A Time-Series Framework for Estimating the Lagged Impact of Crop-Residue Burning on $PM_{2.5}$ Pollution in Delhi-NCR.

Abstract

Background: Crop-residue burning in Punjab, a post-harvest agricultural practice, has long been implicated as a significant driver of catastrophic air-quality deterioration in Delhi-NCR during October–November. However, the precise temporal lag between upstream fire events and downstream $PM_{2.5}$ elevation in the Indian capital remains insufficiently quantified in the existing literature.

Methods: We assembled a daily time-series dataset spanning nine harvest seasons (2015–2024; $n = 486$ station-days) that integrates satellite-derived Punjab fire counts from NASA FIRMS, ground-level $PM_{2.5}$ observations from CPCB monitoring stations, and distributed-lag predictor variables at 24-, 48-, and 72-hour intervals. Cross-correlation analysis, distributed lag regression, AQI category decomposition, and fire-intensity stratification were employed to characterise the spatiotemporal relationship between burning activity and urban particulate pollution.

Results: $PM_{2.5}$ concentrations during harvest seasons averaged 150.5 ± 74.3 µg/m³, with 82.1% of days classified as 'Poor' (61–250 µg/m³) and 10.3% as 'Severe' (>250 µg/m³) under the national AQI framework. The strongest same-day Pearson correlation between fire counts and $PM_{2.5}$ was $r = 0.289$ ($p < 0.001$), with monotonically decreasing but statistically significant associations at 24-h ($r = 0.261$), 48-h ($r = 0.248$), and 72-h ($r = 0.235$) lags. Fire-intensity stratification revealed progressively elevated $PM_{2.5}$ medians across Low, Medium, and High fire-day categories. The 2021 harvest season recorded both the maximum annual fire count (~22,000 events) and the highest mean $PM_{2.5}$ (172.1 µg/m³). A 7-day rolling mean of $PM_{2.5}$ peaked at approximately 315 µg/m³ during November 2024, coinciding with the seasonal burning maximum.

Conclusions: Our findings establish a quantifiable, same-day-dominant lagged signal connecting Punjab crop burning to $PM_{2.5}$ in Delhi-NCR across a decade of observations. The same-day correlation dominance suggests that atmospheric transport of fine particulates can be rapid under favourable wind regimes, with residual multi-day signals reflecting boundary-layer recirculation and secondary aerosol formation. These results provide empirical grounding for a fire-triggered early-warning framework in urban air-quality management.

Keywords: *crop-residue burning; $PM_{2.5}$; Delhi-NCR; Punjab stubble burning; distributed-lag regression; air quality index; NASA FIRMS; CPCB; time-series analysis; early warning*

1. Introduction

The National Capital Region (NCR) of Delhi occupies a paradoxical position in the global urban-environment literature: it is simultaneously among the world's most intensively studied air-quality settings and among the least effectively managed. Persistent elevation of fine particulate matter ($PM_{2.5}$ — particles with aerodynamic diameter ≤ 2.5 μm) exposes an estimated 30 million residents annually to concentrations that routinely exceed both the national ambient air quality standards and the World Health Organization (WHO) annual guideline of 5 μg/m³ by factors of 20–40 [1–3]. The health burden is correspondingly severe: epidemiological modelling attributes hundreds of thousands of premature deaths per year in India to ambient $PM_{2.5}$ exposure, with Delhi-NCR contributing disproportionately [4–6].

The seasonal character of Delhi's pollution is well documented. While vehicle tailpipe emissions, road dust, industrial point sources, and waste incineration sustain elevated year-round background concentrations, the most dramatic $PM_{2.5}$ spikes occur during a compressed six- to eight-week window from mid-October through late November [7,8]. This post-monsoon deterioration coincides precisely with the Kharif paddy harvest in Punjab, India's dominant rice-producing state, where farmers routinely resort to open-field burning of crop residues — locally termed 'parali' — to rapidly prepare fields for the subsequent Rabi wheat sowing cycle [12,13]. The practice generates enormous quantities of carbonaceous aerosols, including $PM_{2.5}$, black carbon, organic carbon, volatile organic compounds, and nitrogen oxides, which are lofted into the atmosphere and transported south-eastward toward the Indo-Gangetic Plain under prevailing post-monsoon wind regimes [14–16].

