CLIMATE INDEX

ACTUARIAL SOCIETY OF SOUTH AFRICA

Technical Documentation



Rendani Mbuvha and
Nadine van der Walt

September 2026

email: climate_index@actuarialsociety.org.za

Contents

1 Introduction **2**

2 Data Sources **2**
2.1 Rationale for gridded data over national station networks . . . 3
2.2 Applicability of CHIRPS and AgERA5 to South Africa . . . 3
2.3 Terms of use . . . 4

3 Overview of the components **4**

4 Extreme Precipitation **5**

5 Drought **6**

6 Extreme Maximum Temperature **6**

7 Extreme Minimum Temperatures **7**

8 Hail **7**
8.1 Input data . . . 8
8.2 0–6 km AGL deep-layer wind shear . . . 8
8.3 Daily hail-environment field and filter . . . 9
8.4 CAPE observing-system homogenisation . . . 9
8.5 Standardisation pipeline . . . 13
8.6 Event-based validation . . . 14
8.7 Methods investigated and rejected . . . 15
8.8 Long-term behaviour . . . 16

9 Wildfire **19**
9.1 Input data . . . 20
9.2 Hot-Dry-Windy Index (HDWI) . . . 21
9.3 Canadian Fire Weather Index (FWI) . . . 22
9.4 Validation of the FWI implementation against GEFF . . . 23
9.5 Normalised-maximum composite . . . 24
9.6 Standardisation pipeline . . . 25
9.7 Event-based validation . . . 26
9.8 Methods investigated and rejected . . . 28
9.9 Long-term behaviour . . . 30

10 Composite Index **32**

References **34**

1 Introduction

The frequency and severity of insurance losses caused by extreme weather events have increased over the past decade. This trend is consistent with the growing evidence of global warming and increased climate variability. The Actuarial Society of South Africa (ASSA) Climate Index is an industry initiative to develop straightforward and consistent metrics for tracking and quantifying the frequency of climate extremes across South Africa. Its development was informed by leading international efforts including the European Extreme Events Climate Index (E3CI), the Actuaries Climate Index (ACI) for the United States and Canada, and the Australian Actuaries Climate Index (AACI), and was led by Prof Rendani Mbuvha and Dr Nadine van der Walt with expert input from Prof Willem Landman from the University of Pretoria. The index serves actuaries, policymakers, and researchers by providing a transparent measure of how often and how intensely extreme conditions occur, with a focus on temperature extremes, heavy precipitation, and drought. It supports risk assessment, pricing and capital modelling work, climate scenario planning, and disclosures across sectors sensitive to weather and climate variability.

The ASSA Climate Index comprises component indices for extreme maximum temperature, extreme minimum temperature, extreme precipitation, and drought, together with the hail and wildfire components described in Sections 8 and 9. It also provides a composite measure that integrates the temperature, precipitation and drought signals. Each component is produced as a time series derived from historical observations and expressed as standardised departures from a long-term baseline, facilitating clear interpretation and comparison over time and across regions. Updated on a seasonal cycle, the index enables users to monitor emerging patterns, detect shifts relative to historical norms, and assess implications for environmental and economic stability in the South African context.

The index is built from two gridded, observation-based datasets that provide consistent national coverage: CHIRPS (9, 15) satellite-derived precipitation ($\approx$ 4 km grid cells) and AgERA5 (10) agro-meteorological fields derived from reanalysis ($\approx$ 11 km grid cells). All component calculations are performed at the native grid-cell level, including the construction of baselines and standardised departures, to preserve local climatology and extremes before any aggregation. This bottom-up approach ensures that regional signals are not dominated by stations with uneven coverage (this would require spatial interpolation, leading to further model risk) and that spatial variability is retained in the intermediate results.

Spatial and temporal aggregations are then applied in a transparent, averaging framework. Regional values are computed as the arithmetic mean of the index values for all grid cells whose centroids fall within the region. Time aggregation follows the same principle: monthly component values are first computed, and seasonal indices are formed by averaging the constituent months of the climatological seasons (DJF, MAM, JJA, SON); other user-specified periods are treated analogously. Planned enhancements include expanding the set of components and refining spatial aggregation methods (for example, alternative weighting schemes) to better reflect exposure patterns and sectoral use cases.

2 Data Sources

To ensure spatially complete, reproducible, and regularly updated coverage over South Africa, the index draws on two widely used gridded products. Precipitation is taken from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), a quasi-global rainfall dataset at 0.05° resolution (about 4–6 km across South Africa) that blends infrared satellite estimates with in situ gauge observations and extends from 1981 to near-real time. Temperature and precipitation are sourced from AgERA5, which includes a suite of indicators derived from the ERA5 reanalysis and curated for agricultural applications by the European Centre for Medium-Range Weather Forecasts (ECMWF), distributed via the Copernicus Climate Data Store (CDS) at 0.1°

resolution (about 10–11 km) from 1979 to present. A reanalysis is a physics-based reconstruction of the historical atmosphere that blends a numerical weather model with all available observations to produce a complete, uniformly gridded record. CHIRPS was downloaded from the Climate Hazards Center, while AgERA5 was obtained from the CDS. Both datasets offer stable, long historical records, consistent national coverage, and routine updates aligned with our seasonal production cycle.

2.1 Rationale for gridded data over national station networks

While South African Weather Service (SAWS) station observations are fundamental, using them directly as the backbone of a national index raises practical and methodological challenges common to many station networks and particularly salient in South Africa: uneven station density (especially in sparsely populated regions), intermittent outages and periods of missing data, station relocations and instrumentation changes that introduce inhomogeneities, variable and sometimes limited metadata and homogenisation, and lags or restrictions in access and redistribution. Licensing and cost barriers can also complicate transparent, reproducible, and open publication of national-scale indices.

By contrast, CHIRPS and AgERA5 provide spatially complete fields with consistent processing across the country, long and regularly updated records and clear, open terms that facilitate reproducibility and broad industry use. CHIRPS's gauge-satellite blending improves performance in data-sparse areas and complex terrain, while AgERA5 supplies internally consistent indicators on a uniform grid, both well suited to component-wise calculations at grid level and aggregation to regions and seasons.

2.2 Applicability of CHIRPS and AgERA5 to South Africa

Evaluations over southern Africa consistently support using ERA5-based reanalyses for temperature, with AgERA5 offering a practical balance of accuracy and spatial coverage for impacts work. Regional and subcontinental studies show good consistency between reference observations and ERA5-derived datasets for seasonal means and the frequency of hot and cold extremes, while noting that event magnitudes can be less precise in some settings, hence the value of higher-resolution ERA5 variants for impact analyses (25). A focused assessment across the Northern Cape further demonstrates that AgERA5 reproduces seasonal means and interannual variability well at station sites, and therefore provides a suitable foundation for temperature components of the index although certain regions such as the Kalahari Desert should be treated with caution due to model limitations (18).

For precipitation, multiple South African evaluations point to strong spatial heterogeneity in dataset accuracy, motivating a dual-source approach. CHIRPS is expected to perform well over interior and mountainous regions. In the Enkangala Escarpment / Drakensberg, CHIRPS showed the best overall statistics versus gauges and resolved orographic gradients more effectively than TerraClimate or TAMSAT, which have been found to over- and under-estimate precipitation, respectively (1). A national assessment found high correlations with monthly rain-gauge totals in the north and parts of the Western / Central catchments, but weaker correlations along the eastern seaboard (13), although CHIRPS have been found to be useful in assessing drought conditions in the Eastern Cape region (20). Complementary evidence from neighbouring Eswatini indicates CHIRPS tracks monthly and seasonal variability well but tends to underestimate totals and exhibits lower correlations in complex topography and transition zones, which are patterns relevant to South Africa's eastern highlands and coastal escarpment (12).

Performance challenges are most pronounced in the winter-rainfall southwest regions of South Africa. A catchment-scale comparison reported that CHIRPS underestimated mean annual precipitation in the Western Cape and recommended product recalibration for the winter rainfall zone, even as it performed reasonably elsewhere and within Kruger National Park (22).



Independent assessments over the Western Cape likewise show that satellite and reanalysis products exhibit mixed skill and are sensitive to station representativeness and spatial aggregation, underscoring the need for local validation (11).

ERA5-derived precipitation adds value by providing spatially and temporally consistent fields and capturing large-scale seasonal signals, but South African analyses highlight biases and timing discrepancies for convective rainfall and heterogeneous local events, reinforcing the case for combining sources (28). Continental and regional comparisons further emphasize that accuracy varies by regime and application. CHIRPS is often recommended for drought monitoring and climatology, while no single product dominates across all conditions. Site-specific validation and, where appropriate, bias correction are recommended (6).

Methodological work based on Southern Africa also shows that observational choices matter. CHIRPS better captures light-to-moderate events relative to some gauge products, yet higher extremes remain challenging even after bias correction, which has implications for how drought and heavy-rain components are constructed (3). Collectively, these findings justify using AgERA5 for temperature, and combining CHIRPS with ERA5-based precipitation for a more robust and regionally responsive precipitation signal across South Africa.

2.3 Terms of use

CHIRPS is in the public domain: the Climate Hazards Center has waived copyright and related rights under the Creative Commons CC0 dedication, so the data may be used, adapted and redistributed for any purpose, including commercial use, without restriction. Citation of the dataset publication (15) remains requested good practice, and the data are provided "as is" without warranty[1].

AgERA5, ERA5 and the other products obtained from the CDS are provided under the Creative Commons Attribution 4.0 International licence (CC-BY 4.0), which replaced the previous Licence to Use Copernicus Products on 2 July 2025. CC-BY 4.0 permits free use, sharing and adaptation for any purpose, including commercial use, provided that credit is given to the Copernicus Climate Change Service (C3S) and ECMWF, a link to the licence is included, and any changes made to the data are indicated; the data are provided without warranty. The attribution and citation requested for each dataset are listed in the References section of its CDS entry[2].

The licence conditions above were last reviewed in July 2026 and may change from time to time. Users of the index who access, redistribute or build on the underlying datasets are encouraged to consult the current licence agreements published by the data providers.

3 Overview of the components

Six components are described, and they fall into two groups that need different amounts of documentation.

For extreme precipitation, drought and the two temperature components, the daily quantity is a measured field. Daily rainfall and the daily maximum and minimum temperature are supplied directly by the source datasets, so each component applies a percentile threshold to an observation and the only choices to justify are the threshold, the baseline and the aggregation.

The hail and wildfire components have no observed daily variable to threshold. Hail is represented by an environmental proxy, the product of convective available potential energy and deep-layer shear under a freezing-level condition, because the reanalysis does not resolve hail itself. Wildfire is a composite of two published fire-weather indices, each a model with its own

[1] https://www.chc.ucsb.edu/data/chirps

[2] https://creativecommons.org/licenses/by/4.0/

inputs and calibration. For both, the daily field is constructed rather than observed, so three further things have to be established: how it is built and why, whether the implementation reproduces an independent reference, and whether it identifies documented events. Both also required a correction to the underlying reanalysis, which the measured variables did not. Sections 8 and 9 therefore carry that additional material, and the difference in length reflects what has to be justified rather than any difference in the standing of the components.

4 Extreme Precipitation

The precipitation component is based on daily gridded values based on the AgERA5 reanalysis product and CHIRPS data. Since each dataset has its own strengths and weaknesses, which results in regional differences in accuracy between the two datasets, the precipitation is provided for both datasets separately. It is imperative that users familiarise themselves with the appropriate dataset to be used based on the region that is of interest. The calculations below are done for each grid cell independently, and this approach is consistent over all climate components.

1. **Reference Climatology:** Let $p(d)$ denote the rolling 5-day total precipitation amount[3] in mm for calendar day d in the year (ignoring leap years, which is dealt with below). Consider a rolling 31-day window centered around $p(d)$ i.e. the window includes the 15 days before and after day d, as well as day d itself). A 31-day window is wide enough to increase sample size, and narrow enough to avoid problems with rainfall seasonality.

   For each d, the precipitation data for all 31-day windows over all years in the base period[4] (1991–2020) are combined. Note that for calendar dates corresponding to early January and late December, this dataset will contain some observations outside of the base period. This yields a dataset of 930 observations (30 years' 31-day observations) per value of d. The climatological threshold $p_{95}(d)$ is then defined as the 95^{th} percentile of this set. To account for leap years, the 29^{th} of February is assigned the average of $p_{95}(59)$ and $p_{95}(60)$ (17).

2. **Frequency Calculation:** For each month m and year k in the timeseries, we then determine $P_{m,k}$, the percentage of days for which the 5-day rolling sum exceeds its corresponding percentile. Let μ_m^{base} denote the mean of this number for calendar month m, where the mean is taken over all years in the base period. By definition, μ_m^{base} is typically close to 5%, although some variation is expected for regions and months with low typical rainfall. Similarly, let σ_m^{base} denote the standard deviation of the exceedance percentages for month m and taken over all years in the base period.

3. **Standardised Anomalies:** The standardised anomaly for month m in year k is then given by

$$\delta_{m,k}^{P} = \frac{P_{m,k} - \mu_m^{\text{base}}}{\sigma_m^{\text{base}}}.$$

4. **Dry grid cells:** Note that some grid cells over South Africa have few wet days during the base period for some calendar months. It is therefore possible that for some grid cells and months σ_m^{base} is zero (or, in practice, below a small threshold of 0.001) because $P_{m,k}$ is (almost) constant across the years k in the base period. To avoid an undefined anomaly in such a month, the monthly departure $P_{m,k} - \mu_m^{\text{base}}$ is retained but divided by the standard deviation pooled over all base-period months of the same season (DJF, MAM, JJA, SON) in place of σ_m^{base}. Since these cases are rare and both forms are standardised, this does not bias the results and is preferred over an incomplete index time series.

[3] Precipitation values less than 1mm are rounded to 0mm before calculating the rolling sum.

[4] `https://wmo.int/media/news/updated-30-year-reference-period-reflects-changing-climate`

5 Drought

The calculation of the Drought component is designed on the same principles as the Extreme Precipitation component. The component is calculated on both AgERA5 and CHIRPS data, and both sets of results are provided for users of the ASSA Climate Index. Similar to the Extreme Precipitation component, users should select the most appropriate dataset based on their requirements.

1. **Reference Climatology:** Let $q(d)$ denote the rolling 365-day total precipitation amount[5] in mm for calendar day d in the year (leap years are handled in the same way as for the Extreme Precipitation component).

   A full annual accumulation is used to reflect the way drought conditions develop and are experienced in practice: as multi-seasonal deficits in rainfall that affect water availability, crop yields, ecosystems, and infrastructure. From the perspective of users such as insurers, reinsurers, and policymakers, an annual horizon provides a more stable and decision-relevant signal than shorter windows, which may be dominated by transient anomalies.

   Consider a rolling 31-day window centered around q_d as above. For each d, the precipitation data for all 31-day windows over all years in the base period (1991–2020) are combined similar to the Extreme Precipitation component, yielding a dataset of 930 observations (30 years' 31-day observations) per value of d. The climatological threshold $Q_{05}(d)$ is then defined as the 5^{th} percentile of this set, and here we are interested in observations below this threshold.

2. **Frequency Calculation:** For each month m and year k in the timeseries, $Q_{m,k}$ denotes the percentage of days for which the 365-day rolling sum is below its corresponding percentile. Let μ_m^{base} denote the mean of this number for calendar month m, where the mean is taken over all years in the base period. By definition, μ_m^{base} is typically close to 5%. Similarly, let σ_m^{base} denote the standard deviation of the exceedance percentages for month m and taken over all years in the base period.

