Reply to Referee comment 1:

We extend our sincere gratitude to the reviewer for the time and effort dedicated to evaluating our manuscript, as well as for the invaluable comments and constructive suggestions provided. Below, we provide a point-by-point response to each comment and concern. All revisions made to the manuscript text are presented in green font for clarity.

---

1. I see that the new algorithm is well able to identify MCS, as it compares well to the threshold-based approach in case studies and also reproduces well global MCS climatologies. What I do, however, not really understand is why the new algorithm is that much faster than the threshold-based algorithm. In fact, I would argue that there is computationally a simpler approach than applying threshold to an input field, as is done for the traditional approach based on brightness temperature and precipitation. As mentioned in the text, the difference is very large: about 3 hours for the traditional approach, but only 2 minutes for the new approach. Is this really only for the identification of the MCS, or does the difference come from tracking the MCS and so attributing to the single MCS time-continuous labels? Possibly, I miss an essential point in the discussion?! It would be helpful to discuss in greater detail the reason for this large difference in computation cost?

Response: We thank the reviewer for this insightful question regarding the computational efficiency of our method. The dramatic speedup of MCSeg stems from fundamental differences in computational paradigms, which we clarify below.

1. Clarification of the Comparison Benchmark

We first clarify that the reported times (approximately 3 hours for the traditional method versus 2 minutes for MCSeg) represent a fair comparison of the MCS identification stage only, conducted on the same hardware, and do not include any subsequent tracking steps. This ensures an equitable comparison.

2. The Hidden Costs of the Traditional Threshold Method

The reviewer is correct that a single threshold operation is computationally cheap. However, a complete and robust traditional MCSs identification pipeline involves a series of computationally intensive post-processing steps:

- Connected Component Labeling (CCL): This constitutes the primary performance bottleneck. The algorithm must scan the entire two-dimensional global grid at each timestep to identify and label all independent, contiguous cloud clusters. The computational complexity of this operation scales with the number of grid points, and its inherent sequential data dependencies make it notoriously difficult to parallelize efficiently.
- Morphological Filtering: Subsequently, each labeled candidate cloud cluster must undergo a series of filters based on physical definitions . This requires expensive feature extraction on thousands of irregularly shaped objects, operations that are largely performed serially.

3. Sources of Efficiency in the Deep Learning Model

In contrast, the efficiency of MCSeg, as an end-to-end deep learning model, derives from several key factors:

- Highly Parallelized Forward Pass: Model inference is essentially a process of data flowing through a fixed sequence of layers. The core operations are transformed into highly regular tensor operations, which are perfectly suited for massive parallelization across the thousands of computing cores available on a GPU.

- Integrated Identification Pipeline: The model implicitly and synergistically performs feature extraction, contextual understanding, and pixel classification within a single, compact forward pass.It bypasses the necessity for the separate, sequential CCL and complex morphological filtering pipeline required by the traditional approach.
- Hardware-Level Optimization: The entire inference process is built upon deep learning frameworks optimized for accelerators like GPUs. The underlying computational libraries are extremely optimized for these fundamental operations, far surpassing the performance of hand-written, sequential CPU code.

In summary, the speed advantage does not come from a simplification of the algorithmic logic, but from replacing an inherently sequential algorithm pipeline with complex post-processing with a highly parallelized, hardware-friendly, and integrated model.

As a specific point: In the introduction (L37-39) it is explicitly written that the traditionl identification of MCS hinders the analysis of MCS climatological characteristics. Is this really true? If we rely on 30 min Imerg and BT data from, say, 2003 to present, that would be 20 years of data that has to be processed. Even with moderate computational resources that should be feasible in a reasonable amount of time and thus does not hinder climatological analysis. In fact, such global analysis have already been done.

Response: We thank the reviewer for this comment and for rightly pointing out the overstatement in our original wording. The reviewer is correct that traditional threshold-based methods have been successfully employed to produce valuable global MCSs climatologies, as evidenced by several key studies in the field. Our intention was to highlight the significant computational burden and inefficiency of these methods, which, while not rendering climatological analysis impossible, does pose a substantial practical barrier to rapid, iterative, and large-scale analysis. To accurately reflect this and to incorporate the reviewer's valid point, we have revised the relevant sentences in the introduction.