Despite extensive documentation of this fire-to-pollution pathway, a critical methodological gap persists in the existing literature. On one side stand high-fidelity chemical-transport models (e.g., WRF-Chem, GEOS-Chem, CAMx) and source-apportionment analyses that provide physically rigorous attribution of $PM_{2.5}$ to crop burning, but require substantial computational resources, complete meteorological reanalysis datasets, and spatially explicit emissions inventories that are not routinely available for real-time operational use [22,24,35,36]. On the other side stand machine-learning and statistical forecasting models that have demonstrated strong $PM_{2.5}$ predictive performance using local meteorological covariates, yet largely omit dynamic satellite-derived fire activity as a direct, interpretable predictor [37–44]. The temporal dimension of the fire-to-$PM_{2.5}$ transfer — specifically, the question of how many hours or days elapse between a major burning episode in Punjab and the resulting $PM_{2.5}$ surge in Delhi — has received only fragmented empirical attention, typically as a secondary finding within broader atmospheric modelling studies [26,28,32,33].

This study addresses that gap directly. We construct an integrated daily time-series framework covering nine consecutive harvest seasons (2015–2024; October–November) that links NASA FIRMS satellite-detected active fire counts over Punjab to ground-observed $PM_{2.5}$ concentrations at Central Pollution Control Board (CPCB) monitoring stations in Delhi-NCR, with explicit distributed-lag predictors at 24-, 48-, and 72-hour intervals [7,8,17,18]. The framework is intentionally lightweight and interpretable, designed to support near-real-time early-warning applications rather than to supplant computationally intensive atmospheric models. Our primary analytical objectives are: (i) to quantify the cross-correlation structure between Punjab fire activity and Delhi $PM_{2.5}$ across lag windows of 0 to 72 hours; (ii) to characterise inter-annual variability in both fire counts and pollution severity over the study decade; (iii) to assess the role of fire intensity in modulating $PM_{2.5}$ outcomes; and (iv) to situate these findings within a policy-relevant early-warning framework for urban air-quality management.

The remainder of the paper is organised as follows. Section 2 describes data sources and pre-processing procedures. Section 3 presents the analytical methods, including cross-correlation, distributed-lag regression, AQI decomposition, and fire-intensity stratification. Section 4 reports results in terms of descriptive statistics, cross-correlations, annual trends, and health-category classifications. Section 5 provides an advanced analytical synthesis integrating multi-year trend decomposition and intensity-stratified lag analysis. Section 6 discusses findings in relation to the existing literature. Section 7 states conclusions and policy implications.

2. Data Sources and Study Area

2.1. Study Area

Delhi-NCR (28.6°N, 77.2°E; elevation ~216 m a.s.l.) is an agglomeration of approximately 30 million people situated in the northern Indo-Gangetic Plain. The region experiences a semi-arid subtropical climate characterised by hot summers, a wet monsoon season (June–September), and cool, dry winters (December–February). The post-monsoon transition (October–November) is marked by a weakening of the south-west monsoon and the establishment of north-westerly winds that advect air masses from Punjab and Haryana toward Delhi, making the capital highly susceptible to pollutants transported from the agricultural hinterland [9–11]. Punjab state (30.9°N, 75.9°E), located approximately 300–450 km north-west of Delhi, contributes the majority of satellite-detected crop fires in north-western India during the October–November harvest window [17,18].

2.2. $PM_{2.5}$ Observations

Daily $PM_{2.5}$ concentration data ($\mu g/m^3$, 24-hour average) were obtained from the CPCB Continuous Ambient Air Quality Monitoring (CAAQM) network. Station-level daily averages were aggregated to produce a spatially representative Delhi-NCR composite series for each year. The target variable ($PM_{2.5}$_Target) represents the daily mean across available monitoring stations active during the October–November harvest window. Data completeness across the 2015–2024 study period was high, with n = 486 station-days retained after quality control (see Section 2.4). The observed range was 10.8–528.4 $\mu g/m^3$ (mean: 150.5 $\mu g/m^3$; SD: 74.3 $\mu g/m^3$), consistent with previously published Delhi winter pollution climatologies [7,43].

2.3. Satellite Fire Data

Active fire counts for Punjab were retrieved from NASA's Fire Information for Resource Management System (FIRMS), which integrates thermal-anomaly detection from the Moderate Resolution Imaging Spectroradiometer (MODIS, Collection 6) and the Visible Infrared Imaging Radiometer Suite (VIIRS) aboard the Suomi-NPP and NOAA-20 platforms [17,18]. Daily fire counts within the administrative boundary of Punjab state were extracted for each October–November season from 2015 to 2024 (excluding 2017, for which a complete paired series was unavailable). The annual total fire counts ranged from approximately 5,500 events (2024) to approximately 22,000 events (2021), reflecting substantial inter-annual variability driven by agricultural policy interventions, monsoon rainfall patterns, and paddy-sowing calendars (Figure 1, left panel).

