Delhi Air: Weekly Deep Analysis

Pollution Roses, Wind-Direction Testing & Weather-vs-Pollution Correlations

Prepared by Ajay Maken MP and Team

Week ending 2026-07-13

Contents

This report at a glance

Over the last 90 days, Delhi's PM2.5 averaged about 3.5 times the WHO 24-hour safe limit.

What this is. This is the weekly deep-analysis report on Delhi's air. It carries the heavier scientific material that cannot be settled in a single day, and presents it once a week, when there is enough fresh data to say something reliable. It answers three practical questions in plain terms: (1) *which wind directions bring Delhi its dirtiest air*, (2) *whether wind direction genuinely matters once we account for how hard the wind is blowing*, and (3) *how closely day-to-day swings in the weather track day-to-day swings in pollution, month by month.* Every number in this report is computed directly from the monitoring data and the weather record at the moment the report is built, and none of it is typed in by hand.

One important limit this week. Delhi's government monitors also report gases (nitrogen dioxide, sulphur dioxide, carbon monoxide, ozone). Since **23 June 2026** the gas feed has been frozen and its numbers are not trustworthy, so this edition deliberately covers only the two particle pollutants, **PM2.5** (fine particles) and **PM10** (coarse dust), whose measurements are unaffected. The gas sections will return in a future edition once the feed is repaired.

Window covered. Analysis ends **13 July 2026**. Directional plots use the last **90 days**; the month-by-month tests cover **January to July 2026**.

This week's headline findings

- **The air is far above the safe limit.** Over the last 90 days Delhi's city-average fine particles (PM2.5) ran about **52 µg/m³** and coarse dust (PM10) about **173 µg/m³**. The World Health Organisation's 24-hour safe levels are just **15** and **45 µg/m³**, so PM2.5 has been roughly **3.5 times** the WHO limit and PM10 about **3.8 times** it.
- **Dirtiest wind, cleanest wind.** Over the same period, days with a dominant **NW** wind (27 such days) carried the highest typical PM2.5, while **SE** winds (16 days) brought the cleanest air. Only sectors with at least 5 days are ranked, so a one- or two-day wind cannot be crowned. These two are the highest and the lowest of the 4 sectors ranked, and no significance test stands behind that ordering, so read them as the ends of this period's range rather than as a measured gap. The pollution roses on the following pages show the full picture by direction.
- **Does direction really matter?** After allowing for wind speed and the depth of the atmosphere (which do most of the mechanical clearing), wind *direction* adds statistically detectable information in the months tested. See the wind-direction test table. Where it is significant, direction is pointing at a real source, not just calmer or windier days.
- **Weather still drives the day-to-day swings.** In June, among the year-on-year comparisons, the strongest link that survived the report's multiple-testing screen was between **Temperature 2 m (°C)** and **PM10**, at **0.76**. This is the strongest of the 12 links tested on that grid, and picking the largest of 12 will always flatter it a little, so read it as the best of the month rather than as a precise measurement. The other grid for the same month is shown beside it and is read separately, because the two grids answer different questions.

Pollution roses (last 90 days)

What this shows and why it matters. A *pollution rose* is a compass built from the last 90 days. Each spoke points to a direction the wind blew *from*. **The length of a spoke shows how *often* the wind came from that direction: a long spoke is a *frequent* wind, not a dirty one.** The **colour bands** stacked along each spoke show how dirty the air was on those days (darker red = more pollution). So read length and colour together: a long, pale spoke is a common but clean wind; a shorter, deep-red spoke is a less frequent but dirtier one. The “dirtiest wind” named on the previous page is the sector with the highest *typical* (median) pollution, which need not be the longest spoke on the rose. Together these point a campaigner to where Delhi’s dirtiest air tends to blow in from, and that is the first step to naming the sources.

