

Point-by-point reply on RC2

Black font = Reviewers' comments

Blue font = Response to reviewers' comments

Red font = Changes in the text of the manuscript

We thank the editor and reviewers for the valuable feedbacks on our manuscript. We have addressed all point-by-point reply below.

This manuscript presents the Multi-method Snow Drought Toolbox (MSDT), an open-source Python framework integrating multiple snow drought identification methods across yearly, monthly, and daily temporal scales. The authors aim to address the fragmentation of existing snow drought analysis workflows by providing a unified platform for SWE data acquisition, snow drought detection, visualization, and standardized output generation.

The development of such a toolbox is timely and potentially valuable for the cryosphere research community. The manuscript demonstrates the capability of MSDT through global-scale analyses using ERA5-Land SWE and several representative snow drought events. The modular design and integration of multiple approaches represent an important step toward improving reproducibility in snow drought studies. Several concerns should be addressed before publication. I recommend a Moderate Revision.

Major Concern 1: The algorithm implementation requires greater transparency

The manuscript integrates seven snow drought identification methods. However, the technical description is insufficient for independent reproduction. Important missing information includes detailed algorithm workflow, input/output structure, handling of missing values, threshold implementation, etc.

For example, SWEI allows multiple probability distributions, but the criteria for selecting or rejecting distributions are unclear. The authors should provide algorithm pseudocode, flowcharts for individual methods, example parameter files if applicable, especially PDM and SWEI.

Response: Thank you for the comment. We have added a supplementary algorithm description that provides method-specific pseudocode, individual flowcharts, input and output structures, missing-value handling, threshold implementation, and example parameter files for all seven snow drought identification methods. In particular, the SWEI section now specifies the supported probability distributions, fitting conditions, fallback rules, and probability clipping procedure, while the PDM section documents the winter-period filtering, 7-day rolling calculation, percentile ranking, drought-type classification, and memory-aware processing strategy.

For more details, please refer to the Text S1 in the Supplement.

.

Major Concern 2: The software contribution lacks sufficient quantitative evaluation of

computational performance

The manuscript highlights several computational improvements, including lazy NetCDF loading, adaptive parallelization, float32 optimization, and compressed NetCDF output. However, these improvements are described qualitatively. The authors should provide quantitative benchmarking which describing the hardware configuration, dataset size, runtime, peak memory consumption, and comparison with traditional scripts. These would significantly strengthen the software contribution.

Response: We thank the comments for requesting standardized quantitative performance evidence, which was metioned by all three reviewers. The supplementary benchmarking attachment now reports all seven methods, temporal and spatial scaling, and lossless output compression.

All tests used the same Intel Core i9-12900K workstation (16 cores, 24 logical processors), 128 GB RAM, Windows 11 Pro, and Python 3.11.7. The full-domain inputs cover 20–84°N at 0.25° over 1980–2024 with the supplied extent mask. Annual methods and PDM use daily inputs, whereas SWEI uses monthly SWE. Each configuration was evaluated in three independent processes. Reported medians include input reading, calculation, compression and output completion; peak RAM is the maximum process resident memory.

The five annual methods required 124.6–319.3 s, with peak RAM medians of 1.76–27.11 GiB. SWEI required 117.3 s and 5.28 GiB using 3-month accumulation and empirical Gringorten plotting positions. PDM required 27.4 min and 28.45 GiB, producing 44 annual files totalling 7.38 GiB. Lossless compression reduced the tested output sizes by 88.2–98.2%, with identical decoded values.

The scaling tests cover 10-, 20- and 45-year records and regional grids at 0.25°, 0.1° and 0.01°. The finer regional grids are computational stress tests derived by nearest-neighbour replication, not native observations or demonstrations of full-domain high-resolution feasibility. Based on the measured peaks, 64 GB RAM is a practical planning recommendation for larger workloads, with headroom for the operating system.

For more details, please refer to Text S3 in the Supplement.