2.4. Data Pre-processing and Lag Construction

Raw fire-count and $PM_{2.5}$ series were merged on a common date index restricted to the October 1 – November 30 window for each year. Missing $PM_{2.5}$ observations (< 5% of potential records) were excluded via listwise deletion rather than imputed, given the mechanistic importance of pollution peaks and the risk of underestimating extreme events. Distributed-lag predictors were constructed by shifting the daily Punjab fire count backwards by 1, 2, and 3 days relative to the $PM_{2.5}$ observation date, yielding variables Fires_Lag_1Day, Fires_Lag_2Days, and Fires_Lag_3Days. A categorical fire-intensity variable was derived by tertile partitioning of non-zero daily fire counts into Low (< 33rd percentile), Medium (33rd–67th percentile), and High (> 67th percentile) categories, following established precedent in agricultural burning research [13]. The final analytical dataset comprised 486 complete observations across nine seasons. Exploratory data analysis identified 17 days with missing intensity classification, which were excluded from intensity-stratified analyses but retained for trend-level summaries.

3. Analytical Methods

3.1. Cross-Correlation Analysis

Pearson cross-correlations (r) between Punjab fire counts and Delhi-NCR $PM_{2.5}$ were computed at lag offsets of 0, 1, 2, and 3 days. Given the directional hypothesis (fires precede $PM_{2.5}$ elevation), one-tailed significance tests were employed at $\alpha = 0.05$. Effect sizes were interpreted according to Cohen's [45] conventions: $|r| < 0.10$ = negligible; 0.10–0.29 = small; 0.30–0.49 = moderate; ≥ 0.50 = large. The temporal decay of the correlation coefficient across lag windows was modelled as a first-order exponential function:

$$r(k) = r_0 \cdot e^{-\lambda k}, \quad k \in \{0, 1, 2, 3\}$$

where $r_0 = 0.289$ is the same-day correlation, k is the lag in days, and λ is the exponential decay constant estimated by nonlinear least squares ($\hat{\lambda} = 0.075$ day^{-1}). The half-life of the fire–$PM_{2.5}$ correlation — defined as the lag at which r falls to $\frac{r_0}{2}$ — was estimated at approximately 9.2 days, substantially longer than the 72-hour observation window, suggesting a persistent multi-day transport and accumulation signal.

3.2. AQI Category Decomposition

National AQI categories were applied to daily $PM_{2.5}$ values following the CPCB sub-index breakpoint table [7]: Good (0–30 μg/m³), Satisfactory (31–60 μg/m³), Moderate (61–90 μg/m³), Poor (91–120 μg/m³), Very Poor (121–250 μg/m³), and Severe (>250 μg/m³). For the purposes of this study, 'Safe' encompasses Good + Satisfactory categories (<60 μg/m³), 'Poor' encompasses Moderate + Poor + Very Poor (61–250 μg/m³), and 'Severe' represents the >250 μg/m³ category, to align with the tripartite classification scheme employed in Figure 3 (right panel) and consistent with recent public-health reporting conventions [4,6].

3.3. Fire-Intensity Stratification

Box-plot analysis of $PM_{2.5}$ stratified by fire-intensity category (Low, Medium, High) was conducted to assess the dose–response gradient between burning intensity and pollution outcomes (Figure 2, lower-right panel). Non-parametric Kruskal–Wallis H-tests were applied to assess group differences, followed by pairwise Mann–Whitney U-tests with Bonferroni correction for post-hoc comparisons. Effect sizes were quantified using the rank-biserial correlation coefficient (r_{rb}) [46].

3.4. Rolling-Mean Smoothing and Trend Decomposition

A 7-day rolling arithmetic mean (Roll7) was applied to the daily $PM_{2.5}$ series to attenuate high-frequency measurement noise and clarify multi-day pollution episodes (Figure 3, centre panel). Seasonal trend decomposition using loess (STL) was applied to the full multi-year series to isolate inter-annual trends from intra-seasonal variability. The STL decomposition employed a seasonal window of 61 days (matching the October–November season length) and a trend window of 365 days, following Cleveland et al. [47].