3. **Standardised Anomalies:** The standardised anomaly for month m in year k is then given by

$$\delta_{m,k}^{D} = \frac{Q_{m,k} - \mu_m^{\text{base}}}{\sigma_m^{\text{base}}}.$$

6 Extreme Maximum Temperature

The Extreme Maximum Temperature component is based on the frequency of *hot spells*. We define *hot spells* as periods of three (3) consecutive days during which the maximum daily air temperature at 2 m height exceeds a climatological percentile threshold. We use daily maximum temperature data ($T_{\max}$) from the AgERA5 reanalysis product.

1. **Reference Climatology:** For each calendar day d of the year, a climatological threshold $T_{95}(d)$ is computed as the 95^{th} percentile of $T_{\max}$ over a baseline reference period (1991–2020). To increase sample size and smooth inter-day variability, a ± 2 day moving window centered on each day of the year is used, giving a total window width of 5 days and 150 days for each percentile estimate.

2. **Frequency Calculation:** For each day of the year d in year k, we define an exceedance mask $M_{d,k}$ as:

$$M_{d,k} = \begin{cases} 1, & \text{if } T_{\max}(d,k) > T_{95}(d), \\ 0, & \text{otherwise,} \end{cases}$$

[5] Precipitation values less than 1mm are rounded to 0mm before calculating the rolling sum.

For each month m in year k, let $T_{m,k}^{Max}$ denote the hot spell frequency, where a hot spell is defined as a sequence of at least three consecutive days with $M_{d,k} = 1$. That is, $T_{m,k}^{Max}$ counts the number of runs of length ≥ 3 in the series of exceedance masks $M_{d,k}$ within the month m. For each calendar month, let μ_m^{base} and σ_m^{base} denote the reference period mean and standard deviation of $T_{m,k}^{Max}$.

3. **Standardised Anomalies:** The standardised anomaly for month m in year k is then given by

$$\delta_{m,k}^{T_{Max}} = \frac{T_{m,k}^{Max} - \mu_m^{\text{base}}}{\sigma_m^{\text{base}}}.$$

7 Extreme Minimum Temperatures

The Extreme Minimum Temperature is based on the frequency of *cold spells*. We *cold spells* define as periods of three (3) consecutive days during which the minimum daily air temperature at 2 m height falls below a climatological percentile threshold. We use daily minimum temperature data (T_{min}) from AgERA5 and proceed as follows:

1. **Reference Climatology:** For each calendar day d, compute the 5^{th} percentile threshold $T_5(d)$ of T_{min} over the baseline period (1991–2020). Similarly to the extreme maximum temperature module, a ± 2 day moving window (width 5) is applied around d to increase the sample size and smooth interday variability.
2. **Frequency Calculation:** For each day of the year d in year k, we define a (negative) exceedance mask $M_{d,k}$ as: as:

$$M_{d,k} = \begin{cases} 1, & \text{if } T_{\text{min}}(d,k) < T_5(d), \\ 0, & \text{otherwise,} \end{cases}$$

For each month m in year k, let $T_{m,k}^{Min}$ denote the cold spell frequency, where a cold spell is defined as a sequence of at least three consecutive days with $M_{d,k} = 1$. That is, $T_{m,k}^{Min}$ counts the number of runs of length ≥ 3 in the series of exceedance masks $M_{d,k}$ within the month m. For each calendar month, let μ_m^{base} and σ_m^{base} denote the reference period mean and standard deviation of $T_{m,k}^{Min}$.

3. **Standardised Anomalies:** The standardised anomaly for month m in year k is then given by

$$\delta_{m,k}^{T_{Min}} = \frac{T_{m,k}^{Min} - \mu_m^{\text{base}}}{\sigma_m^{\text{base}}}.$$

8 Hail

The Hail component is the fifth component of the ASSA Climate Index. It characterises the climatological frequency of large-scale weather conditions favourable to hail over South Africa. The precipitation, drought and temperature components each measure a directly observed (or near-observed) quantity, but no satellite- or reanalysis-derived "hail observation" product exists at the spatial and temporal resolution required for a national, publicly available climate index. The hail component is therefore built from a *convective-environment* proxy: the product of Convective Available Potential Energy (CAPE) and the 0–6 km above-ground-level (AGL) deep-layer wind shear (the change in the wind between near the surface and 6 km up), filtered to retain only the conditions that support hail growth and survival to the ground. Using CAPE, deep-layer shear and a freezing-level constraint together as a hail proxy is well established in the literature (8, 4, 23, 21, 24).

8.1 Input data

The hail proxy uses fields from the ECMWF ERA5 reanalysis (16), distributed via the CDS, at the native ERA5 horizontal resolution of $0.25°$ (≈ 25 km). ERA5 (rather than the AgERA5 used elsewhere in the index) is required because the proxy depends on upper-air winds, the 0°C isotherm height (the height at which the air temperature drops to 0°C) and convective indices, none of which AgERA5 provides. No CHIRPS-based alternative is produced, since CHIRPS contains no convective-environment variables. Throughout this section (x, y) denotes the ERA5 grid cell (longitude, latitude) and d the calendar date. All fields are the daily 12:00 UTC (14:00 South African Standard Time, SAST) snapshot, shortly before the mid-afternoon maximum of convective activity over the South African interior. The variables used are:

- CAPE, $\mathrm{CAPE}(d, x, y)$, in J kg^{-1}[6].
- 10 m horizontal wind components, $u_{10\,\mathrm{m}}(d, x, y)$, $v_{10\,\mathrm{m}}(d, x, y)$, in m s^{-1}.
- Pressure-level wind components u, v and geopotential Z at the levels that bracket the 6 km AGL height across South African terrain (500, 400 and 300 hPa), in m s^{-1} and m^2 s^{-2}. These determine the interpolated 6 km AGL wind (section 8.2).
- Geometric height of the 0°C isotherm above ground level, $z_{\mathrm{AGL}}(d, x, y)$, in m. This is the ERA5 "zero degree level" (`deg0l`)[7].

A static orography field $h_{\mathrm{orog}}(x, y)$ (surface elevation in metres above mean sea level (MSL) from the Shuttle Radar Topography Mission (SRTM) and WorldClim, resampled to the ERA5 grid) is used only to convert pressure-level geopotential heights to AGL for the shear interpolation (section 8.2), not to the freezing level.

8.2 0–6 km AGL deep-layer wind shear

Deep-layer wind shear is the change of the horizontal wind between two heights. It matters for hail because a storm growing in a strongly sheared environment is tilted, so its rain- and hail-laden downdraught falls clear of the warm, moist inflow feeding its updraught. The storm can then persist and, in the strongest cases, develop a rotating, long-lived updraught (a "supercell") that carries hailstones repeatedly through the sub-freezing growth zone. Shear thus governs how organised and long-lived a storm is, while CAPE governs how strong its updraught is. Large hail generally needs both, which is why the daily field uses their product.

The deep-layer shear here is the vector difference between the wind at 6 km AGL and the 10 m wind. The 6 km AGL wind is obtained by interpolating the pressure-level winds in height. Each pressure level's height above ground is

$$z_p(d, x, y) = \frac{Z_p(d, x, y)}{g} - h_{\mathrm{orog}}(x, y), \tag{1}$$

where Z_p is the geopotential at pressure level p and $g = 9.80665\,\mathrm{m\,s^{-2}}$, so Z_p/g is the geopotential height above MSL and subtracting the orography gives height above ground. The 6 km AGL wind $\mathbf{V}_{6\,\mathrm{km}}(d, x, y) = (u_{6\,\mathrm{km}}, v_{6\,\mathrm{km}})$ is then obtained by linear interpolation of the level winds to $z_p = 6\,000$ m, and the shear magnitude is

$$S(d, x, y) = \left\| \mathbf{V}_{6\,\mathrm{km}}(d, x, y) - \mathbf{V}_{10\,\mathrm{m}}(d, x, y) \right\| = \sqrt{(u_{6\,\mathrm{km}} - u_{10\,\mathrm{m}})^2 + (v_{6\,\mathrm{km}} - v_{10\,\mathrm{m}})^2}, \tag{2}$$

where $\mathbf{V}_{10\,\mathrm{m}}(d, x, y) = (u_{10\,\mathrm{m}}, v_{10\,\mathrm{m}})$ is the 10 m wind. The subscripts $6\,\mathrm{km}$ and $10\,\mathrm{m}$ give the two heights explicitly. The 0–6 km AGL layer is the standard deep-layer shear measure for hail and

[6] CAPE measures the energy available to drive vertical motion in a thunderstorm, i.e. how strongly a warm, moist surface parcel would accelerate upward once buoyant. Larger CAPE supports stronger updrafts and larger hailstones.

[7] The height of the 0°C level (the "freezing level") controls how large a hailstone can grow aloft and how much of it survives the fall through warmer air to the ground. ERA5 `deg0l` is already expressed above ground rather than above mean sea level, so it is used directly as z_{AGL} with no orography subtracted. Section 8.3 returns to this point.

supercell environments (8, 24). Interpolating to a fixed *height* above ground keeps the depth of the shear layer constant across terrain, which matters for a climatology that pools the elevated Highveld with the low-lying coast.

8.3 Daily hail-environment field and filter

The daily hail-environment field is the CAPE $\times$ shear product, retained only where the thermodynamic environment supports hail reaching the ground:

$$\Pi(d,x,y) = \begin{cases} \mathrm{CAPE}\,(d,x,y)\; S(d,x,y), & \text{if } 2\,400 \le z_{\mathrm{AGL}}(d,x,y) \le 4\,845 \text{ m}, \\ & \quad \mathrm{CAPE}(d,x,y) \ge 100 \text{ J kg}^{-1}, \\ & \quad S(d,x,y) \ge 1.5 \text{ m s}^{-1}, \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

The daily field Π has units of J kg^{-1} m s^{-1} (CAPE times a wind speed). Only ratios and exceedance frequencies of Π enter the index, so the unit itself never appears in the output. The freezing-level window $2\,400 \le z_{\mathrm{AGL}} \le 4\,845$ m AGL follows Punge et al. (24). The upper bound is a melting limit. Above $\approx 4\,845$ m ($\approx 16\,000$ ft) the fall path through above-freezing air is long enough that growing hailstones melt before reaching the surface. The lower bound is a growth limit. The "warm-cloud layer" is the part of the storm cloud between its base and the 0°C level, where liquid droplets form, and below $\approx 2\,400$ m this layer is too thin to supply the large supercooled droplets that hailstones grow by accreting, so hail growth is starved even though any hail that formed would survive the short fall. Punge et al. set the lower bound empirically: 9.5% of the satellite overshooting-top storm detections in their climatology occurred at a melting level below $2\,400$ m, against only 3.5% of microwave hail detections and 2.5% of insurance claims (24). The growth- and melt-limit readings given here are the physical interpretation of that empirical selection. The CAPE and shear floors exclude environments too weak in energy or organisation to produce hail.

Freezing-level convention. The freezing level z_{AGL} is the ERA5 `deg0l` field used directly. ERA5 already reports the zero-degree level as a height above ground, so no orography is subtracted.

In contrast with the precipitation and drought components, the daily field is used without any temporal accumulation. It is simply the 12:00 UTC value $\Pi(d,x,y)$ of equation (3), with d the calendar date. No spatial smoothing or storm-motion transform is applied (section 8.7). The CAPE that enters equation (3) is the *homogenised* CAPE defined next, in section 8.4.

8.4 CAPE observing-system homogenisation

ERA5 CAPE over the South African hail belt carries a spurious downward step across the satellite record. The *hail belt* means, here and in section 8.8, the fixed box 22–30°S, 24–32°E. It covers the interior summer-rainfall region where most damaging hail occurs (Gauteng, the Free State, North West and eastern Mpumalanga), and also admits parts of Limpopo, KwaZulu-Natal and the neighbouring countries. Statistics in this section and in section 8.8 are labelled by the area they are computed over: the hail belt, the $1\,812$ South African cells, or the full 56×76-cell ERA5 grid (the *ERA5 domain*, which also covers the surrounding ocean and neighbouring countries). The three are not interchangeable. Pre-2007 CAPE is systematically inflated relative to the years that follow, because the network of observations that the reanalysis draws on densified abruptly around then. Dense Global Positioning System (GPS) radio-occultation soundings entered from 2006, and the Infrared Atmospheric Sounding Interferometer (IASI) from 2007. Left uncorrected this imprints a false declining trend on the index and suppresses recent-month detection (section 8.8). The step is removed by a single multiplicative correction to CAPE, anchored to independent radiosonde observations and applied before the filter of equation (3). The procedure is as follows.

Rationale for anchoring to radiosondes. A radiosonde is an instrument package carried aloft by a weather balloon, measuring temperature, humidity and pressure directly as it rises through the atmosphere. One such profile is called a sounding. Because these are direct in-situ measurements that use no atmospheric model and assimilate no satellite data, they are unaffected by the 2006–07 change in the satellite observing system. That independence is what makes them a suitable reference for detecting an observing-system artefact in a model-based product such as ERA5. A sounding does not measure the same quantity as a gridded reanalysis CAPE, but it changes over time for physical reasons only, so a drift of ERA5 relative to the soundings is evidence about ERA5, not about the weather.

Reference observations. Radiosonde soundings are taken from the Integrated Global Radiosonde Archive, version 2 (IGRA v2; (14)), for the two hail-belt stations with continuous records spanning the index period: Irene/Pretoria (IGRA identifier `SFM00068263`; $-25.91°$, $28.21°$) and Bloemfontein (`SFM00068442`; $-29.10°$, $26.30°$), both 1979–present. ERA5 CAPE is sampled at the grid cell containing each station, on the same days. The correction derived from these two stations is applied to the whole South African grid, not only to the hail belt. The grounds for this, including a check against stations outside the hail belt, are set out under "Application to the whole grid" below.

Radiosonde CAPE. A sounding reports the air temperature and the dew point (the temperature at which the air would become saturated, a measure of its moisture) at a series of pressure levels from the ground upward. CAPE is not itself reported; it is calculated from this profile in four steps.

1. *Selecting soundings.* Only ascents launched between 10:00 and 14:00 UTC, around the routine 12 UTC launch, are used. An ascent is discarded if its lowest reported level lies above the 780 hPa pressure level (about 2.2 km above sea level), because the profile then misses the air near the ground, or if it ends before reaching 250 hPa (about 10 km), because it then misses the upper part of a storm's depth. Where a day has more than one usable ascent, the last is kept.
2. *Choosing the parcel.* Thunderstorms often draw their air not from the ground itself but from the warmest and moistest layer near it. The calculation therefore starts from the level, within the lowest 300 hPa (roughly the lowest 3–4 km) of the profile, with the highest equivalent potential temperature, a single quantity that combines a parcel's warmth and its moisture (computed following (7)). CAPE computed from this starting point is called *most-unstable* CAPE.
3. *Lifting the parcel.* The parcel is lifted in small steps, in calculation only, up to the 100 hPa level. While it is unsaturated it cools at the fixed rate of dry air as it expands. The level at which it becomes saturated and cloud begins to form is the lifting condensation level (LCL), found with the formula of (7). Above the LCL, condensation releases heat, so the parcel cools more slowly; the condensed water is assumed to fall out at once (a pseudo-adiabatic ascent).
4. *Adding up the buoyancy.* At each step the parcel is compared with the surrounding air that the sonde measured at the same pressure. Where the parcel is lighter than its surroundings it is buoyant and accelerates upward; CAPE adds up this buoyancy over every layer above the LCL where it is positive. Because moist air is lighter than dry air at the same temperature, the comparison uses the virtual temperature T_v, the temperature that dry air would need to have the same density. In symbols,

$$\mathrm{CAPE} = R_d \int_{p_{\mathrm{top}}}^{p_{\mathrm{LCL}}} \max\!\left(T_v^{\mathrm{parcel}} - T_v^{\mathrm{env}},\, 0\right) \mathrm{d}\ln p, \tag{4}$$

where $R_d = 287\,\mathrm{J\,kg^{-1}\,K^{-1}}$ is the gas constant for dry air, p_{LCL} is the pressure at the LCL and

$p_{\mathrm{top}} = 100\,\mathrm{hPa}$. The result is in $\mathrm{J\,kg^{-1}}$, the energy each kilogram of rising air would gain. Values above $9\,000\,\mathrm{J\,kg^{-1}}$ are physically implausible and are discarded as errors.