.

Minor Concerns

1. The last two paragraphs of the Introduction are somewhat repetitive and should be revised to avoid redundancy.

Response: We revised the last paragraphs to remove repeated statements as follows:

“In the following, we introduce the implemented snow drought identification methods together with the software architecture and graphical interface in Sect. 2. Sect. 3 demonstrates the capability of MSDT using global analyses and representative historical snow drought events. Sect. 4 examines the scientific implications,

computational performance, current limitations, and future development of the toolbox.”

2. The manuscript should clarify whether snow drought intensity has consistent meaning among different methods.

Response: Thank you for the insightful comment. The definition of snow drought intensity varies across different methods. Specifically, the intensity derived from the yearly-scale methods (except the PeakSWEDate method) and the PDM method is quantified as ΔSWE (mm), while the intensity from the PeakSWEDate method is Δday (days). For the monthly-scale method, intensity corresponds to the SWEI value (dimensionless). Descriptions of snow drought intensity are also documented within the output NetCDF (nc) files.

3. The use of ERA5-Land SWE as the primary demonstration dataset requires additional uncertainty discussion.

Response: We added an uncertainty statement that defines ERA5-Land as a demonstration dataset rather than an error-free reference and identifies likely limitations in complex terrain:

“It should be aware that ERA5-Land SWE may be affected by forcing, model-physics, and terrain-representation uncertainties, particularly in mountainous regions. The dataset is used here to demonstrate MSDT functionality; users should evaluate alternative SWE products for region-specific applications.”

4. The snow-covered area filtering approach needs more justification. The choices of 5% MODIS threshold and 75% SWE occurrence threshold appear empirical.

Response: Thank you for the comment. The snow-covered area filtering approach is the method adopted in this paper and has been applied in global studies (Huning and AghaKouchak, 2020; Li and Wang, 2022). The MSDT system does not restrict this process. Users can use this as a reference or determine their own snow-covered area using other methods, and then specify the Extent Mask in the MSDT program in the form of an Excel table (optional) to reduce the computational load. We have added relevant explanations in Section 2.2:

“The MSDT system does not restrict the method for identifying the snow-covered area. Users can use this as a reference or employ other methods to determine their own snow-covered area. Then, in the form of an Excel table (optional), they can specify the Extent Mask within the MSDT program to reduce the computational load.”

References:

- Huning, L.S., and AghaKouchak, A.: Global snow drought hot spots and characteristics, *Proc. Natl. Acad. Sci.*, 117, 19753-19759, <https://doi.org/10.1073/pnas.2005022117>, 2020
- Li, X., and Wang, S.: Recent Increase in the Occurrence of Snow Droughts Followed by Extreme Heatwaves in a Warmer World, *Geophys. Res. Lett.*, 49, e2022GL099925, <https://doi.org/10.1029/2022GL099925>, 2022.

5. The manuscript allows different probability distributions for SWEI but does not discuss their influence on results.

Response: Thank you for the comment. In Section 2.3.2, we have introduced several probability distributions supported by MSDT and their respective applicability:

“The choice of distribution can influence the probability transformation of accumulated snow water equivalent and therefore affect the estimated spatial extent, severity, and duration of snow drought events. The empirical Gringorten approach does not assume a predefined distribution and is useful when sample sizes are limited or the data distribution is complex; however, it has limited ability to represent probabilities beyond the observed record (Gringorten, 1963; Huning and AghaKouchak, 2020). Parametric distributions may provide smoother probability estimates, but their performance depends on how well the assumed distribution matches the observed snow water equivalent characteristics. Gamma, Lognormal, and Weibull distributions are often suitable for non-negative and right-skewed hydroclimatic variables, whereas Pearson type III has been widely recommended in standardized drought-index calculations because of its flexibility in representing skewed distributions (Guttman, 1999; Stagge et al., 2015). The Normal distribution is simple and computationally stable, but it may underestimate extremes in regions with strong skewness or frequent zero-snow conditions. Therefore, probability distribution selection should be treated as a source of methodological uncertainty. Users are encouraged to compare multiple distributions, especially in regions with short records, intermittent snow cover, or complex snow regimes.”