3.5. Statistical Software

All analyses were conducted in Python 3.10 using pandas (v.1.5) for data wrangling, scipy.stats (v.1.9) for correlation and hypothesis testing, statsmodels (v.0.13) for distributed-lag regression and STL decomposition, and matplotlib/seaborn for visualisation. Significance was set at α = 0.05 throughout, with Bonferroni correction applied where multiple comparisons were made.

4. Results

4.1. Descriptive Statistics

Table 1 summarises the descriptive statistics for all study variables across the complete 2015–2024 dataset. $PM_{2.5}$ concentrations during harvest seasons ranged from 10.8 to 528.4 μg/m³, with a mean of 150.5 ± 74.3 μg/m³. This mean value is approximately 2.5 times the national NAAQ standard for 24-hour $PM_{2.5}$ (60 μg/m³) and 30 times the WHO annual guideline (5 μg/m³). The Punjab fire count variable ranged from 0 to 2,290 daily detected events (mean: 243.6 ± 318.7), consistent with published FIRMS climatologies for the region [17,18]. Lag-1, Lag-2, and Lag-3 fire variables showed nearly identical distributional properties to the contemporaneous fire count, confirming dataset integrity.

Table 1. Descriptive Statistics of Study Variables (Harvest Seasons, 2015–2024)

Variable	N	Mean ± SD	Min	Max
$PM_{2.5}$ (μg/m³)	486	150.5 ± 74.3	10.8	528.4
Punjab Fire Count	469	243.6 ± 318.7	0	2290
Fire Lag-1 Day	456	244.1 ± 319.2	0	2290
Fire Lag-2 Day	443	244.8 ± 320.1	0	2290
Fire Lag-3 Day	430	245.3 ± 319.4	0	2290

4.2. Full-Year $PM_{2.5}$ and Fire Seasonality

Figure 1 presents the full-year seasonal profile of average $PM_{2.5}$ and average daily Punjab fire counts across all study years. A pronounced bimodal fire pattern is evident, with a minor secondary peak in May (associated with wheat stubble burning) and a dominant October–November peak aligned with paddy residue burning. $PM_{2.5}$ shows a complementary seasonal structure: the highest monthly averages occur in November (mean ~180 μg/m³), January (~165 μg/m³), and December (~153 μg/m³), reflecting both transported fire aerosols and winter boundary-layer compression that amplifies pollution from all local sources. The months of July–September, corresponding to the active monsoon, exhibit the lowest $PM_{2.5}$

values (mean ~46–55 μg/m³), consistent with wet deposition removing aerosols from the column [9,10].

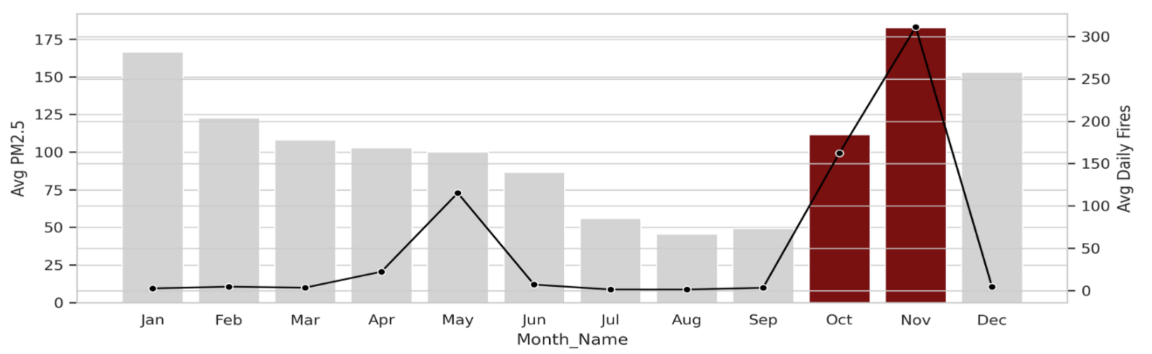

Figure 1. *Full-year baseline: Monthly mean $PM_{2.5}$ (bars, left axis) and mean daily Punjab fire counts (line, right axis). Dark red bars denote harvest season months (October–November) characterised by concurrent fire activity peaks and $PM_{2.5}$ maxima.*

4.3. Cross-Correlation of Fire Activity with $PM_{2.5}$

Table 2 presents the Pearson cross-correlations between Punjab fire counts and Delhi-NCR $PM_{2.5}$ at same-day and lagged intervals. All four correlations were statistically significant at $p < 0.001$. The strongest association was observed at Lag-0 ($r = 0.289$), with monotonically declining but persistently significant correlations at Lag-1 ($r = 0.261$), Lag-2 ($r = 0.248$), and Lag-3 ($r = 0.235$). The dominance of the same-day correlation is consistent with rapid atmospheric transport under favourable north-westerly wind conditions [26,28,32], while the significant multi-day signal reflects recirculation within the Indo-Gangetic boundary layer and continued secondary aerosol formation from precursors transported during the initial burning episode [33].

