis data. This approach allows for targeted improvements in cloud representation without compromising the overall physical consistency of the reanalysis data. Thus, DNN model can provide valuable insights into the local characteristics of BLCs, which can be used to inform the development of more accurate cloud parameterizations in large-scale models for this region. By integrating these insights in a controlled and incremental manner, it is possible to enhance cloud representation while preserving the integrity of these models. These discussions have been incorporated into the revised Section 4.3.

4. *Relation to mixed-layer (single-column) models – what is the added benefit of the rather complex DNN approach in comparison to simple low-order mixed-layer models which also capture the daily evolution of the boundary layer given an initial state and time series of surface fluxes? From a fundamental viewpoint, these models have a number of advantages vs. the DNN as they are (i) a loat cheaper to run, (ii) contain physical reasoning, (iii) are available for many of the large-scale models and thus do not need to be implemented.*

Response: Thanks for pointing out the comparison between the DNN approach and mixed-layer models for representing BLCs. We agree that mixed-layer models have advantages for simulating boundary layer, including lower cost, inherent physical reasoning, and ease of integration with large-scale models (Pelly and Belcher, 2001; Clayson and Chen, 2002; De Roode et al., 2014).

However, our DNN approach offers several distinct benefits that complement these traditional models. Firstly, the DNN can capture complex, nonlinear relationships between various different parameters that may be difficult for several equation to represent accurately. Meanwhile, the DNN model is trained on a large dataset of observational data, allowing it to learn from real-world examples and potentially identify patterns and

relationships that are not explicitly encoded in physical models. This data-driven approach can help improve the accuracy of cloud representation by leveraging the wealth of observational information available. The capability of DNN model is particularly valuable in situations where the interactions between variables are highly complex and not well understood. The advanced capability of the DNN model is demonstrated in this study, showing strong performance in estimating cloud occurrence, position, and fraction.

On the other hand, while mixed-layer models are based on established physical principles and are designed for estimating the boundary layer, they do not address vertical cloud fraction in the same scope as the proposed DNN model. In addition, DNN models generally offer advantages in computational speed compared to physical models. For example, our DNN approach is cost-effective, capable of producing two decades of BLCs with vertical information over the SGP within 1 minute using a single GPU node.

Finally, we propose using the DNN model alongside traditional physical models, not as a replacement. The current DNN model cannot produce detailed cloud microphysics and turbulence information. A hybrid approach can combine the strengths of both methods, leading to a comprehensive and accurate representation of BLCs.

We have revised the manuscript to highlight these points and clarify the complementary role of the DNN approach in the broader context of atmospheric modeling.

Reference:

Clayson, C.A. and Chen, A., 2002. Sensitivity of a coupled single-column model in the tropics to treatment of the interfacial parameterizations. Journal of climate, 15(14), pp.1805-1831.

De Roode, S.R., Siebesma, A.P., Dal Gesso, S., Jonker, H.J., Schalkwijk, J. and Sival, J., 2014. A mixed-layer model study of the stratocumulus response to changes in large-scale conditions. Journal of Advances in Modeling Earth Systems, 6(4), pp.1256-1270.

Pelly, J.L. and Belcher, S.E., 2001. A mixed-layer model of the well-mixed stratocumulus-topped boundary layer. Boundary-layer meteorology, 100, pp.171-187.

5. *Style and use of English language. The manuscript is full of syntactical and grammatical errors, which partly obscures the scientific contents. It needs to be carefully checked by a language editor before it may be re-considered for publication. I will list some recurring errors in my technical comments, but this list is by no means complete.*

Response: We appreciate the reviewer's feedback regarding the English language in the manuscript. We acknowledge that syntactical and grammatical errors can obscure the scientific content and hinder comprehension. We have reviewed the specific errors highlighted in the technical comments and corrected them accordingly. In addition, we have thoroughly revised the manuscript to check grammatical errors and syntactical issues. We believe these revisions have significantly improved the quality of the manuscript, which meets the high standards expected for publication.

Minor remarks:

l. 86: 'comprehensive data' – what kind of data? Please use a more telling attribute!

Response: Thank you for pointing this out. We have explicitly described the datasets in the introduction, including radiosondes (SONDE) profiles, surface meteorological measurements, and Active Remote Sensing of Clouds (ARSCL), from the Atmospheric Radiation Measurement (ARM) program at the Southern Great Plains (SGP) site.

l.89 'By assimilating morning radiosonde observations' – the procedure of deep learning is not really an assimilation (please check throughout the manuscript!)

