



Casa Blanca Weather

Methods and verification

For hydrologists, meteorologists and anyone who wants the details. Station ISTANF10, Stanford, Western Cape.

Casa Blanca Weather combines four global weather models, corrects each against three years of readings from one farm weather station in Stanford, and blends them. This page describes exactly what is done, how well it has worked on a year of data the corrections never saw, and where it falls short.

1. The station

Instrument. Ambient Weather WS-2902 (station ISTANF10), Casa Blanca Farmstead, Stanford, Western Cape; lat -34.4417 , lon 19.4583 . The integrated sensor array (thermo-hygrometer in a small passive radiation shield, cup anemometer, wind vane, tipping-bucket rain gauge, solar/UV sensor) is mounted at about **2.5 m**, north-oriented, in the most open spot of a $\frac{1}{2}$ -acre farmstead. Windbreak trees line the perimeter, especially to the south. Record from **27 June 2023** (first full day 28 June), 5-minute data from the Ambient Weather Network API.

Comparison with WMO siting. WMO guidance (CIMO Guide) puts air temperature at 1.25–2 m in a ventilated screen over short grass, and wind at 10 m over open, level terrain with obstacles at a distance of at least ten times their height. This station meets neither: temperature is measured at 2.5 m (above the standard height) in a compact passive shield, and the anemometer is at 2.5 m and sheltered by trees and farm buildings. The readings are therefore **site-representative** (what this farmstead experiences) but **not WMO-comparable**. The wind sensor records about $0.4 \times$ the ECMWF 10 m gust and rarely sees south to west afternoon winds, which the trees block or deflect.

Rain gauge. The WS-2902 gauge is a small, uncalibrated tipping bucket on the same mast (rim ≈ 2.5 m rather than the usual 0.3–1 m). Expect **under-catch in wind** (the rim is exposed) and in **heavy rain** (tips are missed while the bucket moves), possible loss of very light drizzle, and no snow/hail capability. We have not compared it with a standard gauge. Daily totals are built from the device's daily counter, corrected for the midnight carry-over and counter resets; days with fewer than 260 of 288 records or where the daily and cumulative counters disagree are excluded, never filled. These verification days are calendar days (midnight to midnight SAST). The live “Now” card instead shows rain on the SAWS rain day, **08:00 to 08:00 SAST**: “Rain since 08:00” and “Last 24 h to 08:00”, from differences of the gauge's cumulative counter, so the latter is comparable with manual gauges read at 08:00.

Pressure. On 28 September 2026 at about 07:25 SAST the console was recalibrated so relative pressure equals absolute pressure + 2 hPa (about 17 m elevation, matching the barometer at about 17–19 m; roughly 0.12 hPa per m). Relative pressure recorded before 28 Sep reads about 8.5 hPa low and should not be compared directly with later values. The live 3-hour pressure-trend arrow is computed from absolute pressure, so the recalibration does not affect it. A cross-check against SAWS Hermanus over calm days is ongoing.

Live wind rose. The live page shows a rolling 24-hour wind rose built from the station's 5-minute readings, using 16 sectors and a calm percentage. It shows the sheltered-site wind as measured; south-west winds are under-represented because of the windbreak. The live page is a gauge dashboard refreshed every minute from the station.

2. Models

- **ECMWF IFS** via Open-Meteo `best_match`. For this location `best_match` resolved to ECMWF IFS HRES for the whole test period (checked by exact-value matches against `ecmwf_ifs`); before 1 Oct 2025 the previous-runs archive of `best_match` is stitched from GFS/ICON, so ECMWF training uses the same-day (historical-forecast) archive. Shown on the site, uncorrected, as “ECMWF forecast”.
- **GFS** (NCEP, `gfs_seamless`), **ICON** (DWD, `icon_seamless`) and **UKMO** (Met Office, `ukmo_seamless`).
- Daily values are aggregated in SAST from hourly output at the grid point nearest the station: high/low = max/min of hourly 2 m temperature, rain = 24-hour sum, gust = daily max of hourly gusts. Archives: Open-Meteo Historical Forecast (same-day runs) and Previous Runs (forecasts issued 1–7 days earlier).

3. Method

3.1 Per-model temperature correction (MOS-style linear regression)

For each model, the error (station – forecast) of the daily high and low is regressed on: a separate intercept per season (DJF, MAM, JJA, SON); forecast cloud fraction (daytime for the high, 00–08 h for the low); for the low, forecast night-time 10 m wind; for the high, (forecast high – 20 °C). The fit is robust least squares (three passes, trimming residuals beyond $3 \times \text{MAD}$). Training: complete station days up to **30 Sep 2025** (high-temperature training days: ECMWF 809, GFS 607, ICON 607, UKMO 160), using day-before forecasts where genuine and same-day forecasts otherwise.