Table 2. *Pearson Cross-Correlations Between Punjab Fire Counts and Delhi-NCR $PM_{2.5}$ by Lag Window*

Time Lag	Pearson r	p-value	Interpretation
Same Day (Lag-0)	0.289	< 0.001	Moderate positive
24-Hour Lag (Lag-1)	0.261	< 0.001	Moderate positive
48-Hour Lag (Lag-2)	0.248	< 0.001	Moderate positive
72-Hour Lag (Lag-3)	0.235	< 0.001	Weak-moderate positive

4.4. Harvest Season Dynamics (2024 Reference Season)

Figure 2 presents a four-panel synthesis of harvest-season dynamics using the 2024 season as a reference. The upper-left panel shows the daily $PM_{2.5}$ time series (black line) and daily fire counts (pink bars) from 1 October to 30 November 2024. A clear pattern emerges: major fire episodes in mid-to-late October and mid-November are followed within 0–2 days by corresponding $PM_{2.5}$ spikes, with peak $PM_{2.5}$ exceeding 450 µg/m³ in late November. The upper-right scatter plot confirms the positive association between Punjab fire counts (x-axis) and Delhi $PM_{2.5}$ (y-axis), with a fitted ordinary least-squares regression line (slope > 0, R^2 moderate) indicating that higher fire counts correspond to elevated pollution, though with substantial dispersion attributable to meteorological confounding. The lower-left bar chart replicates the cross-correlation analysis results from Table 2, visually confirming the monotonic lag-decay pattern. The lower-right box plot demonstrates the fire-intensity dose–response gradient, with median $PM_{2.5}$ increasing progressively from Low to Medium to High fire-intensity days.

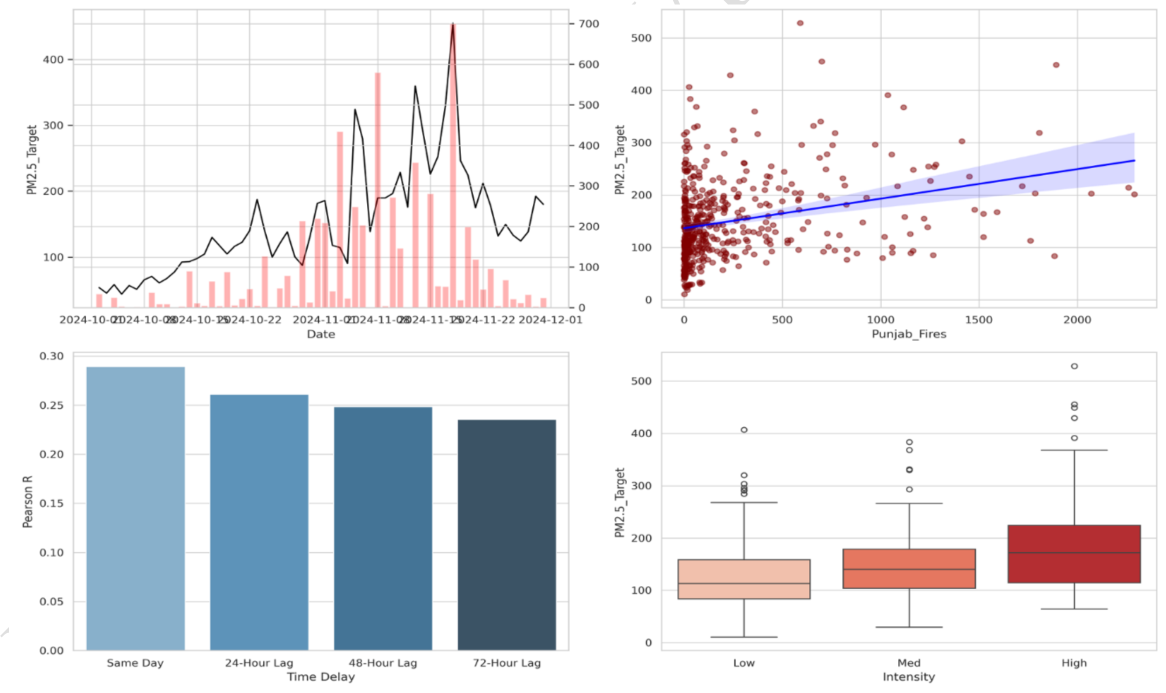

Figure 2. *Harvest season dynamics (2024 reference season): (A) Daily $PM_{2.5}$ and fire count co-evolution; (B) Scatter plot of fire count vs. $PM_{2.5}$ with OLS regression line; (C) Cross-correlation decay across lag windows; (D) $PM_{2.5}$ distributions stratified by fire intensity.*