PM2.5 pollution rose

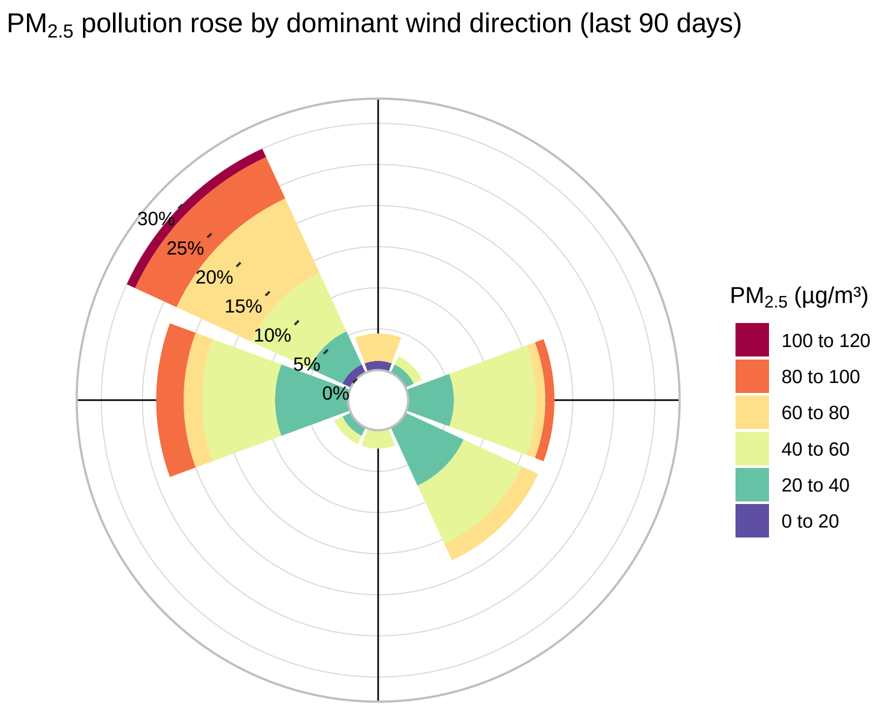

PM10 pollution rose

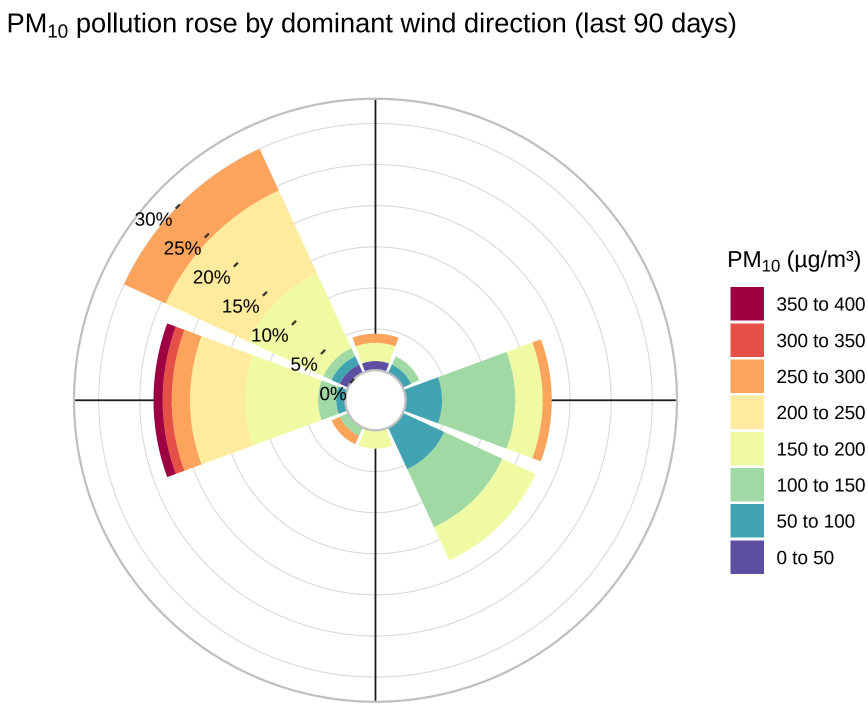

Polar plots (pollution spread across wind speed and direction) need hour-by-hour station data to draw a reliable surface; on daily averages they are too sparse to read, so they are held back until the hourly feed is restored.