This most-unstable CAPE is computed by a different algorithm from the one ERA5 uses internally. As the next paragraph explains, the correction depends only on how the ERA5-to-sonde ratio changes over time, so this definitional difference does not affect it.

The ERA5-to-sonde ratio and its drift. The analysis is restricted to the austral hail season (October–March, with October–December assigned to the following calendar year to form a "season-year"). For each station s and season-year a, define the ERA5-to-sonde ratio

$$R_s(a) = \frac{\sum_{d \in \mathcal{P}_s(a)} \mathrm{CAPE}_s^{\mathrm{ERA5}}(d)}{\sum_{d \in \mathcal{P}_s(a)} \mathrm{CAPE}_s^{\mathrm{sonde}}(d)}, \tag{5}$$

where $\mathcal{P}_s(a)$ is the set of hail-season days in season-year a on which station s has both a sonde CAPE and a co-located ERA5 CAPE. In absolute terms $R_s(a)$ is well below one, typically about 0.15–0.35. A balloon samples a single point while ERA5 reports an average over a $\approx 25\,\mathrm{km}$ cell, and the two CAPE values come from different algorithms, so a large and roughly constant level offset between them is expected and is *not* what the correction uses. Each station's series is divided by its own mean over 2014–2026, $r_s(a) = R_s(a)/\overline{R_s}^{\,2014\text{–}26}$. The reference period starts in 2014, after the last of the major observing-system additions had entered the reanalysis. Normalising to this period aligns the whole record with the observing system that produces the current operational data, so no adjustment is needed from 2014 onward. The normalisation removes every time-invariant mismatch, both the point-versus-cell representativeness and the differing CAPE definitions, and leaves only the drift of ERA5 relative to the sondes over time. Since the sondes are an independent in-situ reference, that drift is attributed to the ERA5 observing system rather than to climate, and the soundings show no comparable trend (at Irene the two most recent of the five observing-system epochs defined in section 8.8 have the highest means; see Figure 1).

Pooling and the single national factor. The two normalised series $r_s(a)$ are averaged into one pooled index $b(a)$ (a season-year with no paired soundings at one station is represented by the other station alone). Writing $b_{1991\text{–}2006}$ for the mean of $b(a)$ over the season-years 1991–2006, the pooled ratio is $b_{1991\text{–}2006} \approx 1.46$ against 1.00 by construction over 2014–2026, i.e. pre-2007 ERA5 CAPE runs about 46% high relative to the 2014–2026 reference period. The two stations individually imply ≈ 1.50 (Irene) and ≈ 1.38 (Bloemfontein). Each 1991–2006 mean rests on roughly $2\,000$–$2\,400$ paired station-days (about 150 hail-season days per station-year, thinning in the sparsest years).

Application to the whole grid. The factor is estimated at the two hail-belt stations but applied to every grid cell. The estimation is confined to the hail belt because that is where CAPE is large enough for the ERA5-to-sonde ratio to be robust to how the season-years are sampled, and where the component carries its signal. Of the other six South African IGRA stations, only Port Elizabeth ($-33.99°$, $25.62°$) and Springbok ($-29.67°$, $17.88°$) have soundings in the 2014–2026 reference period; the remaining four cannot be normalised to it. Three observations support applying the hail-belt factor nationally. First, the cause is a change in the global observing system assimilated by ERA5, which has no regional character. Second, per-cell trends in raw ERA5 hail-season CAPE over 1991–2026 are negative in about 95% of South African cells (section 8.8), which shows that the decline is near-universal in sign, though not that it is uniform in size. Third, on the same definition, restricted to season-years with at least 30 paired days, Port Elizabeth and Springbok give pre-2007 ratios of about 1.8 and 2.2, larger than the 1.46 applied. These two ratios are volatile, however. Mean sounding CAPE is low at both sites (about 330 and $130\,\mathrm{J\,kg^{-1}}$ over the paired hail-season days, against about $1\,400\,\mathrm{J\,kg^{-1}}$ at Irene), so a few thinly sampled season-years can dominate the ratio. Admitting season-years with fewer than 30 paired days

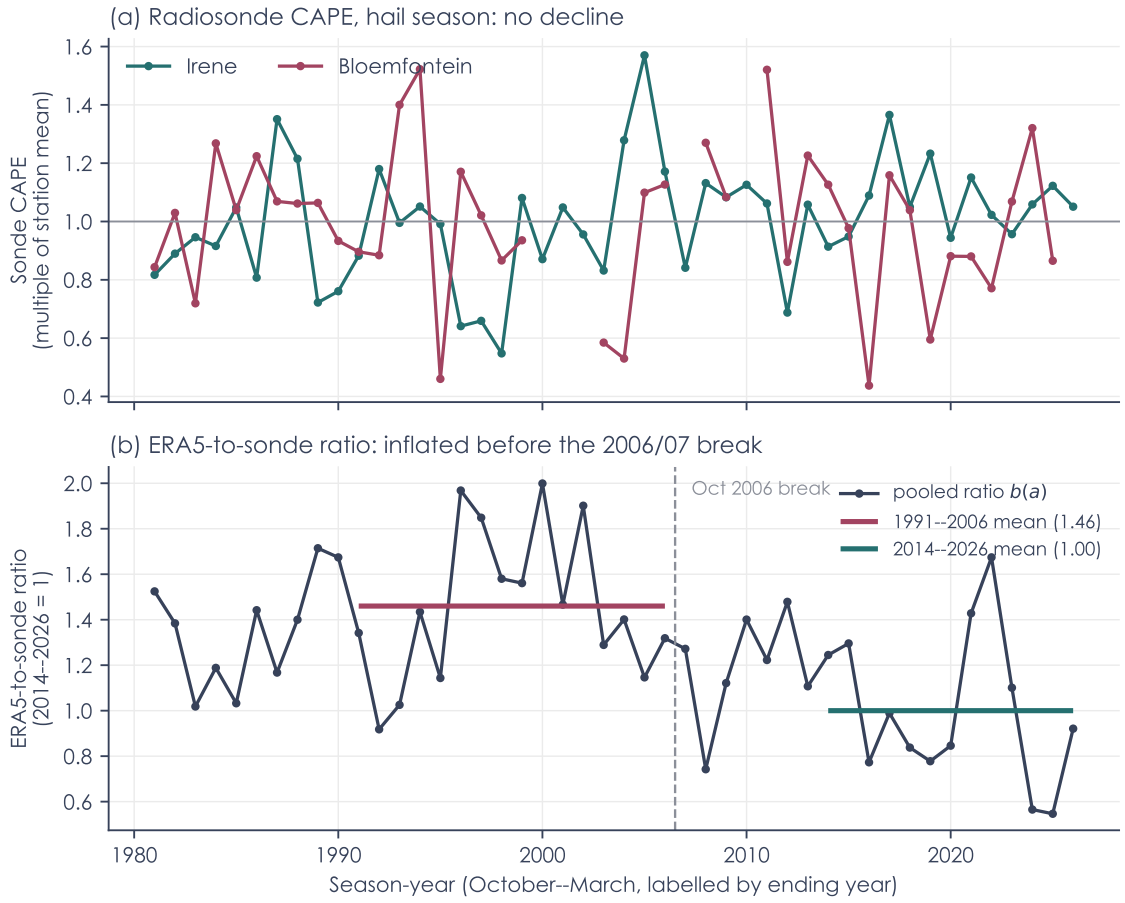


Figure 1: The radiosonde anchor for the CAPE homogenisation (section 8.4). (a) Hail-season mean radiosonde CAPE per season-year at Irene and Bloemfontein, each shown as a multiple of its own record mean; season-years with fewer than 30 paired soundings are omitted from this panel. The soundings show no long-term decline. (b) The pooled ERA5-to-sonde ratio $b(a)$, with each station normalised to its own 2014–2026 mean before pooling. The ratio runs $\approx 46\%$ high over the 1991–2006 estimation window and averages 1.00 by construction over 2014–2026; the dashed line marks the 1 October 2006 break used by the correction.

moves it to about 3.4 at Port Elizabeth and 0.8 at Springbok, whereas the hail-belt ratios change by less than 0.1. They are therefore not used to set the factor. What they do indicate is that, outside the hail belt, the national factor is more likely to under-correct than to over-correct, leaving any residual artefact in the direction of a spurious decline.

The correction (single step). Fitting a per-year curve would over-fit the sparse recent sonde counts and would need continual maintenance, so the bias is reduced to a single constant applied before a fixed break. All ERA5 CAPE before the 2006/07 hail season (every day earlier than 1 October 2006) is multiplied by

$$\kappa = \frac{1}{b_{\,1991\text{–}2006}} = \frac{1}{1.46} \approx 0.685, \tag{6}$$

and CAPE from October 2006 onward (including the operational live stream) is left unchanged. The break coincides with the observing-system densification above and with a cumulative-sum change-point identified independently on the hail-belt hail-season mean ERA5 CAPE series. It falls between the 2005/06 and 2006/07 seasons whether the series starts in 1979, 1981 or 1991,



with a mean of $276\,\mathrm{J\,kg^{-1}}$ before the break against $218\,\mathrm{J\,kg^{-1}}$ after it (series from 1981). The daily record begins in 1979, so the pre-1991 years lie outside the 1991–2006 estimation window. The same κ is applied to them, which the anchor suggests slightly over-deflates the 1980s (where the measured ratio is nearer 1.3–1.4). This affects only the delivered pre-1991 values, not the 1991–2020 climatology. The correction is applied to CAPE *before* the conditions of equation (3) are evaluated and before the CAPE $\times$ shear product is formed. The shear, the freezing level, the filter conditions and the entire downstream pipeline (section 8.5) are unchanged and are recomputed on the homogenised field, so the P95 climatology (section 8.5) and the base statistics all reflect the corrected CAPE.

Rationale and residual. Because the factor is anchored to independent radiosondes, it removes only the ERA5-minus-observation drift and preserves whatever genuine signal the soundings carry. Two residuals remain. First, the smaller inflation over 2007–2012 (a ratio of about 1.2) is left in place. Second, the anchor indicates that ERA5 CAPE runs somewhat *low* in the most recent season-years (the implied factor rises above one, to roughly 1.5 by 2025–26). This later adjustment is deliberately not applied, because those recent factors rest on the fewest paired soundings and are the least stable, so the production index is left marginally conservative in the most recent seasons rather than corrected with a fragile endpoint. The single step also leaves a small discontinuity in the historical series at the 2006/07 boundary, which is the appropriate location for it because the underlying observing-system change was itself a step.

8.5 Standardisation pipeline

The daily field $\Pi(d, x, y)$ is passed through the same standardisation pipeline as the precipitation, drought and temperature components, independently at each grid cell. Here d is a calendar date, m and k its calendar month and year, $\mathrm{DOY}(d) \in \{1, \ldots, 365\}$ its day-of-year computed from the (month, day) in a non-leap calendar (so 1 March is always day 60), and $\mathbf{1}[\cdot]$ the indicator function (equal to 1 when its condition holds and 0 otherwise).

1. **Reference Climatology:** For each day-of-year $d^* \in \{1, \ldots, 365\}$ (29 February handled as for precipitation), pool the daily values on all days lying within ± 15 calendar days of d^* in each base-period year (1991–2020) and take the 95$^{\text{th}}$ percentile, denoted π_{95} (the "P95 threshold"):

$$\pi_{95}(d^*, x, y) = \mathrm{percentile}_{95}\Big[\Pi(d, x, y) : \Delta_{365}\big(\mathrm{DOY}(d), d^*\big) \le 15\Big], \tag{7}$$

   where $\Delta_{365}(i, j) = \min\{|i - j|,\, 365 - |i - j|\}$ is distance around the 365-day year, so the window wraps across the year boundary. As for precipitation, day-of-years in early January and late December therefore draw a few values from just outside the base period (December 1990 and January 2021), giving $31 \times 30 = 930$ pooled values per day-of-year.

2. **Frequency Calculation:** For each month m of year k, the cell exceedance frequency $H_{m,k}$ is the fraction of that month's days on which Π exceeds the day-of-year P95 threshold:

$$H_{m,k}(x, y) = \frac{1}{|\mathcal{D}_{m,k}|} \sum_{d \in \mathcal{D}_{m,k}} \mathbf{1}\Big[\Pi(d, x, y) > \pi_{95}\big(\mathrm{DOY}(d), x, y\big)\Big], \tag{8}$$

   where $\mathcal{D}_{m,k}$ is the set of calendar days in month m of year k and $|\mathcal{D}_{m,k}|$ its number of days.

3. **Standardised Anomalies:** Let $\mu_m^{\mathrm{base}}(x, y)$ and $\sigma_m^{\mathrm{base}}(x, y)$ be the mean and standard deviation of $H_{m,k}(x, y)$ over the base-period years $k = 1991, \ldots, 2020$ (the standard deviation in its sample form, with denominator $n - 1$). The standardised anomaly is

$$\delta^H_{m,k}(x, y) = \frac{H_{m,k}(x, y) - \mu_m^{\mathrm{base}}(x, y)}{\sigma_m^{\mathrm{base}}(x, y)}. \tag{9}$$

Where σ_m^{base} is effectively zero (below 0.001, a value the discrete exceedance frequencies can reach without being exactly zero), the numerator $H_{m,k} - \mu_m^{\text{base}}$ is retained but divided by the standard deviation pooled over all base-period months of the same season (DJF, MAM, JJA, SON) in place of σ_m^{base}, exactly as for precipitation.

Off-season behaviour of the threshold. Because the filter of equation (3) sets Π to zero on most days outside the convective season, the day-of-year P95 is itself exactly zero for 43% of all (day-of-year, cell) combinations (67% in April–September, 19% in October–March). Where the threshold is zero the exceedance criterion $\Pi > \pi_{95}$ reduces to counting any day whose environment satisfies the filter at all, so in the off-season the index counts hail-favourable days by their occurrence, while in the core season it counts days exceeding the local seasonal P95. Both are sensible anomaly semantics, since a hail-favourable environment in mid-winter is anomalous by construction. The regime switch is noted here for transparency.

The monthly grid-cell anomaly $\delta_{m,k}^H$ is computed on the native ERA5 grid over South Africa (56×76 cells, ≈ 25 km) and spatially aggregated to districts, provinces and national level by the same arithmetic-mean convention as the other components.

8.6 Event-based validation

The methodology was evaluated against fifteen documented severe-hail events (Table 1). Two complementary diagnostics are reported for each event, because they answer different questions:

- the **monthly standardised anomaly** $\delta_{m,k}^H$ at the event cell, which is the actual index output. A "hit" is $\delta_{m,k}^H > +1$;
- the **event-day ratio**, the daily field on the event day divided by that cell's day-of-year P95, which tests whether the daily field flags the storm day itself as a top-5% (hail-favourable) environment.

For robustness against grid-versus-event mis-registration (the reported event centroid need not coincide with the strongest ERA5 grid cell), both are taken as the maximum over a 5×5 cell box centred on the event.