6. The role of temperature, radiation, and rain–snow transition in snow drought development should be discussed more explicitly.

Response: Thank you for the comment. The main objective of this article is to provide a unified framework for existing snow-drought identification methods. In our revised discussion, we envisioned that in the future, snow-drought identification methods that take into account non-temperature and non-precipitation factors should be integrated into the MSDT:

“An important direction is to integrate snow drought identification methods that account for drivers beyond air temperature and precipitation, such as radiation, wind redistribution, sublimation, and blowing-snow processes. These factors can strongly influence snow accumulation, redistribution, and ablation, particularly in open, alpine, and high-latitude environments where wind exposure and surface energy balance play major roles.”

7. The manuscript should clarify whether PDM can distinguish warm snow drought and dry snow drought.

Response: Thank you for the comment. We have clarified this in Section 2.1.3 as follows:

“Dry snow drought occurs when accumulated precipitation is below the median and

warm snow drought occurring when it is above the median. Within below-median precipitation conditions, a precipitation percentile exceeding the SWE percentile indicates dry but warm snow drought, suggesting a comparatively greater snowpack deficit in percentile terms. Thus, this approach does not directly apply a temperature threshold but instead uses the combined precipitation and snowpack anomalies to infer potential warming-related influences, such as rainfall replacing snowfall or enhanced snowmelt.”

8. Figure 1 is informative but the text size is too small for publication.

Response: Thank you for the comment. We have enlarged font size in Fig. 1 and Fig. 2 as follows:




Fig. 1




Fig. 2

9. Parallel computing strategy requires more technical details.

Response: Thank you for this comment. We have provided detailed pseudocode and parameter definitions in Text S1 of the Supplement to document the daily PDM parallel computing strategy. The workflow integrates memory-aware input loading, spatially partitioned calculation, and incremental annual output.

After restricting SWE, reference SWE, and precipitation inputs to the selected winter period, MSDT detects available system memory using `psutil.virtual_memory().available`. Input sizes are estimated from the filtered array dimensions assuming float32 storage. The default working-memory budget is approximately 50% of available memory, initially subject to a minimum of 1 GiB and subsequently capped to retain system headroom equivalent to the larger of 2 GiB or 20% of available memory, where possible. The streaming planner allocates an estimated working-memory allowance of 16 times the raw input-block size to accommodate input arrays, intermediate calculations, and block-level results. Small inputs are processed as a single block when they satisfy both the adaptive in-memory limit, which does not exceed 2 GiB by default, and the working-memory budget divided by this allowance. Larger inputs are partitioned into latitude blocks, with nominal input-block targets constrained to 256 MiB–1 GiB. Each block retains the complete filtered temporal dimension required for reference-period percentile estimation.

Block calculations use a `ThreadPoolExecutor` with at most two concurrent workers. Concurrency is determined by the estimated working-memory requirement and the number of active blocks. A bounded submission queue controls the number of outstanding tasks, and replacement tasks are submitted after preceding block results have been written. Dask uses a single-threaded scheduler within this workflow. NetCDF access is protected by a shared library lock, and a single writer incrementally stores each block at its corresponding spatial position in annual output files. Block results are released after writing. Blocks containing no valid extent-mask cells are skipped.