Wind-direction impact and significance testing

What this shows and why it matters. The boxplot below groups the daily city PM readings by the day's dominant wind direction, so you can see directly which sectors tend to sit high and which sit low. But a raw difference between sectors can fool us: in Delhi the cleaner-looking winds are often simply the *stronger* winds, and strong winds clear pollution no matter where they come from. So the table that follows runs the honest test. It asks whether wind *direction* still explains pollution **after** we have already accounted for wind *speed* and the depth of the atmosphere (the boundary-layer height). Only if direction survives that adjustment can we say a particular sector is genuinely bringing pollution in, rather than just being a calm or a windy day with a directional label.

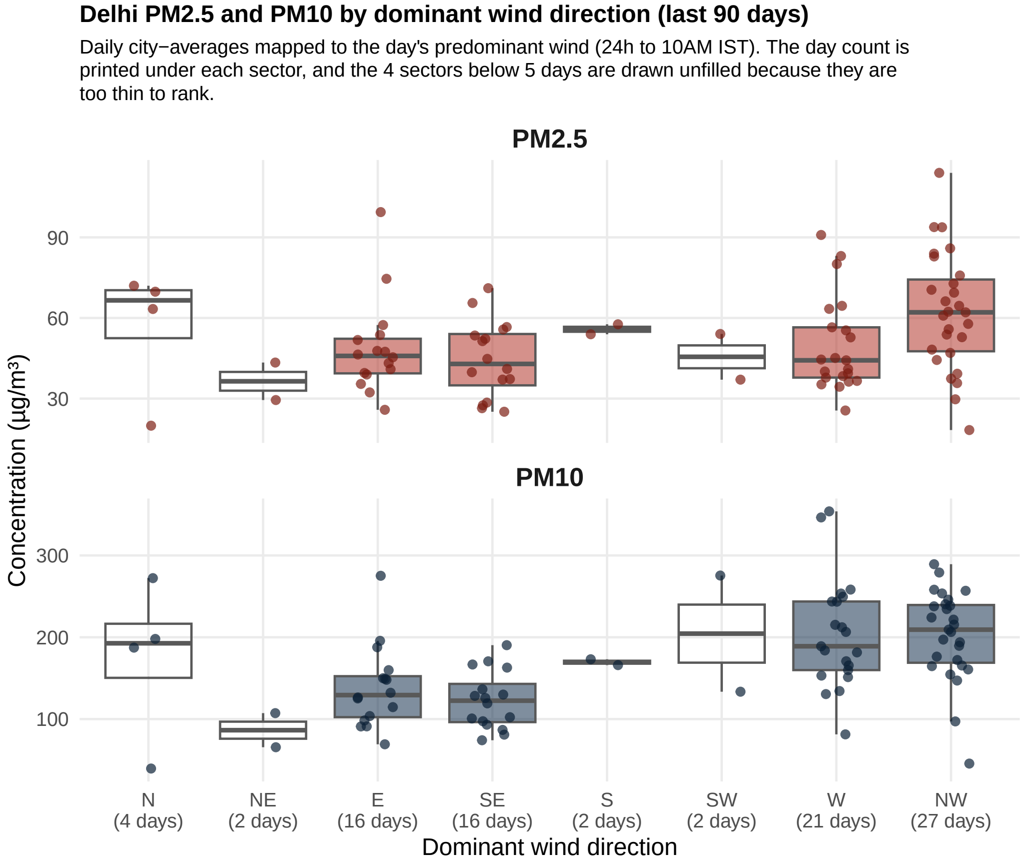

Table 1: Does wind direction matter once wind speed and boundary-layer height are accounted for? Partial-F p = partial F-test on the wind-direction factor in a model of PM on wind direction plus wind speed plus BLH. Partial-F p (BH-adj) applies a Benjamini-Hochberg multiple-testing screen across all the direction tests in this table, the same screen the correlation heatmaps use. Signif? = Yes when a month's BH-adjusted p is at or below 0.10, meaning the report expects no more than about 1 in 10 flagged results to be a false alarm (a 10 percent false-discovery rate); months with too few days to test reliably are marked 'too few days'. KW p = Kruskal-Wallis and ANOVA p = one-way ANOVA are shown as raw, unadjusted checks. A version of this test corrected for day-to-day persistence in the weather was built and tested. At the number of days available in a single month it proved unreliable in both directions, so this table reports the ordinary test with the sample size shown, and the corrected version is held back until it can be made dependable.