Event	Date	Location (lat, lon)	Monthly δ^H	Event-day ratio
Bloemfontein	2001-11-20	$-29.12, 26.21$	$+2.61$	3.32
Potchefstroom	2004-10-27	$-26.71, 27.10$	$+0.25$	1.71
Klerksdorp	2006-11-13	$-26.85, 26.67$	$+0.42$	1.11
Botshabelo	2008-11-10	$-29.27, 26.74$	$+2.18$	2.39
Gillitts (KZN)	2008-11-22	$-29.82, 30.86$	$+1.13$	0.96
Nkomazi	2012-10-18	$-25.62, 31.30$	$+1.93$	4.33
Boksburg	2012-10-22	$-26.21, 28.26$	$+4.20$	1.41
Pretoria (Nov 2012)	2012-11-08	$-25.75, 28.19$	$+2.72$	1.65
Pretoria (Nov 2013)	2013-11-28	$-25.75, 28.19$	$+1.62$	1.84
Tshwane District	2020-11-20	$-25.75, 28.19$	$+2.55$	1.71
Centurion	2024-02-21	$-25.85, 28.19$	$+5.06$	0.32
Lichtenburg	2025-10-19	$-26.15, 26.16$	$+0.59$	1.40
Erasmusrand	2025-11-22	$-25.79, 28.27$	$+1.44$	0.42
Welbekend	2025-11-30	$-26.42, 29.04$	$+1.71$	1.52
Pretoria (Dec 2025)	2025-12-01	$-25.75, 28.19$	$+2.41$	2.03
Hits			12/15	12/15

Table 1: Validation events and production-method scores, on the CAPE-homogenised production index (section 8.4). "Monthly δ^H" is the index output (hit if $> +1$). "Event-day ratio" is the event-day field divided by the cell P95 (top-5% environment if ≥ 1). Both are 5×5 box maxima centred on the event.

Base rates and skill. Because the output is a standardised anomaly, the hit counts must be read against the climatological frequency of each criterion, its base rate. Each base rate is measured at the fifteen event locations over the full 1979–2026 record, restricted to each event's calendar month so that the comparison is against the same time of year. It is the observed proportion of months, or of days, that meet the criterion. Multiplying it by fifteen gives the number of hits that fifteen randomly chosen months, or days, would be expected to produce. A 5×5-box monthly anomaly above $+1$ occurs in 39.5% of months, because hail-belt cells in the hail season are frequently anomalous, so fifteen random months would be expected to yield about 5.9 hits. The event-day criterion is more selective. A 5×5-box daily ratio ≥ 1 occurs on 13.2% of days, so fifteen random days would be expected to yield about 2.0 hits. The observed $12/15$ on each criterion is meaningfully higher than these expectations, and by the wider margin on the event-day criterion. The event-day result is therefore the primary demonstration that the daily field discriminates genuine hail environments. The monthly anomaly, the index's actual output, should be read against its high base rate.

Two caveats qualify both comparisons. The first is spatial. Six of the fifteen events (Pretoria Nov 2012, Pretoria Nov 2013, Tshwane District, Centurion, Erasmusrand and Pretoria Dec 2025) fall within a single ~ 25 km Gauteng neighbourhood, so the validation is spatially Highveld-centric and the events are not fully independent, though their dates are well separated in time.

The second is temporal. Ten of the fifteen fall in 2012 or later, and the production component runs above its fixed 1991–2020 baseline over that epoch (section 8.8), so a hit is easier to come by then than the full-record rate implies. Recomputing both base rates over 2012–2026 alone raises the monthly rate from 39.5% to 52.5% and the event-day rate from 13.2% to 16.4%, which puts the expected counts at 7.9 and 2.5 rather than 5.9 and 2.0. The event-day margin is barely affected. The monthly margin narrows to 12 against 7.9, which reaches the same conclusion by a different route, that the event-day criterion carries the weight of the demonstration.

The daily field flags 12 of 15 storm days as top-5% environments, which supports the proxy as physically reasonable. The three misses have understood, resolution-limited causes:

- **Sub-grid or non-shear-driven event days.** Gillitts (a coastal KwaZulu-Natal (KZN) supercell) is not resolved by ERA5 at 25 km. The ERA5 environment on the day has ample CAPE but weak deep-layer shear at the event cell ($\approx 8\,\mathrm{m\,s^{-1}}$), so its event-day ratio stays just below 1. On the Erasmusrand day the ERA5 environment shows moderate CAPE but very weak deep-layer shear ($\approx 5\,\mathrm{m\,s^{-1}}$, the second-weakest of the fifteen event days after Centurion), so the CAPE $\times$ shear proxy stays low. The storm was evidently organised by a mechanism other than deep-layer shear, consistent with lift along the cold front crossing Gauteng that afternoon, and storms of that kind are outside what the proxy is built to detect. Centurion is the opposite case, with a weak storm-day environment but a strongly anomalous month ($\delta^H = +5.06$).
- **Isolated storms in quiet months.** Potchefstroom, Klerksdorp and Lichtenburg exceed P95 on the event day (event-day ratio > 1) but their monthly anomaly is modest, because each was a lone hail-favourable day in an otherwise unremarkable month. A monthly exceedance-frequency index is, by construction, weaker at flagging isolated single-storm events than month-scale activity.

The CAPE homogenisation (section 8.4) raised both scores by one event relative to the raw-CAPE index. In particular it lifted Pretoria November 2013 from a monthly miss ($+0.88$) to a hit ($+1.62$) by removing the inflated pre-2007 plateau from the baseline against which recent months are standardised.

8.7 Methods investigated and rejected

A range of alternative daily-field constructions was evaluated against the same fifteen events, each with its own climatology rebuilt over 1991–2020, so that no variant benefits from a mis-

matched baseline. The variants were evaluated on raw ERA5 CAPE, before the homogenisation of section 8.4 was introduced, and have not been re-run under it. The ranking is not expected to change, because the homogenisation is a single multiplicative factor on pre-2007 CAPE that scales every variant's daily field, and with it the rebuilt P95, in the same way, whereas the differences between variants arise from how much each enlarges the daily field relative to its own P95. The outcomes are summarised in Table 2. A single mechanism explains the results. Every transform that enlarges the daily field (a spatial, temporal or trajectory maximum) raises its own P95 reference in proportion, so the standardised skill (performance on the fifteen-event validation of section 8.6) does not improve, and for the strongest enlargements it degrades.

Variant	Monthly	Event-day	Assessment
0–6 km AGL shear *(selected)*	11/15	11/15	The reference recipe for this comparison, on raw CAPE. With the CAPE homogenisation of section 8.4 the production index scores 12/15 on both criteria.
Shear from a fixed pressure level (the 10 m to 500 hPa wind difference)	12/15	10/15	Skill is comparable, but the layer this shear measures becomes shallower over high terrain (500 hPa sits ≈ 4.3 km above the Highveld but ≈ 5.7 km above the coast), so cells in different regions measure different physical quantities. The 0–6 km AGL shear measures the same layer everywhere and reads the event days slightly better.
Spatial smoothing (replacing each cell's daily value by the largest value within three cells, ≈ 75 km, in every direction, a 7×7-cell box)	not fully scored		Ruled out by diagnostic before a full re-run. A spatial maximum inflates the field's own P95 climatology in proportion, and the 5×5 verification box already absorbs event-location uncertainty.
Afternoon maximum (hourly fields, taking each day's maximum over 12–18 UTC)	11/15	10/15	Sampling the afternoon peak enlarges both the daily value and its rebuilt P95. Event-day skill drops slightly and monthly is unchanged, for a substantial extra data cost.
Storm-motion adjustment (crediting each cell with the hourly field traced 1–3 h back along the 500 hPa wind, to allow for storms arriving from upstream)	9/15	9/15	The strongest enlargement of the daily field and the largest P95 inflation. Both scores are the worst of all variants.
Gradual freezing-level weighting (a linear ramp replacing the hard edges of the z_{AGL} window)	not fully scored		Evaluated only in an earlier configuration, where it showed no benefit on the event set. The hard window is retained for simplicity.

Table 2: Alternative daily-field constructions evaluated against the production recipe, each scored against its own rebuilt 1991–2020 P95 and all on raw CAPE. "Monthly" is the number of validation events with $\delta^{H}_{m,k} > +1$ and "Event-day" the number with an event-day ratio ≥ 1, both on the 5×5 box. Variants marked "not fully scored" were ruled out by the diagnostic stated in the last column without a full standardised re-run.

8.8 Long-term behaviour

Built on raw ERA5 CAPE, the component trends spuriously downward. The share of hail-belt cell-months with $\delta^{H} > +1$ falls from $\approx 16.5\%$ in 1991–2000 to $\approx 10\%$ in 2021–2026 (Figure 2b, grey bars). Removing this artefact is the purpose of the CAPE homogenisation of section 8.4.

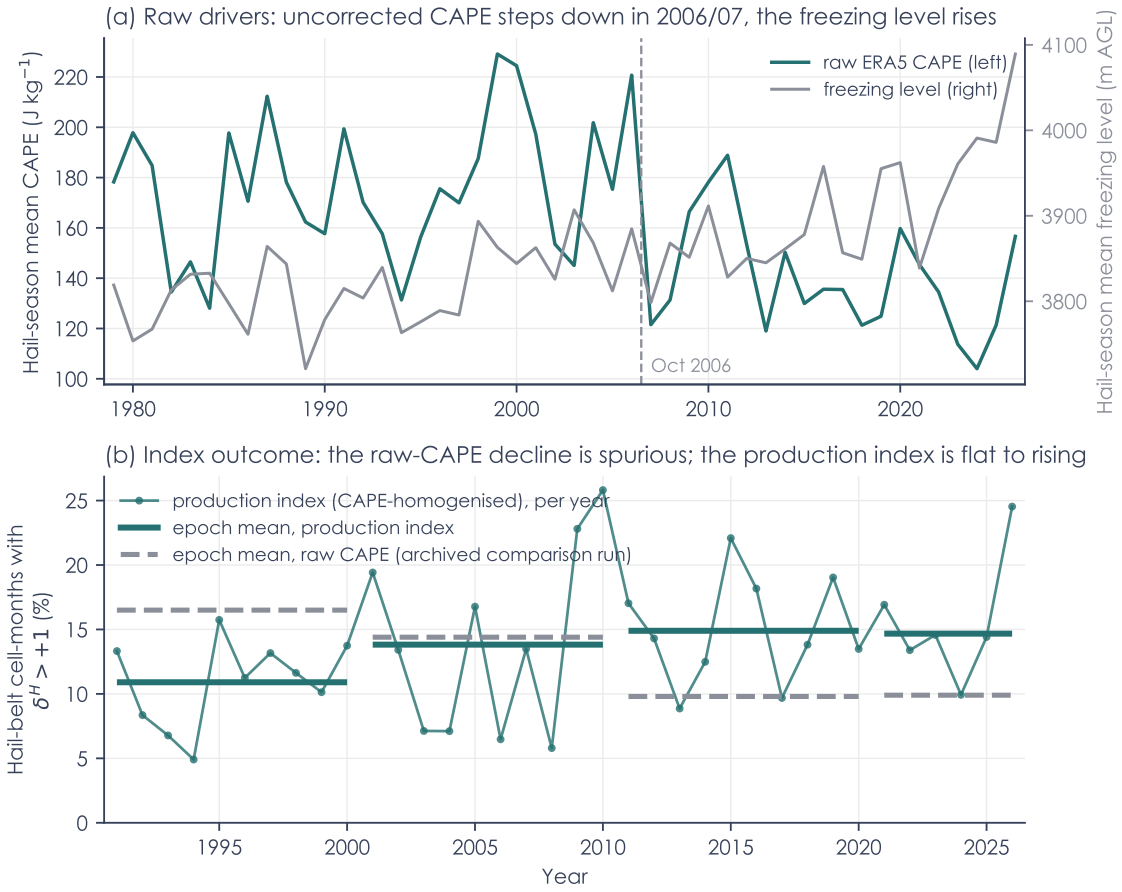


Figure 2: Long-term behaviour of the hail component. (a) Raw (uncorrected) ERA5 CAPE averaged over hail-season (October–March) months of each calendar year over the ERA5 domain (teal, left axis), and the corresponding mean freezing level (grey, right axis); the dashed line marks the October 2006 observing-system break. (b) Share of hail-belt cell-months (22–30°S, 24–32°E, all calendar months) with $\delta^H > +1$, for the production, CAPE-homogenised index per year (teal line, with epoch means as solid bars) and for the same pipeline run on uncorrected CAPE (epoch means only, grey dashed bars). The epochs are 1991–2000, 2001–2010, 2011–2020 and 2021–2026, and each epoch mean is the share pooled over all of that epoch's cell-months, so the partial year 2026 carries the weight of its months only. The uncorrected index falls across the epochs while the production index is approximately flat to rising.

Attribution over the raw drivers, taken as hail-season (October–March) means over the ERA5 domain and compared between 1991–2000 and 2021–2026, identifies two mechanisms, both visible in Figure 2a. The first is a rising freezing level (mean `deg0l` up ≈ 145 m, from $3\,817$ to $3\,962$ m, with the share of days above the $4\,845$ m melt-out bound growing from 6.0% to 9.6%), the expected thermodynamic response to warming (23). This genuinely reduces surface-reaching hail and is retained. The second is a decline in raw ERA5 hail-season CAPE (from 180 to 129 J kg^{-1}, or $\approx -28\%$, against near-neutral shear at 12.7 and 12.6 m s^{-1}). This carries most of the decline in the uncorrected index, and it is substantially a reanalysis artefact, on four lines of evidence:

1. it is step-like rather than gradual, with a change-point between the 2005/06 and 2006/07 seasons (section 8.4) coinciding with major additions to the ERA5 observing system (GPS radio-occultation from 2006 and IASI from 2007; Figure 2a);

2. it is near-universal in sign. Per-cell ordinary-least-squares trends in hail-season CAPE over 1991–2026 are negative in 95% of the $1\,812$ South African cells (94% of hail-belt cells and 98% of the ERA5 domain). The mean trend is $-18.0\,\mathrm{J\,kg^{-1}}$ per decade over the South African cells and $-22.9\,\mathrm{J\,kg^{-1}}$ per decade over the hail belt, the latter -9.4% per decade against a $243\,\mathrm{J\,kg^{-1}}$ hail-belt mean. A circulation-driven decline would carry regional structure, whereas one this widespread in sign points to a systemic observing-system shift. This measures how widespread the decline is, not how uniform its magnitude is;
3. it is not matched by the near-surface moisture that CAPE depends on. The AgERA5 hail-belt afternoon vapour pressure, computed as $e_s(T_{\max}) \cdot \mathrm{RH}_{15}/100$ and averaged over October–March, trends at only $-0.05\,\mathrm{hPa}$ per decade over 1991–2026 (-0.35% per decade on a $15.3\,\mathrm{hPa}$ mean), with RH_{15} falling 0.6 percentage points per decade while $T_{\max}$ rises $0.23\,^{\circ}\mathrm{C}$ per decade. Moisture this flat cannot sustain a CAPE decline of $\approx 9\%$ per decade. This is a consistency check rather than an independent test, because AgERA5 is itself derived from ERA5. It is informative because it follows a different variable chain within the same reanalysis, and the independent evidence is the next item. The decline also fits a wider pattern: Taszarek et al. find ERA5 CAPE decreasing across the tropics and subtropics while rawinsondes show CAPE increasing across the mid-latitudes, and advise caution with ERA5 CAPE trends where reliable long-term soundings are missing (27); and
4. the radiosonde anchor itself (section 8.4; Figure 1) shows no decline. Grouping the paired hail-season soundings by season-year into five epochs aligned with the observing-system changes (1981–1997, 1998–2006, 2007–2012, 2013–2018 and 2019–2026), and expressing each epoch mean as a multiple of the mean of the five, Irene reads 0.91, 1.03, 0.99, 1.04 and 1.04, and Bloemfontein 1.03, 0.87, 1.11, 1.05 and 0.94. Neither shows a sustained decline. Over the same epochs the ERA5-to-sonde ratio at the same cells and days falls by about 16% at Bloemfontein and 38% at Irene between the first epoch and the last.