Response: We appreciate your suggestion. We have replaced "assimilating" with "integrating" throughout the manuscript to accurately reflect the process of using morning radiosonde observations in our deep learning model.

l. 91 ' […] uniquely positioned to unravel the complex initiation […]' Even after careful assessment of the manuscript, I do not agree on this statement: First, the model is not uniquely positioned as other models exist that can cope with the processes in question (mixed-layer models, LES, etc.); Second, I do not agree that it unravels the initiation and evolution (which would correspond to a causal attribution.)

Response: We have revised this statement to acknowledge that the model is not uniquely positioned, as other models also address these processes. The text now reads: "this deep learning model is capable of diagnosing the initiation and evolution of low clouds".

Paragraph lines 82-92. At the end of this paragraph, it remains unclear what is training, input and output data for the DNN. This should be clarified here, at least qualitatively.

Response: Following this comment, we have added a detailed descriptions for the training, input, and output data for the DNN. In particular, the training data for the deep learning model includes SONDE profiles, surface meteorology, and fluxes, while the outputs are the estimations of cloud occurrence, position, and fraction. More detailed description can be found in the revised Section 2.1 and Section 3.1.

Section 2: "Data and instruments" – the section title does not reflect the contents; it also includes the cloud detection algorithm and a regime classification.

Response: We have revised the section title to "Data, Instruments, and Methodology" to more accurately reflect the contents.

l. 149 Why does CBH need to align with the LIDAR-detected PBL top? A BLC can also be initiated far below the PBL top…

Response: Indeed, we added a clear explanation for the criterion, as follows: "Coupled clouds are identified when the cloud base height (CBH), as derived from the ceilometer, aligns with or is below the lidar-detected PBL top height within 0.2 km, and the calculated surface-based Lifting Condensation Level (LCL, Romps 2017) falls within a maximum allowable range of 0.7 km (Su et al. 2022)."

l. 157/8 "typically lasting more than three hours". I suppose, this is a threshold for automatic detection, so is it three hours, or not? ("Typically " is rather confusing here…)

Response: We have revised this statement for clarity. It now reads: "lasting more than three hours," removing the word "typically" to avoid confusion.

l. 153-166 The classification is unclear to me. Is it done per day or per situation? For the stratiform cases, there is a three-hour threshold, but for the cumulus cases, there is no threshold in terms of duration (other than that the clouds emerge after local sunrise). So, how is it possible to characterize days then – these criteria could be evaluated separately for any situation, and certainly there are days in which regime shifts occur.

Response: We appreciate the reviewer's query regarding the classification process. The classification is indeed done per day rather than per individual situation. We recognize that there are situations where transitions between cloud regimes occur, such as the well-known stratiform to cumulus transition. However, in our study, we excluded mixed days and focused solely on days that exhibit purely stratiform or cumulus regimes. This approach ensures that our analysis is clear and specific to each cloud regime without the complexity introduced by transition periods. The revised text now explains the classification criteria more clearly.

Regarding normalization: The target data (cloud trigger, cloud vertical structure and cloud fraction) is already normalized (binary [0,1] , binary vector with elements [0,1] , real vector with real elements from the range [0,1]). Why does normalization need to be applied here? In what sense would it help?

Response: Although normalization is not mandatory, it is a common process used in numerous studies (e.g., Klambauer et al. 2017; Salimans and Kingma, 2016). We apply normalization to both the input and output layers, making them have similar magnitude. This practice offers several benefits: it improves convergence by ensuring features are on a similar scale, stabilizes the training process by preventing issues like exploding or vanishing gradients, and enhances model performance by enabling more consistent weight updates during training. While it is feasible to train without normalization, normalization generally leads to more efficient and effective learning.

Reference:

Klambauer, G., Unterthiner, T., Mayr, A., & Hochreiter, S. (2017). Self-normalizing neural networks. Advances in neural information processing systems, 30.

Salimans, T., & Kingma, D. P. (2016). Weight normalization: A simple reparameterization to accelerate training of deep neural networks. Advances in neural information processing systems, 29.

Fig. 1 need to be improved and is incosistent with part of the main text. The RH profiles / morning SONDE is duplicated information and should be within the box entitled "input". Are "profiles" the same as "soundings"? Or is there a difference? LT and MONTH should be combined to the azimuth angle as this is the actual physical parameter that matters and just gets encoded by Month and LT. What is the difference between surface meteorology and surface fluxes? In my understanding, surface fluxes are part of the surface meteorological parameters… The vertical alignment reflecting the geometry of the boundary layer is not appropriate here and confuses. A schematic focusing on input and output would help. If the models for cloud position, cloud fraction and trigger are independent entities, they should be reflected by separate hidden layers. The tag "cloud position" is potentially confusing here; I suggest to rephrase as "cloud structure" or "vertical position" as position alone might be mistaken for horizontal position. The relation between cloud trigger, cloud position and cloud fraction should be made more clear; If I understand correctly, the trigger is part of the input for the cloud position and cloud fraction models, but this is not appropriately reflected in the schematic. What is meant by time indicators? Are the meteorological parameters / fluxes input as instantaneous values or daily time series? Correspondingly: is the output produced per time instant or rather per day?