Month	Pollutant	Days	Sectors	Partial-F p	Partial-F p (BH-adj)	KW p	ANOVA p	Signif?
April	PM2.5	30	7	0.467	0.533	0.473	0.524	No
April	PM10	30	7	0.575	0.575	0.475	0.500	No
May	PM2.5	31	5	0.081	0.162	0.008	0.015	No
May	PM10	31	5	0.001	0.005	0.000	0.000	Yes
June	PM2.5	30	7	0.001	0.005	0.050	0.000	Yes
June	PM10	30	7	0.193	0.309	0.069	0.050	No
July	PM2.5	13	7	0.428	0.533	0.169	0.132	too few days
July	PM10	13	7	0.057	0.152	0.163	0.004	too few days

Monthly correlation: weather vs pollution (year-on-year and day-on-day)

What this shows and why it matters. Each month gets two heatmaps that ask a single question: *when the weather was unusual, was the pollution unusual in the same way?* The top grid (**year-on-year**) compares each 2026 day with the same calendar window in 2025; the bottom grid (**day-on-day**) compares each day with the one before it, which is the cleaner test because it removes anything that drifts slowly across the month. **Blue** means the two move in opposite directions (for example, stronger wind, lower pollution, the atmosphere doing its job); **red** means they rise and fall together. The number in each cell is the correlation, with a robustness check in brackets.

Reading the weather rows. Each weather row carries a small tag for how it usually helps or hurts: **(hi)** = Delhi's air is usually cleaner when this is *higher* (stronger wind, a deeper mixing layer, more rain); **(lo)** = cleaner when this is *lower* (a strong morning inversion traps pollution near the ground, so lower is better); **(n)** = no fixed direction (temperature).

How we keep it honest. Delhi's pollution *level* carries over from one day to the next, because weather systems last several days, so the daily readings within a month are not fully independent pieces of evidence. That is exactly why this report does not rely on the levels alone. The lower **day-on-day** grid compares each day only with the day before, which subtracts out that slow carry-over and tests the sharper question: when the weather *changed*, did the pollution change with it.

Why one significance test is not enough. Because the days carry over, the ordinary test can credit a month with more confidence than it has earned. So this report also runs a correction published by Whitney Newey and Kenneth West, which measures how much each day still resembles the days before it and rebuilds the margin to match. **That correction does not always make this report stricter, and we will not pretend that it does.** On the day-on-day grid each reading is already a change from the day before, and that subtraction has usually over-corrected for the carry-over, so the rebuilt margin can come out *narrower* than the ordinary one and make a finding easier to reach rather than harder. Across all 160 tests in this edition, the correction moved the answer towards caution in only 50 of them. A correction that can loosen the standard cannot be the only guard, so **this report computes both answers and keeps whichever of the two is the more cautious.** A cell must satisfy both tests before it is allowed to carry colour. That is a stricter standard than either test on its own.

Why a grid needs a second screen. One grid is not one test: for most months it is twelve, being six weather measurements set against two pollutants. Testing twelve things at once and then pointing at the most striking of them is how a report manufactures a finding that was never there. A screen published by Yoav Benjamini and Yosef Hochberg therefore ranks each grid's results and raises the bar for each in turn, set so that among all the cells allowed to carry colour we expect about one in ten to be a coincidence. Each month and each grid is screened on its own, so a weak month cannot borrow credibility from a strong one. A cell computed on fewer than 20 usable pairs is not tested at all, because an autocorrelation-robust margin cannot be trusted on a handful of days. Of the 160 cells in this edition, 28 carry colour.

Cells that fail the screen are printed in grey. Their numbers are still shown, because hiding a figure the reader might want is its own kind of dishonesty, but a grey number describes what the month happened to look like and is not evidence of anything. One honest limit: the screen works cleanly when the tests on a grid do not conspire with one another, and Delhi's weather does not fully oblige, because wind, rain, temperature and mixing depth rise and fall together. Read the ten per cent as a sound working guide rather than as a guarantee.