After the homogenisation, the production share of hail-belt cell-months with $\delta^H > +1$ is approximately flat to rising across the four epochs 1991–2000, 2001–2010, 2011–2020 and 2021–2026, reading $\approx 11\%$, 14%, 15% and 15% (Figure 2b, teal bars). The last epoch is six years rather than ten, and its value sits fractionally below the third. This is consistent with the radiosonde record. It cannot be tested against observed hail, because no observational series is long enough: the satellite hail climatology for South Africa covers 2005–2018 (24), and insured-loss records are confounded by growth in exposure. The genuine freezing-level melt-out effect is preserved. The residual limitations are those set out in section 8.4 and are not repeated here.

Provincial detail. The national series averages provinces that behave differently, so it should not be read on its own. Figure 3 gives the provincial history. Between the first and last defined months of the 60-month mean the Free State moves $+0.49$, KwaZulu-Natal $+0.48$ and the Eastern Cape $+0.38$, against $+0.15$ nationally. Three provinces do not rise. Gauteng ends where it began at -0.01, North West moves -0.04, and Limpopo falls -0.26, the only substantial decline in the set.

Where the rise sits matters, because the validation episodes of section 8.6 are concentrated where it does not. Seven of the fifteen fall in Gauteng and three more in North West. Gauteng is the smallest province on the ERA5 grid at 25 cells, so its provincial mean is the noisiest in the figure, but the flat reading does not depend on that: taking Gauteng, North West and Mpumalanga together, 287 cells, still gives $+0.02$. The union of the 5×5 boxes around all fifteen episodes gives $+0.28$, above the national figure, because that set also holds the two Free State episodes and one in KwaZulu-Natal. The increase in hail-conducive environments is therefore genuine, and it is concentrated away from the Highveld interior where the documented damaging events cluster.

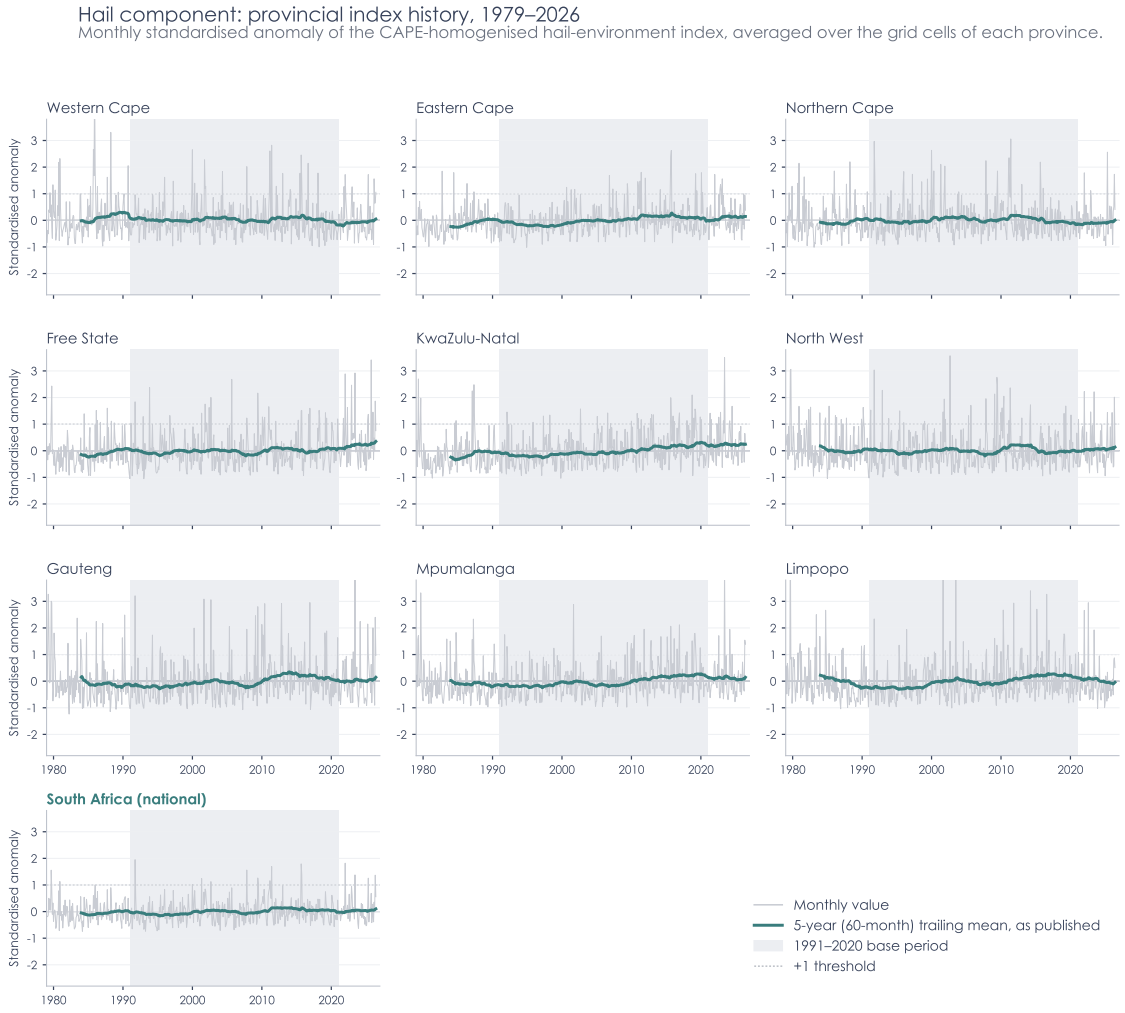


Figure 3: Provincial and national history of the hail component, 1979–2026, on the production CAPE-homogenised output. Thin line: monthly standardised anomaly $\delta^{H}_{m,k}$ averaged over the grid cells of the province (arithmetic mean of all cells whose centroids fall within it, the aggregation used throughout the index). Bold line: its 60-month trailing mean, the moving average published with the index, plotted at the last month of each window and therefore starting five years into the record. Shading: the 1991–2020 base period. Dotted line: $+1$. Gauteng holds 25 of the $1\,812$ South African cells at $0.25°$ and is correspondingly the noisiest panel.

9 Wildfire

The Wildfire component is the sixth component of the ASSA Climate Index. It characterises the climatological frequency of atmospheric conditions conducive to the ignition and spread of wildfire over South Africa. Products that instead count observed fires from satellite hot-spot detection (identifying fires by their thermal signature in satellite imagery) respond to land use, human ignition patterns and satellite viewing geometry as much as to weather. The Wildfire component therefore targets *fire-weather* environments, meaning the joint configuration of temperature, humidity, wind and antecedent moisture that determines whether an ignition, once started, can grow into a damaging fire. Fire-weather indices of this kind underpin most modern operational fire-danger rating systems (29, 26, 30).

The component is constructed as the elementwise daily maximum of two complementary fire-weather indices, each first normalised by its own day-of-year climatological mean. The two indices are the Hot-Dry-Windy Index (HDWI; (26)) and the Fire Weather Index (FWI) of the Canadian Forest Fire Weather Index System (29, 30), the latter computed with day-length tables adapted for the Southern Hemisphere. The HDWI captures short-lived days that are simultaneously hot, dry and windy, while the FWI captures longer-running fuel-drying episodes and carries a memory of antecedent dryness. Taking their normalised maximum lets the composite span both regimes. On any given day it takes the value of whichever index stands at the larger multiple of its own climatological mean. A single index is not used because the two regimes are physically distinct and neither index alone captures both (section 9.8).

9.1 Input data

The temperature, humidity and precipitation fields are taken from AgERA5 (version 2) (10) at the native AgERA5 grid of $0.1°$ (≈ 10 km), consistent with the other AgERA5-based components of the index. Throughout this section d denotes a calendar date, (x, y) the grid cell, and m and k the calendar month and year. The variables used are:

- the daily maximum 2 m air temperature $T_{\max}(d, x, y)$ in °C;
- the 2 m relative humidity, in %. Two fields are built from it, $\mathrm{RH}_{\mathrm{noon}}(d, x, y)$ at 12:00 South African Standard Time for the FWI channel and $\mathrm{RH}_{\mathrm{aft}}(d, x, y)$ at 14:00 SAST for the HDWI channel. Neither is a published AgERA5 field, for the following reason. AgERA5 supplies five fixed local hours (06h, 09h, 12h, 15h and 18h) rather than fixed UTC hours. It does not derive that local hour from longitude continuously, and it does not use civil time zones. It aggregates over eight longitudinal bands $45°$ wide, the first centred on the Greenwich meridian and running from $22.5°$W to $22.5°$E (10). South Africa spans 15 to $34°$E and straddles the boundary at $22.5°$E. West of it the labelled hour is UTC and east of it it is UTC$+3$, so one label denotes two clock times three hours apart inside one country. For the three fields this component uses, the clock times in SAST are as follows:

AgERA5 field	west of $22.5°$E	east of $22.5°$E
09h	11:00	08:00
12h	14:00	11:00
15h	17:00	14:00

  Read down either column and the same two clock times, 11:00 and 14:00 SAST, are available in both zones, from different fields. That is what makes a fixed clock hour reachable everywhere. $\mathrm{RH}_{\mathrm{noon}}$ interpolates between those two, weighting 11:00 by $2/3$ and 14:00 by $1/3$, which uses the 09h and 12h fields in the west and the 12h and 15h fields in the east. $\mathrm{RH}_{\mathrm{aft}}$ is simply the 14:00 SAST value, published as the 12h field in the west and the 15h field in the east, so it needs no interpolation. Where the two zones meet, their values for the same clock hour agree to well within the ordinary difference between neighbouring grid columns, so both constructed fields are continuous across the boundary, and so is the component built from them. The reason the two channels are built differently is given below. Both are near-surface (2 m) values rather than an upper-air (for example 850 hPa) value, because fire-weather indices depend on moisture near the ground;
- the daily precipitation total $P(d, x, y)$ in mm (the same AgERA5 field used by the precipitation and drought components).

The near-surface wind is taken from the ECMWF ERA5 reanalysis (16) at $0.25°$ (≈ 25 km) and regridded to the AgERA5 $0.1°$ grid by bilinear interpolation. The two channels take the wind at different times of day, because the indices they feed are defined differently:

- the FWI channel uses the 10 m wind speed at 10:00 UTC, $W_{\mathrm{noon}}(d, x, y)$. The Canadian system is defined on readings taken at noon local standard time (29), with the observation

hour set to the exact clock hour and not corrected for local solar time (19). Noon South African Standard Time is 10:00 UTC;

- the HDWI channel uses a 24-hour daily-mean 10 m wind $W(d, x, y)$, from the Copernicus post-processed daily-statistics product (`derived-era5-single-levels-daily-statistics`). That product distributes daily means of the two wind *components* over the UTC day, $\overline{u}_{10}$ and $\overline{v}_{10}$, which are combined as

$$W(d, x, y) = \sqrt{\overline{u}_{10}(d, x, y)^2 + \overline{v}_{10}(d, x, y)^2}. \tag{10}$$

Each channel therefore reads one precomputed field per day, so both are equally cheap to update, but the two are different kinds of quantity. The FWI channel's value is a true wind speed at an instant. The HDWI channel's is the magnitude of the daily-mean *vector*, which is not the same as the daily mean of the hourly *speed*. The two are equal only if the wind holds a constant direction all day, and otherwise the vector form is the smaller, by more as the direction turns further. It is thus a conservative measure of a day's wind forcing, and it is one mechanical reason the component under-reads the sub-daily sub-daily wind episodes discussed in section 9.7. Note also that this daily mean is taken over the UTC day, whereas the AgERA5 fields above are aggregated over their zone's local day. West of $22.5°$E the two windows coincide, and east of it the AgERA5 window is shifted three hours earlier. A daily maximum and a daily total are largely insensitive to a shift of that size, which is why neither the temperature nor the precipitation field steps at the boundary.

The choice of ERA5 over the higher-resolution ERA5-Land wind, and of these sampling times over an afternoon peak, is examined in section 9.8. In short, ERA5-Land is masked over the ocean and blanks coastal fire cells, whereas the ERA5 $0.25°$ field gives complete coastal coverage and is consistent with the ERA5 used by the hail component. Unlike the precipitation and drought components, this one has no CHIRPS-based counterpart, because CHIRPS supplies precipitation alone and the component also needs temperature, humidity and wind.

Why the two channels are built differently. The FWI channel has a target hour it must meet. The system it implements is defined on readings taken at noon local standard time (29), and no AgERA5 field falls there in either zone. The nearest available values are 11:00 and 14:00 SAST, so reaching noon requires interpolation.

The HDWI channel has no specified hour. Its authors define the index on the daily maximum over daytime steps rather than at a fixed time (26), so any single-hour sample is already a departure, examined in section 9.8. The hour is therefore open, and 14:00 SAST is both near the climatological peak of the vapour-pressure deficit and a value that each zone publishes exactly.

Record and spin-up. All fields begin on 1 January 1979 and run to the end date of the current release. The three FWI fuel-moisture codes are integrated forward from 1 January 1979, about twelve years before the start of the 1991–2020 baseline, so that the slowest code, the Drought Code, is fully equilibrated before the climatology window. Its drying timelag is 52 days (29), but it also carries moisture across a dry season, so a spin-up of several years rather than several months is required.

9.2 Hot-Dry-Windy Index (HDWI)

The Hot-Dry-Windy Index of (26) is the elementwise product of the vapour-pressure deficit (VPD) and the near-surface wind speed:

$$\mathrm{HDWI}(d, x, y) = \mathrm{VPD}(d, x, y) \cdot W(d, x, y), \tag{11}$$

with VPD in hPa and W in m s^{-1}. The vapour-pressure deficit is the gap between the amount of water vapour the air could hold when saturated and the amount it actually holds. It is the drying demand the air places on live and dead vegetation, and multiplying it by the wind speed, which sets how fast that demand is replenished over a fire, gives a hot-dry-windy intensity. It is computed from $T_{\max}$ and $\mathrm{RH}_{\mathrm{aft}}$ using the saturation vapour pressure $e_s(T)$ of (7), with T in °C:

$$e_s(T) = 6.112 \exp\!\left(\frac{17.67\,T}{T + 243.5}\right) \text{ hPa}, \tag{12}$$

$$\mathrm{VPD}(d,x,y) = e_s\big(T_{\max}(d,x,y)\big)\left(1 - \frac{\mathrm{RH}_{\mathrm{aft}}(d,x,y)}{100}\right), \tag{13}$$

where $e_s(T)$ is the maximum water-vapour pressure the air can sustain at temperature T (it rises roughly exponentially with T). Pairing the daily maximum temperature with the 14:00 SAST humidity (section 9.1) targets the part of the day when fire weather is most severe. The implementation uses the surface form of the index (2 m VPD and 10 m wind) rather than the lowest-50 hPa layer maximum of (26). In the standardised-anomaly framework the systematic offset between the two levels cancels, because each field is compared against its own climatology. What the surface form loses is the occasional day on which the strongest wind in the lowest few hundred metres sits well above the ground, where turbulence can still carry it down to a fire but a 10 m reading does not see it. Such episodes last hours rather than days, and the component already uses a 24-hour mean wind, so they fall outside what a monthly-frequency index sets out to resolve (the design envelope of section 9.7).