Response: We have revised Figure 1 to address the reviewer's comments and improve consistency with the main text.

- **The input box now includes the morning SONDE, and it is clarified that "morning SONDE" refers to the morning meteorological profiles.**
- **While we agree that using the azimuth angle makes physical sense, local time and month already imply this information, so there is no need to add an extra input parameter.**
- **Surface meteorology typically includes wind, temperature, pressure, and humidity. Although fluxes may be considered part of meteorology, we separate them here for clarity. The detailed input parameters are listed in Table 1.**
- **The idea is to use deep learning as a hidden layer to directly estimate BLCs from surface conditions, rather than resolving the complex PBL processes. Thus, we have marked the PBL in the diagram to indicate that the DNN model bypasses the turbulent and complex PBL processes to directly obtain BLCs from surface conditions.**
- **We have now indicated that there are three separate hidden layers, sharing a similar structure, to reflect the independent entities for cloud position, cloud fraction, and trigger.**
- **The term "cloud position" has been revised to "cloud base" to avoid confusion.**
- **The relationship between the models has been clarified: the trigger value serves as input for the cloud position and cloud fraction models.**
- **The caption now explains that the time indicators refer to local time and month.**
- **It is clarified in the caption that the meteorological parameters and fluxes are instantaneous hourly values.**