January 2026

Year−on−year (n = 30 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.08 (−0.06)	−0.08 (−0.04)
Temperature 2 m (°C) (n)	0.43 (0.24)	0.39 (0.18)
Boundary layer height (m) (hi)	−0.31 (−0.22)	−0.35 (−0.30)
Rainfall total (mm) (hi)	0.08 (0.18)	−0.03 (−0.02)
Safdarjung inversion strength (°C) (lo)	−0.37 (−0.19)	−0.28 (−0.04)
Safdarjung inversion base (m AGL) (hi)	0.08 (0.07)	0.04 (−0.01)

inversion base computed on 29 pairs, not 30.

Day−on−day (n = 29 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.63 (−0.67)	−0.73 (−0.77)
Temperature 2 m (°C) (n)	0.18 (0.15)	0.20 (0.11)
Boundary layer height (m) (hi)	−0.53 (−0.56)	−0.60 (−0.59)
Rainfall total (mm) (hi)	−0.01 (0.06)	−0.22 (−0.12)
Safdarjung inversion strength (°C) (lo)	−0.11 (0.08)	−0.05 (0.20) *
Safdarjung inversion base (m AGL) (hi)	−0.09 (−0.19)	−0.07 (−0.08)

inversion base computed on 27 pairs, not 29. * marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

February 2026

Year−on−year (n = 28 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.23 (−0.28)	−0.24 (−0.22)
Temperature 2 m (°C) (n)	0.45 (0.47)	0.58 (0.62)
Boundary layer height (m) (hi)	0.40 (0.34)	0.34 (0.28)
Rainfall total (mm) (hi)	−0.27 (−0.28)	−0.27 (−0.25)
Safdarjung inversion strength (°C) (lo)	0.07 (0.03)	0.10 (0.08)
Safdarjung inversion base (m AGL) (hi)	−0.54 (−0.16) *	−0.48 (−0.08) *

inversion base computed on 27 pairs, not 28. * marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

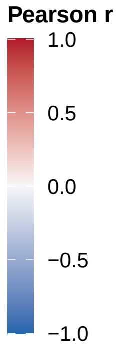

Day−on−day (n = 28 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.46 (−0.41)	−0.55 (−0.56)
Temperature 2 m (°C) (n)	0.00 (0.00)	0.19 (0.10)
Boundary layer height (m) (hi)	−0.35 (−0.24)	−0.39 (−0.31)
Rainfall total (mm) (hi)	0.35 (−0.09) *	0.17 (−0.12) *
Safdarjung inversion strength (°C) (lo)	0.40 (0.04) *	0.30 (0.01) *
Safdarjung inversion base (m AGL) (hi)	−0.26 (0.10) *	−0.16 (0.18) *

* marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

March 2026

Year−on−year (n = 31 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.26 (−0.28)	−0.01 (−0.06)
Temperature 2 m (°C) (n)	0.70 (0.75)	0.67 (0.70)
Boundary layer height (m) (hi)	0.39 (0.38)	0.49 (0.47)
Rainfall total (mm) (hi)	−0.15 (−0.31)	−0.26 (−0.39)
Safdarjung inversion strength (°C) (lo)	−0.03 (−0.00)	−0.05 (−0.00)
Safdarjung inversion base (m AGL) (hi)	−0.00 (−0.02)	0.02 (0.04)

Day−on−day (n = 31 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.60 (−0.57)	−0.65 (−0.64)
Temperature 2 m (°C) (n)	0.18 (0.09)	0.31 (0.16)
Boundary layer height (m) (hi)	−0.36 (−0.38)	−0.45 (−0.42)
Rainfall total (mm) (hi)	−0.15 (−0.44) *	−0.23 (−0.45)
Safdarjung inversion strength (°C) (lo)	0.25 (0.30)	0.28 (0.29)
Safdarjung inversion base (m AGL) (hi)	−0.29 (−0.10)	−0.45 (−0.21)