HDWI has units of hPa · m s^{-1} and is non-negative whenever $\mathrm{RH}_{\mathrm{aft}} \le 100\%$. AgERA5 caps relative humidity at 100%, so the deficit cannot go below zero.

9.3 Canadian Fire Weather Index (FWI)

The Fire Weather Index is the head of the six-code Canadian Forest Fire Weather Index System (29), an operational index adopted globally (for example by the EU Copernicus Emergency Management Service (CEMS) Global ECMWF Fire Forecast, GEFF (30)). The system is driven by four daily inputs and propagates three fuel-moisture state variables (FFMC, DMC and DC, defined below) forward in time. These are combined through two intermediate spread indices (ISI and BUI, also defined below) into the final scalar FWI. The four daily drivers used here are the daily maximum temperature $T_{\max}$, the clock-noon relative humidity $\mathrm{RH}_{\mathrm{noon}}$, the 10:00 UTC wind W_{noon} (converted to km h^{-1} for the FFMC and ISI), and the 24-hour precipitation total P. The system specifies noon readings of temperature, humidity and wind together with the rainfall accumulated over the preceding 24 hours. The wind and the humidity both meet it here, the wind as a 10:00 UTC reading and the humidity as an interpolation to that same instant (section 9.1). Precipitation is a 24-hour total accumulated over the local day rather than from noon to noon, which shifts a rain event by up to half a day relative to the code that responds to it. Temperature is the daily maximum rather than the noon value, a deliberate deviation discussed in section 9.4.

State variables. The three fuel-moisture codes represent fast-, medium- and slow-drying fuel layers:

FFMC (Fine Fuel Moisture Code). Moisture of surface litter, time scale ≈ 16 h. Sensitive to T, RH, W_{noon} and P.

DMC (Duff Moisture Code). Moisture of the loosely compacted organic layer (5–10 cm), time scale ≈ 12 days. Sensitive to T, RH and P. Uses a month-dependent day-length factor L_e.

DC (Drought Code). Moisture of the deep compact organic layer, time scale ≈ 52 days (29). Sensitive to T and P. Uses a month-dependent day-length factor L_f.

Each state variable $S \in \{\mathrm{FFMC}, \mathrm{DMC}, \mathrm{DC}\}$ is initialised at the standard value ($\mathrm{FFMC}_0 = 85$, $\mathrm{DMC}_0 = 6$, $\mathrm{DC}_0 = 15$; (29)) and updated each day by the deterministic recursion of (29), which applies the effect of rain first and then drying,

$$S(d, x, y) = f_S\Big(S(d-1, x, y),\, T_{\max},\, \mathrm{RH}_{\mathrm{noon}},\, W_{\mathrm{noon}},\, P,\, L(m(d))\Big), \tag{14}$$

where $L(m)$ is the relevant month-of-year day-length factor (L_e for DMC, L_f for DC).

Southern-Hemisphere day-length tables. Day length enters these codes because fuel dries during daylight hours, so the daily drying rate is scaled by a month-dependent day-length factor. The published Canadian factors are calibrated for boreal Northern-Hemisphere latitudes and must be adapted for South Africa. The two factors are different quantities and are treated separately:

- For the **Drought Code**, L_f is the published Southern-Hemisphere row of (19), Table A3.2, which is the Canadian row with the seasons reversed,

  $$L_f = (6.4,\, 5.0,\, 2.4,\, 0.4,\, -1.6,\, -1.6,\, -1.6,\, -1.6,\, -1.6,\, 0.9,\, 3.8,\, 5.8)$$

  for months January $\rightarrow$ December. L_f is an additive seasonal adjustment in the Drought Code's evaporation term (not itself a length of time). Its negative winter values simply mean that in the short, cool days of winter the deep layer is treated as retaining rather than losing moisture.

- For the **Duff Moisture Code**, L_e is the published $20°$ S row of (19), Table A3.1 (in hours),

  $$L_e = (10.1,\, 9.6,\, 9.1,\, 8.5,\, 8.1,\, 7.8,\, 7.9,\, 8.3,\, 8.9,\, 9.4,\, 9.9,\, 10.2),$$

  for months January $\rightarrow$ December. The factor is an *effective* day length, defined as the theoretical mean day length less three hours, rather than the astronomical photoperiod. That source tabulates it by reference latitude and states which row applies where. Its general rule assigns the $20°$ S row to latitudes 10–$30°$ S and the $40°$ S row to 30–$55°$ S. It then names South Africa as an explicit exception, prescribing the $20°$ S values for the whole country, and notes that only the most southerly portion reaches $35°$ S. The whole-country $20°$ S row is used here as prescribed.

Spread indices and final FWI. The Initial Spread Index combines the Fine Fuel Moisture Code with wind,

$$\mathrm{ISI} = 0.208\, f_W(W_{\mathrm{kmh}})\, f_F(\mathrm{FFMC}), \qquad f_W(W_{\mathrm{kmh}}) = \exp(0.05039\, W_{\mathrm{kmh}}), \tag{15}$$

where $W_{\mathrm{kmh}} = 3.6\, W_{\mathrm{noon}}$ is the wind in km h^{-1}. The constant 0.05039 is defined for wind in km h^{-1}, not m s^{-1}. The Build-Up Index combines the two deeper codes as $\mathrm{BUI} = g(\mathrm{DMC}, \mathrm{DC})$, and the final index is $\mathrm{FWI} = h(\mathrm{ISI}, \mathrm{BUI})$, where f_F, g and h are the standard functions of (29). FWI is non-negative and dimensionless.

9.4 Validation of the FWI implementation against GEFF

Because the FWI is computed in-house, its implementation was cross-validated against the canonical GEFF reanalysis (30), cell by cell for 1985, 2005, 2017 and 2024 over the whole model grid, which extends beyond South Africa into the neighbouring countries. The four years are spread across the record, so the result is not a property of the recent era alone. Agreement is high for every code (Table 3): median per-cell correlations run from 0.91 to 0.99, and at least 98% of cells exceed $r = 0.8$ for every code. That is consistent with the recursions being faithfully implemented.

Both this component and GEFF read the wind at local noon, so the wind contributes little to the residual difference. The humidity is also read at noon on both sides (section 9.1), so two differences remain. The temperature driver is the daily maximum rather than a noon value, and the temperature, humidity and precipitation fields are AgERA5, which is derived from ERA5 but downscaled to $0.1°$ with an elevation correction, whereas the reference product is driven by ERA5 directly on its own grid. That product is the CEMS fire-danger reanalysis `cems-fire-historical-v1`, which runs the GEFF model of (30) on ERA5. The two causes are not separated here, because the correlations above already establish that the recursions are implemented correctly, and because both act as level offsets. In the standardised-anomaly framework (section 9.6) each cell is compared against its own climatology, so a level bias that is steady through the year shifts the baseline and the monthly value together and largely cancels.

The daily maximum temperature is retained deliberately. AgERA5 supplies it at $0.1°$ with an elevation correction, and temperature is the most terrain-sensitive of the four drivers, so that resolution is worth more over the escarpment and the Cape mountains than the timing refinement would be. Van Wagner's own testing bounds the size of a timing effect of this kind. He compared the Fine Fuel Moisture Code worked from noon observations against an hourly computation carried forward to 1600 h, the afternoon peak the code is meant to stand for. The gap was judged too small to warrant delaying the official observation hour (29). It widens with latitude because the daily maximum temperature arrives progressively later, so at South African latitudes it should sit at or below the author's tested range.

Agreement with GEFF tests this implementation of the equations. It is not a target. The component's threshold is computed on its own composite, so it never needs to reproduce GEFF's values.

Code	Median per-cell r	Cells with $r > 0.8$	Median bias (ours − GEFF)
DC	0.990	99.7%	$+34.8$
DMC	0.967	98.8%	$+1.5$
FWI	0.957	100.0%	-0.2
ISI	0.942	99.7%	-0.3
FFMC	0.909	99.8%	-0.1

Table 3: Cell-by-cell agreement of the in-house FWI System codes with the canonical GEFF reanalysis (30) over the model grid, pooled over 1985, 2005, 2017 and 2024, on the production configuration. Agreement is uniform across cells. In the poorest tenth the correlation is still 0.97 for DC, 0.93 for FWI, 0.92 for DMC, 0.91 for ISI and 0.88 for FFMC. The DMC row depends on the day-length table of section 9.3.

9.5 Normalised-maximum composite

The HDWI and FWI carry different units and different climatological amplitudes, and each has a pronounced but distinct intra-annual cycle (HDWI peaks in the dry, windy months, while the FWI cycle is governed by accumulated drying). A naive elementwise maximum $\max(\mathrm{HDWI}, \mathrm{FWI})$ would therefore be dominated by whichever index happens to have larger absolute values, not by whichever is the more anomalous. To address this, each index is first divided by its own day-of-year climatological mean and only then combined.

For each cell (x, y) and day-of-year $d^* \in \{1, \ldots, 365\}$, let

$$\overline{\mathrm{HDWI}}(d^*, x, y) = \frac{1}{|\mathcal{B}(d^*)|} \sum_{d \in \mathcal{B}(d^*)} \mathrm{HDWI}(d, x, y), \tag{16}$$

$$\overline{\mathrm{FWI}}(d^*, x, y) = \frac{1}{|\mathcal{B}(d^*)|} \sum_{d \in \mathcal{B}(d^*)} \mathrm{FWI}(d, x, y), \tag{17}$$

where $\mathcal{B}(d^*)$ is the set of baseline (1991–2020) days whose day-of-year lies within ± 15 days of d^*, the same 31-day window used elsewhere in the index. Day-of-year is taken on the 365-day convention used throughout, in which each slot is identified by its (month, day) in a non-leap year and 29 February is interpolated from 28 February and 1 March, so the channel means and the percentile climatology of section 9.6 index dates identically. The daily fire-weather composite is the elementwise daily maximum of the two day-of-year-normalised ratios:

$$F(d,x,y) = \max\left\{ \frac{\mathrm{HDWI}(d,x,y)}{\overline{\mathrm{HDWI}}\big(\mathrm{DOY}(d),x,y\big)},\ \frac{\mathrm{FWI}(d,x,y)}{\overline{\mathrm{FWI}}\big(\mathrm{DOY}(d),x,y\big)} \right\}. \tag{18}$$

Because the two ratios are non-negative, share the same dimensionless scale (a multiple of the local climatological mean) and each fluctuate around 1, their maximum selects, per cell and day, whichever channel is the larger multiple of its own climatological mean. The two ratios have different dispersions, so this is a pragmatic scale-free combination rather than a formal ranking of which channel is the more unusual. Because the percentile threshold of section 9.6 is computed on the composite itself, the baseline exceedance rate does not depend on how the two channels are weighted against each other. Where a channel's day-of-year mean is zero (for example the FWI in a cell that is permanently wet in that part of the year), that channel's ratio is undefined and the composite reduces to the other channel. A day is treated as missing only if both means are zero. In practice both means are strictly positive over South Africa.

9.6 Standardisation pipeline

The daily composite $F(d,x,y)$ is passed through the same standardisation pipeline as the precipitation, drought and hail components, independently at each grid cell. As in the hail pipeline (section 8.5), $\mathrm{DOY}(d) \in \{1,\dots,365\}$ is the day-of-year of date d in a non-leap calendar, $\mathcal{D}_{m,k}$ the set of calendar days in month m of year k, and $\mathbf{1}[\cdot]$ the indicator function (1 when its condition holds, 0 otherwise).

1. **Reference Climatology:** For each day-of-year $d^* \in \{1,\dots,365\}$, pool the daily composite values on all days lying within ± 15 calendar days of d^* in each base-period year (1991–2020) and take the 95$^{\text{th}}$ percentile, denoted F_{95} (the "P95 threshold"):

$$F_{95}(d^*,x,y) = \mathrm{percentile}_{95}\Big[F(d,x,y) : \Delta_{365}\big(\mathrm{DOY}(d),d^*\big) \le 15\Big], \tag{19}$$

   where $\Delta_{365}(i,j) = \min\{|i-j|,\ 365-|i-j|\}$ is distance around the 365-day year, so the window wraps across the year boundary and, as for precipitation, day-of-years in early January and late December draw a few values from just outside the base period (December 1990 and January 2021). The 29$^{\text{th}}$ of February is handled as for precipitation, by averaging the percentiles for day-of-year 59 and 60.

2. **Frequency Calculation:** For each month m of year k, the cell exceedance frequency $F_{m,k}$ is the fraction of that month's days on which the composite exceeds the day-of-year P95 threshold:

$$F_{m,k}(x,y) = \frac{1}{|\mathcal{D}_{m,k}|} \sum_{d \in \mathcal{D}_{m,k}} \mathbf{1}\Big[F(d,x,y) > F_{95}\big(\mathrm{DOY}(d),x,y\big)\Big]. \tag{20}$$

   Over the base period this averages 5.1% of days per month, as the 95$^{\text{th}}$-percentile threshold implies. The average is taken over the cells where the composite is defined. AgERA5 is a land product, so the 22.9% of the rectangular grid that falls over the ocean carries no value at any date and enters no aggregate of the index.

3. **Standardised Anomalies:** Let $\mu_m^{\mathrm{base}}(x,y)$ and $\sigma_m^{\mathrm{base}}(x,y)$ be the mean and standard deviation of $F_{m,k}(x,y)$ over the base-period years $k = 1991,\dots,2020$ (the standard deviation in



its sample form, with denominator $n-1$). The standardised anomaly is

$$\delta^{F}_{m,k}(x,y) = \frac{F_{m,k}(x,y) - \mu^{\text{base}}_{m}(x,y)}{\sigma^{\text{base}}_{m}(x,y)}. \tag{21}$$

Where σ^{base}_{m} is effectively zero (below 0.001), the numerator $F_{m,k} - \mu^{\text{base}}_{m}$ is retained but divided by the standard deviation pooled over all base-period months of the same season (DJF, MAM, JJA, SON) in place of σ^{base}_{m}, exactly as for precipitation.

The day-of-year channel means of equations (16)–(17) are part of the saved baseline climatology (alongside the P95 threshold and the base mean and standard deviation), so that seasonal updates standardise new months against exactly the same reference as the launch.

9.7 Event-based validation

The component was evaluated against thirteen documented severe-wildfire episodes in South Africa, listed in Table 4. Each is scored at the spatial and temporal scale on which it actually occurred, because the episodes are not all of one kind. Some were single fires at a known location. One was a complex of several fires spanning two provinces. One was a two-day front across half a province. Three were whole fire seasons of five to seven months. In the two North West seasons, hundreds of thousands of ha reportedly burned in thousands of separate burn scars. A single test geometry cannot represent all of these, and applying one would misrepresent most of them.

Four scales are used for the event validation as set out below.

- **Point.** A single fire with a known location. The score is the maximum monthly anomaly $\delta^{F}_{m,k}$ over a 5×5 cell box (≈ 55 km) centred on it, the box absorbing the mismatch between a reported centroid and the strongest model cell.
- **District.** A fire complex too broad for a box but narrower than a province. The score is the mean anomaly for the event month over the district municipalities the complex occupied, on the same arithmetic-mean aggregation the delivered product uses. Where a complex spans more than one district, all of them are included.
- **Province.** A single synoptic episode affecting a whole province. The score is the province-mean anomaly for the event month.
- **Season.** A reported fire season rather than a fire. The score is the province-mean anomaly averaged over the months of the reported season window.

A hit is a score above $+1$ at every scale.

The strongest single result in the set is the Mpumalanga 2024 season, the highest of the 47 seasons on record for that province, with July 2024 alone reaching $+5.8$.