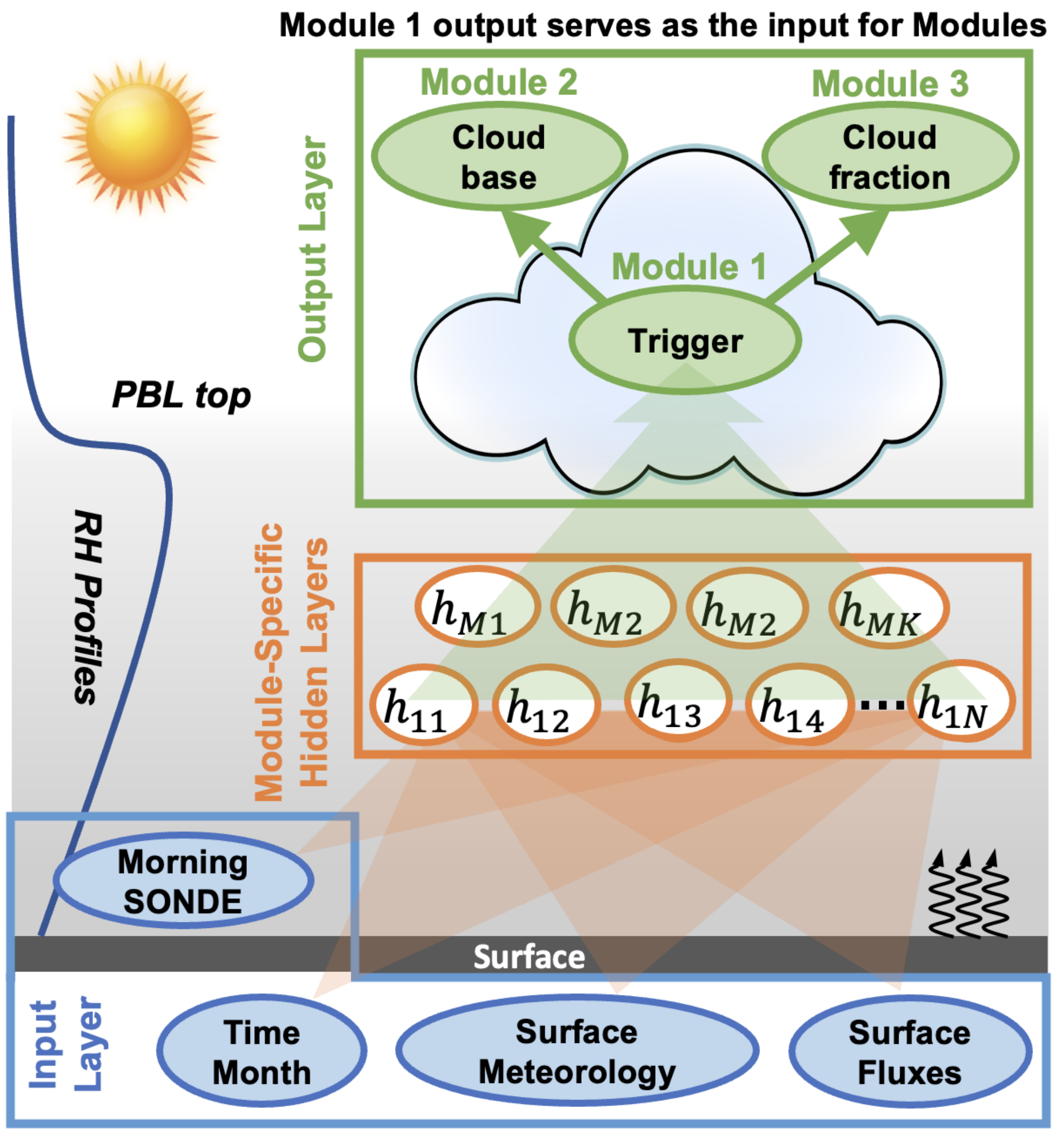

Figure R2 (The revised Figure 1). Conceptual diagram of the deep learning framework for simulating boundary layer cloud (BLC) characteristics over the US Southern Great Plains. Inputs for the deep neural networks (DNN) include morning meteorological profiles from radiosonde (SONDE), time indicators (i.e., local time and month), and surface conditions such as fluxes (curved black arrows) and meteorological data. The relevance of relative humidity (RH) profiles and the planetary boundary layer (PBL) top is emphasized due to their critical role in BLCs development. These variables are processed through multiple layers of hidden neurons (h11 to hMK). Both input and output parameters are hourly, except for the morning SONDE. Separate DNN modules are constructed for each task: Module 1 handles the initiation (trigger) of BLC; Module 2 estimates the cloud base; and Module 3 estimates cloud fraction and thickness. Together, these models synergize to predict the presence, altitude, and stratification of BLC.

Tab. 1 / Text Why is data used if it contributes a negative feature importance?

Thank you for pointing this out. We do filter out the input parameters based on their importance scores, as noted in the original lines 304-305. To clarify, we have extended the description: "After determining the importance scores from the test run, in refining the model, features contributing a negligible or negative effect on performance (i.e., importance scores less than zero) are excluded to ensure only beneficial data is used." In practice, filtering these parameters does not notably affect the results, as the negative values have a magnitude of -0.001.

l. 338/339 "measurements" and "surface meteorology" – what exactly is mean here?
Response: We have revised the text to specify that "measurements" refer to observed surface meteorology and fluxes.

l. 398' "Parameterization" – the DNN model has not parameterization for the cloud top.
Response: We have revised this statement as follow for better accuracy: "This can introduce inherent limitations, as the tops of many clouds may be decoupled from surface influences despite a coupled base, potentially leading to gaps in the DNN's ability to accurately define and estimate the cloud top".

l. 491/ l.559 ' (also compare major point #3) "a more accurate representation" – the DNN is employed here as an offline, a posteriori analysis tool. While you convincingly argue that the DNN has skill to yield a better cloud field, it is misleading to talk about a "better representation" as the DNN is run offline. In fact, the cloud field modified by the DNN is most likely inconsistent with the reanalysis! So, there is no better representation of clouds by just applying the DNN.
Response: We appreciate the reviewer's feedback and have revised the text to clarify that the DNN provides improved cloud field estimations rather than an integrated representation within reanalysis data. These statements have been revised as follow:
"……we achieve a more accurate estimations of cloud fractions for both stratiform and cumulus clouds."
"The implementation of this model within reanalysis datasets like ERA-5 and MERRA-2 has resulted in enhanced cloud field estimations for stratiform and cumulus clouds."

l. 577/8 "advancing our understanding of BLC dynamics" – this is not true. While we get improved cloud fields, the DNN tells little about the dynamics; in fact, the point of ML / deep learning is that we can have forecasts without understanding of the dynamics.
Response: We appreciate the reviewer's insight and have deleted the statement. Instead, we focus on the DNN's capability as a powerful diagnostic tool for improving cloud field estimations.

l. 577/8 "improving the representation of low clouds – see above point for lines 491/559.
Response: We have deleted this statement as noted above.

Technical comments / Typos:
l. 97: 'this model' – which model?
Response: We specified it as the deep learning model.

l. 99: 'we strive to narrow the gaps in boundary layer clouds' – bad style, please rewrite!

Response: We revised it as "we aim to help bridge the existing gaps in representing boundary layer clouds…"

l. 139/140 Syntax incorrect.
Response: We revised this statement as: "We treat BLCs as synonymous with land-coupled clouds, in contrast to clouds that are decoupled from the PBL and land surface."

l. 147/148 please cite the data by their DOI (which is provided on the ARM website)
Response: We added the DOI as suggested.

l. 148 abbreviation CBH for cloud base height is introduced to late; The phrase has been used before.
Response: We defined CBH for the first appearance.

l. 183 'a advanced' – please correct!
Response: Revised as suggested.

l. 581 What are "synoptic regions"?
Response: Revised it as "synoptic patterns".