Pearson r: 1.0, 0.5, 0.0, −0.5, −1.0

* marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

April 2026

Year−on−year (n = 30 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.18 (−0.15)	−0.06 (−0.04)
Temperature 2 m (°C) (n)	0.65 (0.45)	0.74 (0.47) *
Boundary layer height (m) (hi)	0.12 (0.14)	0.03 (0.01)
Rainfall total (mm) (hi)	−0.45 (−0.32)	−0.40 (−0.20)
Safdarjung inversion strength (°C) (lo)	0.22 (0.37)	0.12 (0.29)
Safdarjung inversion base (m AGL) (hi)	−0.45 (−0.38)	−0.27 (−0.16)

* marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

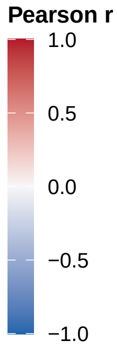

Day−on−day (n = 30 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.33 (−0.34)	−0.11 (−0.20)
Temperature 2 m (°C) (n)	0.38 (0.46)	0.45 (0.57)
Boundary layer height (m) (hi)	−0.05 (−0.03)	−0.01 (−0.04)
Rainfall total (mm) (hi)	−0.04 (−0.21)	−0.08 (0.01)
Safdarjung inversion strength (°C) (lo)	0.22 (0.17)	0.11 (0.01)
Safdarjung inversion base (m AGL) (hi)	−0.15 (−0.16)	0.02 (0.07)

May 2026

Year−on−year (n = 31 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.18 (−0.18)	−0.21 (−0.19)
Temperature 2 m (°C) (n)	0.66 (0.62)	0.78 (0.71)
Boundary layer height (m) (hi)	0.30 (0.25)	0.34 (0.28)
Rainfall total (mm) (hi)	−0.16 (−0.29)	−0.06 (−0.23)
Safdarjung inversion strength (°C) (lo)	0.06 (0.18)	0.08 (0.06)
Safdarjung inversion base (m AGL) (hi)	−0.36 (−0.48)	−0.25 (−0.41)

inversion base computed on 27 pairs, not 31.

Day−on−day (n = 31 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.46 (−0.40)	−0.24 (−0.16)
Temperature 2 m (°C) (n)	0.32 (0.32)	0.39 (0.41)
Boundary layer height (m) (hi)	−0.18 (−0.19)	−0.31 (−0.22)
Rainfall total (mm) (hi)	−0.32 (−0.28)	−0.29 (−0.22)
Safdarjung inversion strength (°C) (lo)	−0.01 (0.02)	−0.03 (−0.01)
Safdarjung inversion base (m AGL) (hi)	−0.27 (−0.18)	−0.06 (0.07)

inversion base computed on 29 pairs, not 31.

Pearson r
1.0
0.5
0.0
−0.5
−1.0

June 2026

Year−on−year (n = 30 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.21 (−0.18)	−0.15 (−0.02)
Temperature 2 m (°C) (n)	0.70 (0.68)	0.76 (0.76)
Boundary layer height (m) (hi)	0.43 (0.45)	0.51 (0.52)
Rainfall total (mm) (hi)	−0.27 (−0.31)	−0.32 (−0.44)
Safdarjung inversion strength (°C) (lo)	−0.09 (0.00)	−0.01 (0.10)
Safdarjung inversion base (m AGL) (hi)	−0.33 (−0.40)	−0.29 (−0.22)

inversion strength and inversion base computed on 15/26 pairs, not 30.

Day−on−day (n = 30 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.56 (−0.53)	−0.27 (−0.16)
Temperature 2 m (°C) (n)	0.21 (0.22)	0.51 (0.36)
Boundary layer height (m) (hi)	−0.08 (−0.10)	0.23 (0.03)
Rainfall total (mm) (hi)	−0.39 (−0.26)	−0.32 (−0.17)
Safdarjung inversion strength (°C) (lo)	−0.16 (−0.09)	−0.16 (−0.24)
Safdarjung inversion base (m AGL) (hi)	−0.14 (0.29) *	−0.15 (0.29) *