Extra terms were tested out of sample and **not kept** because they did not improve the test year: a calm-clear night term in the low (forecast night wind < 10 km/h and cloud < 30 %); a hot-day term in

the high (forecast high ≥ 30 °C); an offshore (NW–N–NE) warm-day term (forecast high ≥ 25 °C). Their coefficients are fixed at zero.

3.2 Adjustments applied to the blend

- **Rolling 30-day high bias:** the mean error of the blended high over the last 30 complete station days (at least 10) is subtracted, at every lead.
- **Error carry-over (same-day issue only):** $0.115 \times$ yesterday's blend error is subtracted from today's high and $0.242 \times$ from the low (lag-1 persistence of error; the lag-2 version for day 1 was tested and dropped).
- **Morning nowcast:** for today's forecast (issued 07:37 SAST), low = min(forecast low, station minimum since midnight).

3.3 Blend

“Our forecast” is the **equal-weight mean of the corrected models** available at each lead (all four to day 5, then ECMWF, GFS, ICON on day 6 and ECMWF, GFS from day 7). No weights are fitted. The likely range drawn in the charts is the spread of the corrected models.

3.4 Rain

Rain is locally corrected, per model: a day is wet if the model's rain reaches a seasonal threshold chosen on training data to maximise the Heidke skill score for station ≥ 0.5 mm; wet-day amounts are then mapped by seasonal quantile mapping (1 % quantiles, station vs model, training years). The blend amount is the mean of the corrected models. The **chance of rain** shown is the mean of (a) the share of models above their wet threshold and (b) ECMWF's precipitation probability, recalibrated with a reliability table from the same-day archive.

3.5 Wind gusts

Station daily max gust is regressed linearly on each model's daily max 10 m gust (ECMWF: **$0.407 \times \text{gust} - 0.92 \text{ km/h}$** ; GFS $0.371 \times \text{gust} + 5.65$; ICON $0.394 \times \text{gust} + 1.17$; UKMO $0.332 \times \text{gust} + 3.62$), then averaged. This is the “Sheltered” gust. Because the regression is fitted on mostly ordinary days it shrinks towards typical values: on the 5 test-year days with a station gust ≥ 40 km/h the blend was **13.8 km/h too low** on average. The strong-wind flag therefore fires at a tuned 35 km/h (test year: 2 of 5 events caught, 0 false alarms). “Open ground” is the raw ECMWF gust.

3.6 Confidence percentages

A day counts as a **hit** when the corrected high **and** low are both within ± 2 °C of the station. Hit rates were measured by lead on 1 Oct 2025 – 26 Sep 2026 (out of sample) by re-issuing the forecast from archived model runs (day 1: 71%, day 7: 34%). Each day's value is adjusted by model spread: the days at each lead are split into terciles of corrected-model spread and the tercile's hit rate is used (bins of ≥ 30 days). Days 8–10 have no archive and are **extrapolated** (weighted least squares, $\text{logit}(\text{hit rate}) =$

a + b*lead over measured leads 1-7); they are marked with an asterisk. Values are rounded to 5 % and clamped to 10–95 %.

4. Verification

Train/test split. All corrections were fitted on data up to 30 Sep 2025 and are verified on the following year, 1 Oct 2025 – 25 Sep 2026, which they never saw. Scores are bias / MAE / share within $\pm 2^\circ\text{C}$ of the station on complete days, day-ahead forecasts (issued the day before).

4.1 Temperature, day-ahead, test year

Model		Raw			Corrected			n
		bias	MAE	$\pm 2^\circ$	bias	MAE	$\pm 2^\circ$	
ECMWF	hi	-1.26	1.77	61%	+0.26	1.27	82%	359
ECMWF	lo	+1.40	1.75	62%	+0.17	1.29	78%	359
GFS	hi	-1.30	1.71	64%	+0.42	1.23	80%	359
GFS	lo	+1.67	2.27	50%	-0.19	1.46	70%	359
ICON	hi	-3.00	3.07	27%	+0.34	1.36	77%	355
ICON	lo	+1.68	1.86	59%	-0.18	1.27	81%	355
UKMO	hi	-1.37	1.84	59%	+0.18	1.21	81%	314
UKMO	lo	-0.40	1.59	70%	-0.33	1.55	71%	315
Mean of 4	hi	-1.76	1.96	54%	+0.25	1.08	86%	312
Mean of 4	lo	+1.05	1.52	70%	-0.13	1.22	80%	313

$^\circ\text{C}$; hi/lo = daily high/low; $\pm 2^\circ$ = share of days within $\pm 2^\circ\text{C}$. “Corrected” = per-model correction (3.1) only. Mean of 4 = equal-weight mean on days with all four models. Source: summary.md.