What the episodes do and do not test. Both channels are exercised, and which one carries an episode follows its character. Over the eight point episodes the FWI channel drives 76% of the exceedance cell-days inside the scoring box, rising to 94% at Knysna, the clearest drought-driven fire in the set. Over the district, provincial and season episodes the HDWI channel drives 72 to 95%, because those aggregate many separate windy dry spells rather than one accumulated drying signal. The fuel-moisture chain is therefore tested here rather than merely carried along.

The four episodes not flagged. St Francis Bay in November 2012 was a structural fire in which 76 buildings reportedly burned, with no vegetation extent reported, and its cell meets the hit criterion in 44.7% of Novembers on record, the loosest test in the set. Wupperthal in December 2018 reportedly destroyed 52 houses in a Cederberg village that is sub-grid at $0.1°$, and the

Episode	Period	Scale and region	Score	Hit
St Francis Bay	Nov 2012	point	-0.39	
Knysna	Jun 2017	point	$+4.21$	✓
Wupperthal	Dec 2018	point	$+0.41$	
Betty's Bay	Jan 2019	point	$+1.37$	✓
Table Mountain	Apr 2021	point	$+1.74$	✓
Kluitjieskraal	Jan 2024	point	$+1.51$	✓
Overstrand	Jan 2026	point	$+4.42$	✓
Mossel Bay	Jan 2026	point	$+1.83$	✓
Hertzogville complex	Oct 2020	districts, Lejweleputswa and Dr Ruth Segomotsi Mompati	$+1.08$	✓
KwaZulu-Natal front	Jul 2024	province, KwaZulu-Natal	$+4.10$	✓
North West season	May–Sep 2021	season, North West	$+0.20$	
North West season	May–Oct 2022	season, North West	$+0.54$	
Mpumalanga season	Apr–Oct 2024	season, Mpumalanga	$+1.32$	✓
Hits				9/13
Disaggregation of the two North West seasons, not counted above				
2021 season	Jun 2021	district, Bojanala	$+2.08$	(✓)
2022 season	Oct 2022	province, North West	$+1.67$	(✓)

Table 4: Validation episodes and production scores. Each is scored at the scale on which it occurred (see text), and a hit is a score above $+1$. The Overstrand entry covers the Stanford, Pearly Beach and Gansbaai fires of January 2026 as one event, scored over the union of their three 5×5 boxes, which overlap to 45 cells. The Hertzogville entry spans two districts because the complex reached Christiana in North West as well as Hertzogville and Hoopstad in the Free State. The final two rows are not episodes and are not counted in the hits. Each breaks a North West season down to the month the component reads most strongly, shown because both seasons are misses as provincial averages while the month inside them is not. The two are given at different scales because the seasons differ that way. In 2021 the signal was confined to Bojanala, the district holding 41% of that season's burnt area, and in 2022 it was province-wide.

surrounding month was not anomalous. Both sit inside the design envelope described below rather than contradicting it.

The two North West seasons score $+0.20$ and $+0.54$ as provincial averages over five and six months. Neither is a single anomalous event. In both, the component reads some months well above normal and others below, and the disaggregation below separates them.

Design envelope. The component is a fire-weather *month* index. It characterises whether a month offered anomalously frequent fire-conducive weather, the quantity most relevant to month-aggregated insured losses, and it is deliberately not built to resolve individual fire days. Its purpose is to track how the frequency of loss-conducive conditions shifts relative to a fixed baseline, for use in trend, pricing-cycle and scenario work. Rating or detecting a single fire is a different task, served by event-level tools.

That distinction is visible in the three point episodes the daily composite does not flag on the fire day itself, where it reaches only 0.41 to 0.82 of its own P95 at the event cell. The three are not one mechanism. Wupperthal is hot and dry on an ordinary daily wind. Betty's Bay is of median warmth on the strongest daily wind of any episode in the set. St Francis Bay is hot but humid, at an above-median 82% at noon, and the fire spread from roof to roof through thatch in the late afternoon instead of through vegetation (strictly not a wildfire event). What they share is that the decisive forcing lasted hours rather than a day, and neither channel's wind can represent a spike of that length. The FWI channel reads a single noon value and the HDWI channel a 24-hour mean. A month containing more such days still registers through the month's broader fire-weather environment even where the single afternoon spike is averaged away. The hail component meets the same limit, under-reading isolated single-storm days in otherwise quiet



months (section 8.6). A monthly-frequency design is weaker at flagging an isolated day than a month of elevated activity, in both components.

Spatial resolution sets a second limit at the same episodes. The wind is an ERA5 $0.25°$ field, so the Cape Peninsula is sub-grid and the component cannot resolve the wind contrast across Table Mountain. The peninsula carries substantial insured exposure, so this is a real limitation rather than a technicality, and it is not removed by the 5×5 box, which absorbs mis-registration but cannot manufacture detail the wind field does not hold.

Matching the scale to the extent matters as much as matching the month. The Hertzogville event of October 2020 burned about $100\,000$ ha across Hertzogville and Hoopstad in the Free State and Christiana in North West, so it occupied two district municipalities either side of a provincial boundary. Scored over both it reads $+1.08$. Scored over the Free State district alone it reads $+0.27$ and would be recorded as a miss, because the anomaly sat in the North West half, where 87% of cells exceeded $+1$ that month against 21% on the Free State side.

The season episodes show the same design working in its intended direction. A season score is an average over five to seven months, so a month that stands out is diluted. Disaggregating recovers it, as the last two rows of Table 4 show. The North West 2022 season scores $+0.54$ as a provincial average, but October 2022 alone reaches $+1.67$, the third highest of the 47 Octobers on record, and in the two western districts 82 and 92% of cells exceed $+1$ that month against climatological rates of 14 and 16%. The 2021 season is genuinely quieter as a provincial average, ranking 13th of 47, but the disaggregation recovers the district and month it reads most strongly. Bojanala held 41.3% of the province's burnt area over 1 May to 31 August 2021, and in June that district reads $+2.08$, the second highest of the 47 Junes on record, with 98% of its cells above $+1$ against a climatological 17%. Over that same 1 May to 31 August window Bojanala averages $+0.67$, sixth of the 47 years. The season score of $+0.20$ is what averaging June across five months and eight districts leaves.

9.8 Methods investigated and rejected

Two choices define the component, namely how the two channels are combined and which wind drives each. Both were selected against alternatives. Every construction in Table 5, and the wind alternative for the FWI channel, was rebuilt through the production pipeline, with the production inputs, the 1979 spin-up and its own rebuilt 1991–2020 climatology, then scored by the scale-matched validation of section 9.7. Three alternatives were not rebuilt, and each is identified as such below.

Combining the two channels. Table 5 lists the constructions tested. The two channels can be used alone, combined by taking the larger of the two normalised ratios (the production choice), combined as a disjunction by counting days on which either channel exceeds its own threshold, or combined multiplicatively.

The composite has the best score in the table, but only just. The standard-deviation product reaches the same nine hits and four further constructions reach eight, so thirteen episodes weighted towards the Western Cape cannot settle the ordering on counts alone. What the counts do show is the shape of the trade. The composite catches Betty's Bay, which the FWI channel alone reads at -0.24 and the ISI at $+0.28$, because the HDWI channel carries that episode. It misses Wupperthal at $+0.41$, which the ISI alone catches at $+1.13$, and the ISI discards the fuel-moisture memory entirely.

The two constructions that reach nine differ in which episodes they take. The composite catches Table Mountain at $+1.74$ and misses St Francis Bay at -0.39. The standard-deviation product does the reverse, reading Table Mountain at $+0.97$ and St Francis Bay at $+1.20$. Table Mountain is a documented wildland fire and St Francis Bay is a structural fire with no vegetation extent reported, so the composite takes the better nine.

Construction	Hits	Assessment
Larger of the two normalised ratios *(selected)*	**9/13**	The production component. It takes the most episodes of any construction here and spans both fire-weather regimes; see below and the channel evidence.
ISI alone	8/13	FFMC and wind only. One hit behind the composite, and it carries no Drought Code memory at all.
FWI channel alone	7/13	Structurally blind to wind-driven fire weather that is not preceded by drought. It reads Betty's Bay at -0.24.
HDWI channel alone	8/13	Carries no fuel-moisture memory. Misses Table Mountain and Overstrand, both carried by accumulated drying.
Either channel exceeding its own threshold (union)	8/13	Two thresholds instead of one, applied to each channel rather than to the combined value. A day qualifies on either channel alone, which is the looser test, so the count takes in days that clear one channel's own 95$^{\text{th}}$ percentile without standing out on the combined measure.
Product of the two normalised ratios	8/13	A conjunction, so it suppresses episodes carried by one channel while the other is near or below its normal level, and it reads Table Mountain at $+0.83$. Normalising each channel by its day-of-year standard deviation instead of its mean reaches nine hits, but still misses Table Mountain.

Table 5: Constructions evaluated, each rebuilt on the production configuration and scored by the scale-matched validation of section 9.7. "Hits" is out of the thirteen episodes.

The case for carrying both channels rests instead on what each contributes across the whole record. Over the South African cells there are 9.05 million cell-days on which the composite exceeds its threshold. Of those, 58.0% are driven by the HDWI channel, and 31.4% would not be flagged at all by the FWI channel alone. That share is strongly regional:

Province	Missed	Province	Missed
Northern Cape	38.7%	Mpumalanga	26.7%
Western Cape	33.7%	Free State	26.5%
Limpopo	31.7%	KwaZulu-Natal	25.0%
North West	30.0%	Gauteng	24.7%
		Eastern Cape	23.6%

"Missed" is the share of that province's composite exceedance cell-days, 1979–2026, that the FWI channel alone would not flag.

A single fuel-moisture channel would therefore under-read the arid west and the winter-rainfall fynbos belt, the fire-prone shrubland of the Cape coast, by about 1.6 times as much as the eastern interior, and the Western Cape carries the bulk of insured wildfire loss. For a product whose purpose is to compare regions and track them over time, a regional bias of that size is a more serious defect than a one-episode difference in a hit count.

The composite does pay a price for the maximum. Because the P95 threshold is computed on the composite, and a maximum of two fields is larger than either, the threshold sits above both channels' own. An episode can therefore clear one channel's threshold without clearing the composite's, which is what happens at Wupperthal.

Multiplying the HDWI channel by a standardised drought anomaly (AgERA5- or CHIRPS-based) was also tested, on the earlier prototype configuration, and collapsed the signal to essentially



one episode. Cell-level drought anomalies at fire locations are typically modest, and satellite-gauge precipitation products are known to perform poorly over the winter-rainfall Western Cape, where many episodes occur (22), so the drought factor suppresses almost every event day. That construction was not rebuilt, because the mechanism is understood and the outcome is not in doubt.

Wind for the FWI channel. The Canadian system is defined on readings taken at noon local standard time, so the FWI channel is driven by the ERA5 10 m wind at 10:00 UTC, which is noon South African Standard Time. The alternative tested was the 24-hour daily mean previously used for both channels.

The two winds agree on twelve of the thirteen validation episodes. The daily mean reaches 10/13 and the noon wind 9/13, and the single episode between them is St Francis Bay, read at $+1.33$ under the daily mean and -0.39 under the noon wind. That episode is the structural fire with no reported vegetation extent, and the loosest test in the set, its cell meeting the hit criterion in 44.7% of Novembers on record. Aside from this event the two winds are level, each taking 9 of the remaining twelve episodes. The event set therefore does not separate them, and the choice rests on agreement with the reference implementation, which improves sharply. Against GEFF the Initial Spread Index rises from a median per-cell correlation of 0.848 to 0.942, its bias falls from -1.93 to -0.27, and the share of cells above $r = 0.8$ rises from 78% to 99.7%. The FWI itself rises from 0.910 to 0.957 with its bias falling from -3.92 to -0.24.

Sampling at 12:00 UTC was rejected. That is 14:00 South African Standard Time, one to two hours after the humidity sample, and the standard is explicit that the observation hour is clock noon and is not corrected for local solar time (19).

Wind for the HDWI channel. The 24-hour daily mean is retained. Two alternatives were tested. The ERA5-Land daily mean at 0.1° is masked over the ocean and blanks coastal fire cells, forcing the nearest valid cell inland at several coastal episodes, which is disqualifying for a component whose largest losses are coastal.

The index's authors specify the daily maximum over daytime time steps rather than a fixed hour (26), so the daily mean is a deviation. The component already takes the vapour-pressure deficit at its daily peak, through maximum temperature and mid-afternoon humidity, and pairs it with a daily-mean wind, so two of the three inputs target the afternoon and one does not. The deviation is retained because the sub-daily alternatives above were tested and did not improve the result, and because the index measures exceedance of its own threshold, so a systematic difference in level cancels.

9.9 Long-term behaviour

Figure 4 shows the provincial and national history of the component. The thin line is the monthly anomaly $\delta^{F}_{m,k}$ averaged over the grid cells of each province. The bold line is its 60-month trailing mean, which is the moving average published alongside every component of the index, plotted at the last month of each window.

Against the fixed 1991–2020 baseline the component rises nationally. The national 60-month mean moves from -0.29 at its first defined value to $+0.14$ at its last. The mean number of exceedance days per cell-month doubles over the same span, from 0.95 per month in 1979–1990 to 1.90 from 2011 onward, passing through 1.21 in 1991–2000 and 1.42 in 2001–2010.

The rise is not uniform. It is largest in the northern interior and the arid west, at -0.47 to $+0.26$ in Gauteng, -0.41 to $+0.31$ in North West, -0.33 to $+0.32$ in Limpopo and -0.35 to $+0.24$ in the Northern Cape. KwaZulu-Natal does not rise at all, ending lower than it starts at $+0.03$ to -0.05, with a 2021–2026 mean of -0.06. The Western Cape rises least of the remaining provinces, at -0.32 to -0.06, and is the only one to end below its own baseline while the country ends above



it. That matters for interpretation, because the Western Cape carries the bulk of insured wildfire loss.

The vapour-pressure deficit drives the change, not the wind. Averaged over the South African grid cells, the deficit rises 12.6% between 1979–1990 and 2011–2025, from 20.9 to 23.5 hPa, as the daily maximum temperature gains $0.93\,^{\circ}$C and the mid-afternoon relative humidity loses 4.2 percentage points. Over the same interval both wind fields are flat, gaining 0.9% for the HDWI channel and 1.4% for the FWI channel. The Hot-Dry-Windy Index gains 13.7%, which is the deficit carried through almost unchanged.

The rise is not an artefact of the shared construction. Every component of the index is built from a day-of-year percentile threshold, a monthly frequency and a standardisation against the same fixed 1991–2020 baseline, so a trend common to all of them would point at the method rather than at the weather. They do not move together. The comparison below runs over Jan 1981–May 2026, with each aggregated over South Africa on its own grid. On that window the wildfire component rises $+0.42$. The extreme-maximum-temperature component, which counts hot spells, rises $+0.27$, and the extreme-minimum-temperature component, which counts cold spells, falls -0.48, both moving in the direction a warming climate implies. The moisture-driven components stay smaller. Drought moves $+0.01$ on AgERA5 and -0.23 on CHIRPS, precipitation -0.00 and $+0.20$ on the same two sources, and hail $+0.20$.

The wildfire component rises further than maximum temperature alone, $+0.42$ against $+0.27$, which is consistent with what drives it. It compounds warming with drying rather than restating temperature, as the vapour-pressure deficit figures above show.

Two properties of the output distribution follow. Because the recent climate sits above the fixed baseline, values above $+1$ are common in recent fire seasons. In 2024, the highest year on record, the national mean anomaly was $+0.85$ and 37% of cell-months exceeded $+1$.