Further details are provided in Text S1 of the Supplement. The relevant manuscript text reads:

“For the computationally intensive daily PDM method, MSDT uses a memory-aware calculation and output workflow. Inputs are first restricted to the selected winter period. Small datasets are processed as a single spatial block, whereas larger datasets are partitioned into latitude blocks that retain the complete temporal sequences needed for percentile estimation. Block size and concurrency are determined from the filtered input dimensions and available system memory, with explicit parallel execution limited to two workers. A bounded task queue controls the number of resident block results, and blocks outside the extent mask are omitted. Floating-point intermediate arrays use float32 precision where appropriate. A single writer incrementally stores block results in annual NetCDF files, and these results are released after writing. Outputs use lossless NetCDF4 compression at level 4, with shuffle enabled and storage chunks aligned with spatial blocks. This workflow controls memory demand throughout calculation and output and reduces storage requirements.”

10. The NetCDF output structure should be documented according to CF conventions

Response: Thank you for this suggestion. We have documented the NetCDF output structure in the Supplement (Text S2) and revised the shared output metadata to follow CF conventions, version 1.8. The documentation specifies variable dimensions, stored data types, units, coordinate definitions, missing-value encoding, classification codes, and method-specific threshold representations. Annual reporting periods are explicitly identified using date-based time coordinates and time bounds, while PDM outputs retain actual daily dates and record their hydrological reporting period within each file. Missing or masked records are distinguished from non-drought conditions. Classification variables carry consistent flag_values and flag_meanings attributes, and custom descriptions and classification rules are preserved. Both the general writer and incremental PDM writer use the revised schema. Sixteen regression tests passed, including comparisons with pre-update finite-valued results on synthetic data. Twelve representative output files, covering all seven methods, extent masking, and incremental PDM export, also passed the CF-1.8 checks in IOOS Compliance Checker 6.1.0 using strict criteria. We provide the output specification and representative file headers to facilitate independent interpretation and reuse.

11. Table 2 should include data types and units for each output variable.

Response: Thank you for the comment. We have included Table S2 in the Supplementary Information for clarification:

“Table S2. Data type, units, and notes of Snow drought characteristics			
Variable	Data type	Unit	Notes
sd_flag	integer, preferably int8	unitless / 1	Binary snow drought record: 1 = snow drought, 0 = non-drought. In current code, SWEI and PDM are explicitly int8; some yearly methods may be written as default integer unless further standardized.
sd_type_flag	int8	unitless / category	Only used by PDM. 0 = non-drought, 1 = warm snow drought, 2 = dry but warmer, 3 = dry but snowier.
intensity	floating point, usually float32 or float64	method-dependent	SWE anomaly for SWE-based methods, days for PeakSWEDate, and standardized index value for SWEI.
intensity_flag	integer, preferably int8	unitless / category	Severity class of snow drought. Flag values follow Table S1.

swe	floating point, usually float32 or float64	same as input SWE, commonly m or mm	SWE-related metric used by the selected method. Written mainly for yearly methods; SWEI and current PDM output do not write raw swe.
threshold	floating point, usually float32 or float64	method-dependent	Threshold used to identify snow drought. Unit follows the compared variable: SWE unit for SWE-based methods, days for PeakSWEDate, and not written as a separate variable for SWEI.

”

12. The visualization module should support modern mapping libraries instead of relying only on Basemap, which is becoming outdated.

Response: Thank you for the comment. Basemap was used during the initial stage of development. In the version submitted, Matplotlib and Cartopy were employed. We have revised the relevant statements:

“Interactive geospatial visualization is implemented using Matplotlib and Cartopy, ...”

13. The manuscript should provide example workflows in the Supplementary Material.

Response: Thank you for this helpful suggestion. In the revised manuscript, we have added example workflows to the Supplementary Material, including step-by-step procedures for data preparation, parameter setting, execution of the seven snow drought identification methods, and interpretation of the output NetCDF variables. Example parameter templates are also provided for each method, with particular detail for SWEI and PDM because they involve probability distribution fitting, rolling accumulation, percentile ranking, and method-specific classification rules.

For details, please refer to Text S1 in the Supplement.

14. Claims regarding “standard desktop computers” are too general without hardware specification.

Response: Thank you for this helpful suggestion. We have deleted “standard desktop computers”.