As published (blend + adjustments of 3.2, test year): day-ahead high MAE **1.01** $^\circ\text{C}$ (bias +0.04), low **1.21** (bias -0.13), n = 359; same-day issue high 0.98, low 0.86 (the morning nowcast gives the low a small cold bias, -0.41). Raw ECMWF day-ahead: high 1.77, low 1.75.

4.2 By lead time (published blend)

Lead (days)	0	1	2	3	4	5	6	7
High MAE $^\circ\text{C}$	0.91	1.01	1.12	1.21	1.34	1.51	1.94	2.19
Low MAE $^\circ\text{C}$	0.89	1.21	1.27	1.34	1.40	1.54	1.72	1.81
Both within $\pm 2^\circ\text{C}$	77%	71%	68%	64%	59%	51%	43%	34%
n days	360	360	360	360	360	360	360	358

Wind: station daily max gust, day-ahead, test year: blend MAE **3.7** km/h (bias -2.4) vs raw ECMWF 22.8 km/h (bias +22.8, mostly the site effect), n = 359.

4.3 Rain

Test year 1 Oct 2025 – 25 Sep 2026: **359** valid station days (1 excluded, no gap-filling), 731 mm observed, 81 days ≥ 1 mm, 42 days ≥ 5 mm, 19 days ≥ 10 mm. Lead 0 is the archived same-day run. Leads with an archive so far: day 0, 1, 2, 3, 5, 7.

≥ 1 mm	Lead	POD	FAR	Freq. bias	ETS	HSS	events / n
Our forecast	0	0.86	0.27	1.19	0.57	0.72	81 / 359
Our forecast	1	0.84	0.28	1.16	0.55	0.70	74 / 337
Our forecast	2	0.83	0.36	1.29	0.46	0.63	69 / 323
Our forecast	3	0.77	0.37	1.23	0.42	0.60	75 / 338
Our forecast	5	0.66	0.51	1.35	0.27	0.42	68 / 322
Our forecast	7	0.57	0.55	1.27	0.20	0.33	81 / 359
ECMWF raw	0	0.79	0.27	1.09	0.52	0.68	81 / 359
ECMWF raw	1	0.82	0.26	1.11	0.56	0.72	74 / 337
ECMWF raw	2	0.74	0.31	1.07	0.46	0.63	69 / 323
ECMWF raw	3	0.68	0.35	1.05	0.39	0.56	75 / 338
ECMWF raw	5	0.65	0.46	1.19	0.31	0.47	68 / 322
ECMWF raw	7	0.63	0.55	1.38	0.22	0.36	81 / 359
Persistence	0	0.46	0.54	1.00	0.17	0.30	81 / 358
Persistence	1	0.29	0.72	1.03	0.04	0.08	73 / 336
Persistence	2	0.23	0.78	1.03	0.01	0.01	69 / 322
Persistence	3	0.20	0.79	0.96	-0.01	-0.02	75 / 337
Persistence	5	0.19	0.83	1.12	-0.03	-0.06	68 / 321
Persistence	7	0.23	0.77	1.02	0.00	0.00	81 / 358

Wet-day contingency at ≥ 1 mm. Persistence = last station day known at issue time.

≥ 5 mm	Lead	POD	FAR	Freq. bias	ETS	HSS	events / n
Our forecast	0	0.64	0.21	0.81	0.51	0.68	42 / 359
Our forecast	1	0.58	0.25	0.78	0.45	0.62	36 / 337
Our forecast	2	0.54	0.32	0.80	0.39	0.56	35 / 323
Our forecast	3	0.57	0.38	0.92	0.37	0.54	37 / 338
Our forecast	5	0.36	0.43	0.64	0.25	0.40	33 / 322
Our forecast	7	0.36	0.50	0.71	0.21	0.35	42 / 359
ECMWF raw	0	0.48	0.29	0.67	0.36	0.53	42 / 359
ECMWF raw	1	0.50	0.10	0.56	0.44	0.61	36 / 337
ECMWF raw	2	0.46	0.24	0.60	0.36	0.53	35 / 323
ECMWF raw	3	0.41	0.29	0.57	0.31	0.48	37 / 338
ECMWF raw	5	0.33	0.31	0.48	0.26	0.41	33 / 322
ECMWF raw	7	0.26	0.62	0.69	0.13	0.24	42 / 359
Persistence	0	0.31	0.69	1.00	0.12	0.22	42 / 358
Persistence	1	0.22	0.79	1.06	0.06	0.12	36 / 336
Persistence	2	0.09	0.91	0.97	-0.01	-0.02	35 / 322
Persistence	3	0.14	0.86	0.95	0.02	0.04	37 / 337
Persistence	5	0.18	0.84	1.12	0.04	0.07	33 / 321
Persistence	7	0.14	0.86	1.00	0.01	0.03	42 / 358

Wet-day contingency at ≥ 5 mm (small sample: about 35–42 events).