2. The authors convincingly show that their new algorithm performs very well compared to other ones, both in accuracy of MCS identification and in computational cost. This is shown in case study figures, in global climatologies and also in tables. I wonder whether all details are actually needed in the manuscript. So, for example, I can imagine that table 2 and its continuation table 3 are too detailed and at least part of them could be provided in supplementary material. Possibly, the same applies to part of figures 6-8.

Response: We thank the reviewer for the suggestion regarding the presentation of our detailed results. We agree that streamlining the main manuscript in this manner enhances its readability for a broad audience, while still ensuring that comprehensive data remain available to interested researchers. In accordance with this advice, we have relocated the extensive experimental results from Tables 2 and 3 to the Supplementary Information. This allows the main text to focus on the high-level conclusions drawn from these comparisons without being encumbered by the full datasets. Furthermore, we have carefully assessed Figures 5-8. While we believe a curated selection of these panels is essential in the main text to visually support our central claims regarding identification accuracy, we have moved the more specific regional analysis into the Supplementary Information, retaining only the global-scale visualization results in the main body. All relocated items are explicitly cited in the main text. We believe these changes have significantly improved the focus and presentation of our work and thank the reviewer again for this constructive recommendation.

3. In Section 7, some global characteristics of MCS are listed, e.g., their frequency and link to precipitation. The previous sections of the paper discuss algorithmic aspects of the MCS identification, or they compare the new algorithm's performance with other ones. This section, however, is much more strongly focused on meteorology and it also states that many of the findings are already known from existing literature. I think that this meteorological discussion only partly fits into the overall 'structure' of the manuscript. So, either the meteorological analysis should be extended and so bring new insights that can be gained based on the new, more efficient (faster) MCS identification. Since this is not the main focus of the study, I instead suggest to 'frame' this section also more strongly to show that the new method is able to reproduce existing climatologies of MCS.

Response: We sincerely thank the reviewer for this observation. We agree that the primary focus of the manuscript should remain on the validation of our novel deep learning algorithm. Following the reviewer's excellent suggestion, we have completely restructured the manuscript to reframe the climatological analysis as a key component of the algorithm's validation. Specifically, we have merged the original Section 6 ("Compared to the Threshold Method") and Section 7 ("Global MCSs Characteristics") into a single, comprehensive new section titled "Comprehensive Validation of the MCSeg Algorithm". Within this new section: The regional comparisons with the threshold-based method are presented as the first tier of validation. The global-scale analysis is now explicitly positioned as the second tier of validation. Instead of presenting meteorological findings, we now explicitly state that the purpose is to verify whether our algorithm can successfully reproduce well-established climatological patterns from the literature (e.g., spatial distributions, precipitation contributions, and seasonal cycles). We consistently use language that frames the agreement with prior studies as direct evidence of our algorithm's accuracy and physical realism.

4. There are some aspects that need to be clarified or improved in paper structure:

- Parts of the abstract are rather technical, at least if one is not too familiar with the machine learning approaches. As an example: many readers will not immediately understand what a 'significance learning strategy' and/or a 'multi-scale feature extraction methjod' is. Are these pieces of information really necessary in the abstract?

Response: We thank the reviewer for this suggestion. Following the reviewer's advice, we have revised the abstract to replace the technical terms 'significance learning strategy' and 'multi-scale feature extraction method' with more general descriptions of their functions.

- L147: The edges of the image are expanded and filled with non-MCS values of brightness temperature. Does this mean that the domain is not periodic in zonal direction, and if so: how does this lead to artificial MCS artefacts near the dateline?