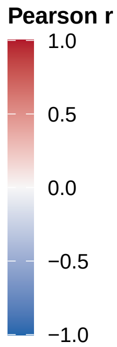

inversion strength and inversion base computed on 24/26 pairs, not 30. * marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

July 2026

Year−on−year (n = 13 days)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	−0.08 (−0.12)	0.47 (0.14) *
Temperature 2 m (°C) (n)	0.65 (0.78)	0.78 (0.93)
Boundary layer height (m) (hi)	0.02 (−0.05)	0.53 (0.28)
Rainfall total (mm) (hi)	−0.49 (−0.74)	−0.45 (−0.85) *

* marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

Day−on−day (n = 13 day−to−day changes)

	PM2.5	PM10
Wind speed 10 m (km/h) (hi)	0.13 (0.13)	0.31 (0.12)
Temperature 2 m (°C) (n)	0.61 (0.65)	0.78 (0.82)
Boundary layer height (m) (hi)	0.12 (0.18)	0.46 (0.36)
Rainfall total (mm) (hi)	−0.22 (−0.53) *	−0.20 (−0.59) *

* marks a cell where the two measures differ by more than 0.25; on a variable with many zero days, such as rainfall, the bracketed rank figure is the more trustworthy of the two.

Pearson r

1.0

0.5

0.0

−0.5

−1.0

Annexure: methodology in plain English

The daily building block. Every pollution number is a **24-hour average ending at 10 AM Indian time**, the same window Delhi's pollution control board uses, so a "day" in this report always means yesterday-morning to this-morning. A monitoring station only counts on a given day if it reported at least **16 of the 24 hours**, which keeps a half-broken sensor from skewing the city figure. The **city value** each day is the plain average across all stations that passed that test.

Pollutants covered, and why only two. This edition reports **PM2.5** and **PM10** only. Delhi's monitors also measure four gases, but the gas feed froze on **23 June 2026** and its numbers cannot be trusted, so the gas roses, gas tests and gas correlations are held back until the feed is fixed. The two particle measurements are independent of that fault.

Wind and weather. Wind speed, wind direction, temperature, rainfall and the boundary-layer height (how deep the mixing layer of the atmosphere is) come from the ERA5 / Open-Meteo weather record for Delhi, averaged over the same 24-hour-to-10-AM window. The morning temperature-inversion strength and height come from the Safdarjung radiosonde (weather balloon). For directional plots, each day's *dominant* wind direction is its most common hourly sector, and the day's mean wind angle is a proper circular average.

The wind-direction test. The honest question is whether *direction* explains pollution beyond what *speed* and atmospheric depth already explain. We fit a model of PM on wind direction, wind speed and boundary-layer height, and compare it against the same model without direction (a partial F-test). If direction adds nothing, its effect was really a speed effect wearing a directional badge. Kruskal-Wallis and ANOVA are shown alongside as the raw, unadjusted comparisons.

The correlation heatmaps. For each month we take the day-to-day *residual* (how far each day sat above or below normal) for every weather variable and every pollutant, and measure how tightly the two move together (Pearson correlation, with a rank-based Spearman check in brackets to guard against a few extreme days driving the result). The year-on-year residual compares 2026 against a smoothed 2025 baseline (the 2025 pollutant side uses a 48-hour window to damp single-day spikes); the day-on-day residual compares each day with the previous one. Two guards decide whether a cell is trustworthy: a drift-tolerant significance test (Newey-West) and a multiple-comparisons screen (Benjamini-Hochberg at 10%). Failing cells are greyed.

Two weather variables to read with care. Boundary-layer height is a coarse-grid reanalysis value for the whole city, so its link to pollution is a *regional* signal, not a neighbourhood one. Inversion strength and height come from a single morning balloon, a good gauge of regional stability but not of conditions at each monitor. The 2026 inversion record currently ends around late June, so the most recent month's inversion rows may be blank.

Provenance. Pollution: CREA compilation of CPCB / DPCC station data (raw hourly, long format). Weather: ERA5 / Open-Meteo hourly reanalysis for Delhi plus the Safdarjung radiosonde. All figures in this report were recomputed from those files at build time; none are carried over or hand-entered.

Prepared by Ajay Maken MP and Team.