The distribution is also asymmetric, with a heavier right tail than a Gaussian and a lighter left one. The anomaly standardises a count of exceedance days whose base-period mean is about 1.5 days per cell-month against a standard deviation of about 1.7 days, and 38% of base-period cell-months contain no exceedance day at all.

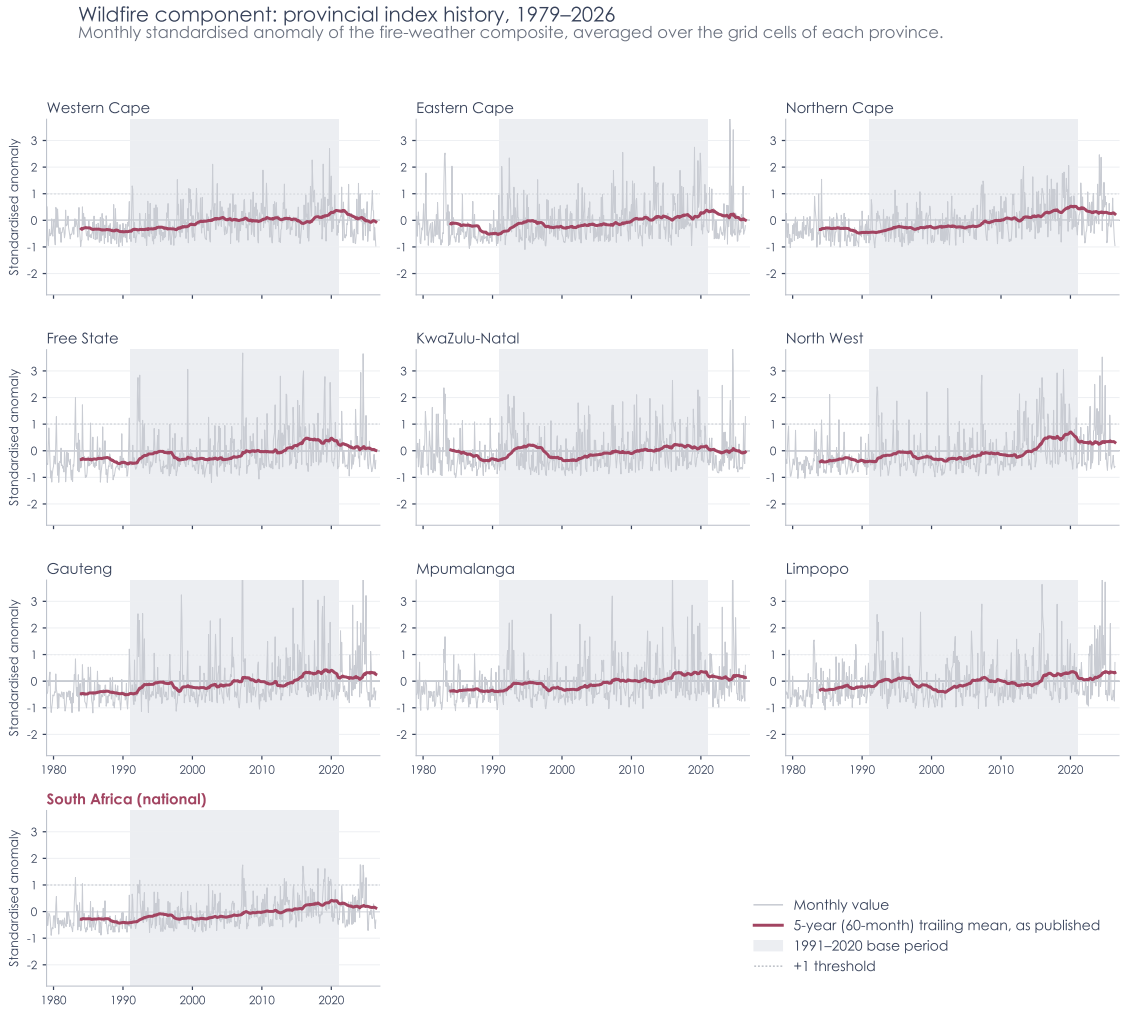


Figure 4: Provincial and national history of the wildfire component, 1979–2026, on the production output. Thin line: monthly standardised anomaly $\delta^{F}_{m,k}$ averaged over the grid cells of the province (arithmetic mean of all cells whose centroids fall within it, the aggregation used throughout the index). Bold line: its 60-month trailing mean, the moving average published with the index, plotted at the last month of each window and therefore starting in December 1983. Shading: the 1991–2020 base period. Dotted line: $+1$.

10 Composite Index

The composite combines four signals: the shift in the temperature distribution, extreme precipitation, drought and hail. The two temperature components enter as a single term,

$$\delta^{T}_{m,k} = \frac{\delta^{T_{Max}}_{m,k} - \delta^{T_{Min}}_{m,k}}{2},$$

and the composite before its own standardisation is the arithmetic mean of the four,

$$\tilde{\delta}^{C}_{m,k} = \frac{\delta^{T}_{m,k} + \delta^{P}_{m,k} + \delta^{D}_{m,k} + \delta^{H}_{m,k}}{4}.$$

Why the cold-spell term is subtracted. Hot-spell and cold-spell frequency describe the two ends of one temperature distribution. A shift towards warmer conditions raises the first and lowers the second, so entering both with the same sign measures how often temperature is extreme in either direction, which such a shift leaves close to unchanged. The two then offset each other and the composite records little. Subtracting the cold-spell term makes δ^T a measure of the displacement of the distribution, which is the quantity of interest. The Actuaries Climate Index and the Australian Actuaries Climate Index adopt the same convention for the same reason (5, 2).

Why the two temperature components form one term. Entered separately in a four-term mean they would carry half the composite between them, because both respond to the same shift. Combined into δ^T they carry a quarter, equal to each of the other three signals. The choice concerns weight rather than information, since the two components are close to uncorrelated from month to month and neither is redundant.

Which components enter. The hail component (section 8) enters. It is built from convective available potential energy and deep-layer shear, quantities no other component uses, and it is largely independent of the other three. The wildfire component (section 9) does not. It is constructed from the daily maximum temperature and the relative humidity, and its national monthly series correlates with the extreme-maximum-temperature component at $r = 0.64$, so including it would restate a signal the composite already carries through δ^T rather than add an independent one. It is published as a standalone index.

Which grid the components are combined on. Two composites are published, one for each precipitation dataset. CHIRPS carries precipitation alone, so the CHIRPS composite takes its extreme-precipitation and drought terms from CHIRPS and its two temperature terms from AgERA5. The AgERA5 composite takes all four of those terms from AgERA5. The hail component enters both from ERA5. The three sources are on different grids, CHIRPS at $0.05°$, AgERA5 at $0.1°$ and ERA5 at $0.25°$, so each component is bilinearly interpolated onto the grid of the precipitation dataset that names the composite, and the four signals are combined there cell by cell. Bilinear interpolation of a coarse field onto a finer one adds no information. It gives each fine cell a value consistent with its neighbours, and avoids the discontinuities that nearest-neighbour assignment would leave along the edges of the coarse cells. The effective resolution of the composite is set by its coarsest input, so structure finer than $0.25°$ comes from the precipitation, drought and temperature terms rather than from hail.

Standardisation of the composite. The mean of four weakly correlated standardised anomalies has a smaller standard deviation than any of them, so $\tilde{\delta}^C$ is not on the same scale as its components and a value of $+1$ would not carry the meaning it carries in a component. The composite is therefore standardised a second time, per calendar month and per cell, against the same 1991–2020 base period used throughout the index,

$$\delta^C_{m,k} = \frac{\tilde{\delta}^C_{m,k} - \mu^{\text{base}}_m}{\sigma^{\text{base}}_m},$$

where μ^{base}_m and σ^{base}_m are the mean and standard deviation of $\tilde{\delta}^C_{m,k}$ over the base-period years for calendar month m in that cell. Each component is already centred on that period, so the subtraction changes little and the division carries the effect.

Aggregation to the reporting regions. The standardised composite is then averaged over the cells falling within each district, each province and the country, which is the same step and the same arithmetic applied to every component. Standardising before this aggregation, rather than after it, holds the meaning of a given value fixed across the reporting levels. A month at

$+1$ is a month one base-period standard deviation above the norm whether it is read for a cell, a district or the country. Standardising after aggregation would instead force a spread of one onto whichever level the standardisation was applied to, and the composite would then read as more variable than any of the components it is built from, which are standardised at cell level. Composite values at cell level serve only as the input to this step. They are not published, and the composite is delivered for districts, provinces and the country.

References

(1) H. B. Abubakar. Evaluation of high-resolution precipitation datasets CHIRPS, TerraClimate and TAMSAT over the Enkangala Escarpment of South Africa. Preprint, Research Square. URL https://doi.org/10.21203/rs.3.rs-4365508/v1. Version 1.

(2) Actuaries Institute. Australian actuaries climate index: Design documentation. https://www.actuaries.asn.au/climate-index, 2018. Accessed 10 September 2026.

(3) A. A. Addisuu, G. M. Tsidu, and L. V. Basupi. Improving daily CMIP6 precipitation in southern africa through bias correction — part 1: Spatiotemporal characteristics. 13:95. doi: 10.3390/cli13050095.

(4) J. T. Allen, M. K. Tippett, and A. H. Sobel. An empirical model relating U.S. monthly hail occurrence to large-scale meteorological environment. *Journal of Advances in Modeling Earth Systems*, 7(1):226–243, 2015. doi: 10.1002/2014MS000397.

(5) American Academy of Actuaries, Canadian Institute of Actuaries, Casualty Actuarial Society, and Society of Actuaries. Actuaries climate index: Development and design. https://actuariesclimateindex.org/about/, 2019. Accessed 10 September 2026.

(6) J. Bagiliko, D. Stern, D. Ndanguza, and F. F. Torgbor. Validation of satellite and reanalysis rainfall products against rain gauge observations in Ghana and Zambia. 156:241. doi: 10.1007/s00704-025-05462-7.

(7) David Bolton. The computation of equivalent potential temperature. *Monthly Weather Review*, 108(7):1046–1053, 1980. doi: 10.1175/1520-0493(1980)108<1046:TCOEPT>2.0.CO;2.

(8) H. E. Brooks, J. W. Lee, and J. P. Craven. The spatial distribution of severe thunderstorm and tornado environments from global reanalysis data. *Atmospheric Research*, 67–68:73–94, 2003. doi: 10.1016/S0169-8095(03)00045-0.

(9) Climate Hazards Center, University of California, Santa Barbara. CHIRPS: Climate Hazards Group Infrared Precipitation with Station data. Data portal, 2025. URL https://www.chc.ucsb.edu/data/chirps.

(10) Copernicus Climate Change Service (C3S). Agrometeorological indicators from 1979 to present derived from reanalysis (AgERA5). Copernicus Climate Data Store, 2025. URL https://cds.climate.copernicus.eu/datasets/sis-agrometeorological-indicators.

(11) J. de Waal, M. Maranan, A. H. Fink, A. Watson, A. Bennett, and J. Helmschrot. Validation of satellite precipitation estimates over the Western Cape region, South Africa. 26:759–779. doi: 10.1175/JHM-D-24-0123.1.

(12) W. M. D. Dlamini and S. S. Tfwala. The performance of a high-resolution satellite-derived precipitation product over the topographically complex landscape of Eswatini. 12:e278. doi: 10.1002/gdj3.278.

(13) J. A. du Plessis and J. K. Kibii. Applicability of CHIRPS-based satellite rainfall estimates for South Africa. 63(3):43–54. ISSN 1021-2019. doi: 10.17159/2309-8775/2021/v63n3a4.

(14) Imke Durre, Russell S. Vose, and David B. Wuertz. Overview of the Integrated Global Radiosonde Archive. *Journal of Climate*, 19(1):53–68, 2006. doi: 10.1175/JCLI3594.1.

(15) C. Funk, P. Peterson, M. Landsfeld, D. Pedreros, J. Verdin, J. Rowland, L. Bo, G. Husak, J. Michaelsen, and J. Verdin. The climate hazards infrared precipitation with stations – a new environmental record for monitoring extremes. *Scientific Data*, 2:150066, 2015. doi: 10.1038/sdata.2015.66.

(16) H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: 10.1002/qj.3803.

(17) R. J. Hyndman and Y. Fan. Sample quantiles in statistical packages. *The American Statistician*, 50(4):361–365, 1996.

(18) J. A. Kruger, S. J. Roffe, and A. J. van der Walt. AgERA5 representation of seasonal mean and extreme temperatures in the Northern Cape, South Africa. 120(3/4):16043. doi: 10.17159/sajs.2024/16043.

(19) B. D. Lawson and O. B. Armitage. Weather guide for the Canadian Forest Fire Danger Rating System. Technical report, Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Edmonton, Alberta, 2008. Appendix 3, latitude considerations in adapting the FWI System for use in other countries, by M. E. Alexander.

(20) P. T. Mahlalela, R. C. Blamey, N. C. G. Hart, C. J. C. Reason, et al. Drought in the Eastern Cape region of South Africa and trends in rainfall characteristics. 55:2743–2759. doi: 10.1007/s00382-020-05413-0.

(21) S. Mohr and M. Kunz. Recent trends and variabilities of convective parameters relevant for hail events in Germany and Europe. *Atmospheric Research*, 123:211–228, 2013. doi: 10.1016/j.atmosres.2012.05.016.

(22) W. V. Pitman. Can CHIRPS fill the gap left by the decline in the availability of rainfall stations in southern Africa? *Water SA*, 47(2):162–171, 2021. doi: 10.17159/wsa/2021.v47.i2.10912.

(23) A. F. Prein and G. J. Holland. Global estimates of damaging hail hazard. *Weather and Climate Extremes*, 22:10–23, 2018. doi: 10.1016/j.wace.2018.10.004.

(24) Heinz Jürgen Punge, Kristopher M. Bedka, Michael Kunz, Sarah D. Bang, and Kyle F. Itterly. Characteristics of hail hazard in South Africa based on satellite detection of convective storms. *Natural Hazards and Earth System Sciences*, 23:1549–1576, 2023. doi: 10.5194/nhess-23-1549-2023.

(25) S. J. Roffe and A. J. van der Walt. Representation and evaluation of southern Africa's seasonal mean and extreme temperatures in ERA5-based reanalysis products. 284:106591. doi: 10.1016/j.atmosres.2022.106591.

(26) A. F. Srock, J. J. Charney, B. E. Potter, and S. L. Goodrick. The hot-dry-windy index: A new fire weather index. *Atmosphere*, 9(7):279, 2018. doi: 10.3390/atmos9070279.

(27) Mateusz Taszarek, John T. Allen, Mattia Marchio, and Harold E. Brooks. Global climatology and trends in convective environments from ERA5 and rawinsonde data. *npj Climate and Atmospheric Science*, 4:35, 2021. doi: 10.1038/s41612-021-00190-x.

(28) D. Terblanche, A. Lynch, Z. Chen, and S. Sinclair. ERA5-derived precipitation: Insights from historical rainfall networks in southern Africa. 61:1473–1484. doi: 10.1175/JAMC-D-21-0096.1.

(29) C. E. Van Wagner. Development and structure of the Canadian Forest Fire Weather Index System. Forestry Technical Report 35, Canadian Forestry Service, Petawawa National Forestry Institute, Ottawa, 1987.

(30) C. Vitolo, F. Di Giuseppe, B. Krzeminski, and J. San-Miguel-Ayanz. A 1980–2018 global fire danger re-analysis dataset for the Canadian Fire Weather Indices. *Scientific Data*, 6:190032, 2019. doi: 10.1038/sdata.2019.32.