Amounts	Lead	MAE wet	bias wet	total ratio	n wet
Our forecast	0	4.40	-2.48	0.81	81
Our forecast	1	4.07	-1.94	0.88	74
Our forecast	2	4.23	-1.42	0.98	69
Our forecast	3	5.09	-2.91	0.84	75
Our forecast	5	4.98	-3.33	0.83	68
Our forecast	7	7.06	-5.32	0.72	81
ECMWF raw	0	4.89	-3.64	0.69	81
ECMWF raw	1	4.52	-3.13	0.72	74
ECMWF raw	2	5.11	-3.22	0.75	69
ECMWF raw	3	5.29	-3.55	0.71	75
ECMWF raw	5	5.60	-4.04	0.69	68
ECMWF raw	7	6.98	-5.16	0.75	81
Climatology	0	7.00	-6.64	0.97	81
Climatology	1	6.78	-6.40	1.00	74
Climatology	2	7.00	-6.62	1.01	69
Climatology	3	6.78	-6.39	1.02	75
Climatology	5	5.62	-5.24	1.19	68
Climatology	7	7.00	-6.64	0.97	81
Persistence	0	8.60	-4.09	1.00	81
Persistence	1	9.23	-5.01	1.03	73
Persistence	2	8.30	-7.31	0.85	69
Persistence	3	9.47	-6.16	0.83	75
Persistence	5	6.58	-6.17	1.32	68
Persistence	7	8.63	-6.83	1.01	81

mm. Wet day = station ≥ 1 mm; total ratio = forecast total / observed total over all days. Climatology = station monthly mean daily rain (training years). Individual GFS, ICON and UKMO scores, raw and corrected, are in rain_verification.md.

Chance of rain (event: station ≥ 0.5 mm, same-day, n = 359): Brier score **0.066** for the site's chance, 0.088 for raw ECMWF probability and 0.193 for station monthly climatology (Brier skill +0.66 and +0.54). The Open-Meteo previous-runs archive carries no probabilities, so the chance cannot yet be verified beyond day 0; the recalibration table overlaps the test year, so this score is partly in-sample.

Forecast chance	n days	mean forecast	observed frequency
0-10%	217	0.01	0.03
10-20%	19	0.17	0.16
20-30%	4	0.25	0.00
30-40%	13	0.33	0.00
40-50%	8	0.44	0.25
50-60%	6	0.56	0.50
60-70%	12	0.65	0.50
70-80%	18	0.76	0.61
80-90%	9	0.83	0.89
90-100%	53	0.96	0.98

Reliability, day 0. Mid-range bins hold only 4–19 days each.

Reading the rain results. The corrected blend detects more wet days than raw ECMWF but also over-forecasts them at ≥ 1 mm (frequency bias above 1); skill falls with lead time, and persistence has little skill beyond day 0. Totals are under-forecast overall (total ratio 0.72–0.98 across days 0–7), but not uniformly: by season the day-1 ratio runs from 0.73 to 1.19; and heavy days are clearly too low: on the 15 day-1 days with ≥ 10 mm the gauge averaged 25.5 mm against 18.3 mm forecast. With a single uncalibrated gauge that probably under-catches, the true under-forecast of heavy rain may be larger than shown.

5. Limitations

- **One station.** Every correction and score comes from a single point. Errors in the sensors (shield, siting, gauge catch) are learned as “truth”.
- **Short record.** About three years of data, one independent test year. Extremes (heat ≥ 33 °C, gusts ≥ 40 km/h, rain ≥ 10 mm) number in the tens or fewer; their scores are indicative only.
- **Rain.** Corrected with a simple threshold and quantile mapping from a tipping bucket that has never been checked against a standard gauge. Heavy-rain amounts remain under-forecast.
- **Representativeness.** The corrections describe this farmstead: its shelter, exposure and gauge. They do not transfer to the wider village, the Klein River valley or catchments further inland, where rain and wind can differ markedly.
- **Operational differences.** Verification re-issues forecasts from archived runs at the grid point; the live forecast uses the current Open-Meteo run at 07:37 SAST, so live scores can differ slightly.

6. Data and contact

Help us check our rain gauge. If you run a standard or calibrated rain gauge in or near Stanford, we would love to compare daily totals. Co-located comparisons (a standard gauge a few metres from ours) are especially welcome, as are any checks of our temperature and wind against reference instruments.

Data: 5-minute station records from 27 June 2023 and all archived forecasts and verification files are can be shared as CSV/JSON on request. Please contact Casa Blanca Farmstead, Stanford.

Page last updated Fri 9 Oct 2026, 20:17 SAST. Numbers are regenerated from the verification files at each daily rebuild.