Response: The reviewer rightly pointed out that our initial preprocessing approach, which involved expanding image edges and filling them with a constant BT value of 300 K, did not account for zonal periodicity and could indeed introduce artificial boundaries along the dateline. In response, we have revised the data preprocessing procedure. Specifically, we have eliminated the step of expanding the image edges to 2048×5632 and filling the peripheral regions with 300 K. Instead, we now explicitly incorporate zonal periodicity during the tiling phase. When generating sub-blocks for model input, any tile that crosses the dateline (180°W/180°E) is seamlessly wrapped by incorporating data from the opposite side of the domain. This ensures that the model is never exposed to hard, non-physical boundaries at the dateline during either training or inference.

- L133: Here, it is mentioned that ‘spurious MCS’ are to be excluded. What are spurious MCS, and if they are still MCS, why should they be excluded?

Response: We thank the reviewer for this comment. The term "spurious MCSs" was indeed potentially misleading, and we agree that it required a more precise explanation. Our intention was to convey that an overly lenient threshold (e.g., 250 K) would identify a large number of cold cloud regions, but not all of these regions represent robust, long-lived MCSs as defined by our study's objectives. These non-robust systems, which we previously referred to as "spurious". To address this comment and improve the clarity of our manuscript, we have revised the relevant sentence in the "Data Processing". The change shifts the focus from excluding "spurious MCSs" to the balanced selection of thresholds to minimize both false negatives and false positives.

- COD vs SOD: I am not completely sure whether I understand the distinction between the two? In Figure 1, it is written that the MCS in tropical and extratropical regions (Region 1 and 2 in the figure) are different in structure? Is it that the MCS differ in their degree of spatial clustering, or do all the individual MCS differ in their structure? For readers less familiar with the distinction between tropical and extratropical MCS some further background information (and references) could be helpful?

Response: We sincerely thank you for raising these critical points. We have thoroughly revised the manuscript to provide a clearer explanation. The difference is indeed structural at the individual MCSs level, stemming from their distinct formative environments. Tropical MCSs, developing in a more uniform environment, tend to have a coherent and compact structure (like a well-defined "blob"). Extratropical MCSs, often embedded in frontal systems, exhibit a more amorphous and diffuse structure, where the intense convective cores are embedded within a larger, less convective cloud shield. Based on the above, we have refined our analogy: Identifying the compact, well-defined tropical MCSs is analogous to Salient Object Detection (SOD), where the target is the most prominent and distinct object in the image. Identifying the embedded convective cores within the diffuse extratropical cloud shield is analogous to Camouflaged Object Detection (COD), where the target is visually similar to its background and lacks clear boundaries. As suggested, we have added a sentence acknowledging the meteorological drivers (e.g., vertical wind shear) behind these structural differences and have included relevant references (e.g., Galarneau Jr et al. (2023);

Muetzelfeldt et al. (2025); Paul et al. (2025)) to provide further background for readers. We thank the reviewer again for helping us strengthen this part of our work.

- Most likely related to the previous point: The title of Section 2.2 is somewhat ‘misleading’. I would not have doubted that machine learning can be used to identify MCS. The real point of Section 2.2 is that the authors suggest to use different approaches (SOD vs. COD) for low and mid latitudes. This should be reflected also in the title of the Section 2.2.

Response: We thank the reviewer for this suggestion. Following the reviewer's advice, we have changed the title of Section 2.2 to " The Latitudinal Challenge in Deep Learning-Based MCS Identification" to better reflect our core proposition of employing distinct strategies (SOD for tropics, COD for extratropics) for MCS identification across different latitudinal regions.

- L67-76: This part is less about MCS identification, but more about background information (occurrence frequency, link to precipitation). I think it would better fit into the introduction or, possibly, the discussion.

Response: We thank the reviewer for this suggestion. As the reviewer rightly pointed out, the core theme of this section should remain focused on the identification methodology itself. Following the reviewer's advice, we have evaluated the option of moving this paragraph. However, we found that the connection between MCSs and precipitation is already established in the introduction to provide motivation for the study, and the analysis of this link is further elaborated in the discussion section of our revised manuscript. Therefore, to maintain the conciseness and thematic coherence of the methodology section, we have decided to delete this paragraph (Lines 67-76) entirely. We believe this change significantly improves the focus and clarity of the section. Thank you for the valuable comment.