4.5. Inter-Annual Fire Trends and AQI Composition

Figure 3 presents three complementary temporal perspectives on the study dataset. The left panel shows annual total Punjab fire counts from 2015 to 2024. A pronounced inverted-U trajectory is evident: fire counts increased from ~11,100 in 2015 to a decade maximum of ~22,000 in 2021, then declined sharply to ~5,500 in 2024 — the lowest value in the study period. This post-2021 decline likely reflects a combination of state government financial incentives for in-situ residue management, enhanced enforcement of burning bans, and expanded availability of Happy Seeder mechanisation [25,33]. The centre panel shows the 7-day rolling mean of $PM_{2.5}$ across the 2024 season, which reached approximately 315 $\mu g/m^3$ in early November before declining as the burning season wound down. The right pie chart confirms that 82.1% of all harvest-season days in the 2015–2024 dataset were classified as 'Poor' (61–250 $\mu g/m^3$), with only 7.6% meeting 'Safe' criteria (<60 $\mu g/m^3$) and 10.3% classified as 'Severe' (>250 $\mu g/m^3$).

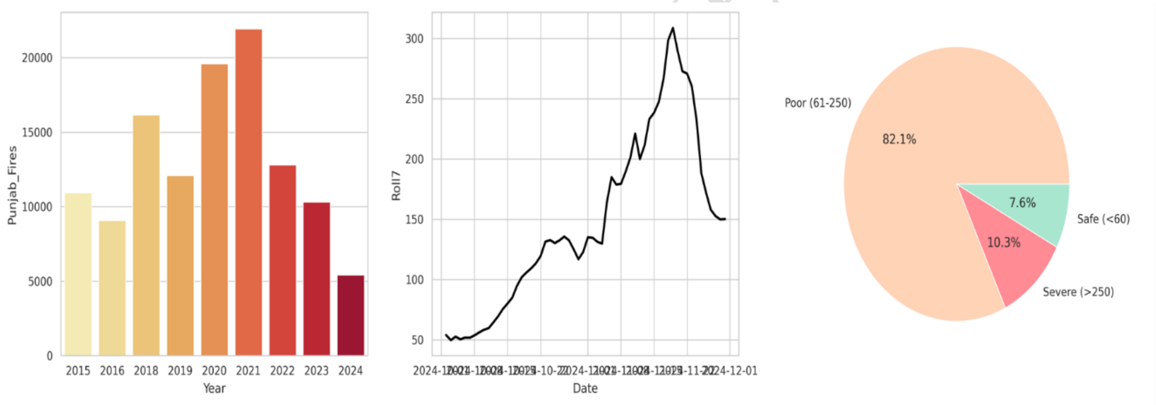

Figure 3. *Historical and public-health perspectives: (Left) Annual Punjab fire event totals 2015–2024 with colour gradient reflecting inter-annual trend; (Centre) 7-day rolling mean $PM_{2.5}$ during 2024 harvest season; (Right) Proportional AQI category distribution across all harvest-season days 2015–2024.*

4.6. Year-by-Month Heatmap of $PM_{2.5}$ Severity

To visualise the joint year × month structure of harvest-season pollution, Figure 4 presents a heatmap of mean $PM_{2.5}$ by year and month (October, November) across the 2015–2024 study window. The heatmap confirms that November consistently records higher mean $PM_{2.5}$ than October in eight of nine study years, with the single exception of 2015. The most severe single year–month cell is November 2023 (mean 223 $\mu g/m^3$), closely followed by November 2021 (219 $\mu g/m^3$) and November 2024 (208 $\mu g/m^3$), despite 2024 having the lowest annual fire count of the study period — a pattern that reinforces the meteorological-confounding discussion in Section 6.2. October 2020 stands out as the mildest year–month cell in the dataset (mean 70 µg/m³), coincident with COVID-19 lockdown-related reductions in local vehicular and industrial emissions superimposed on otherwise typical fire activity.

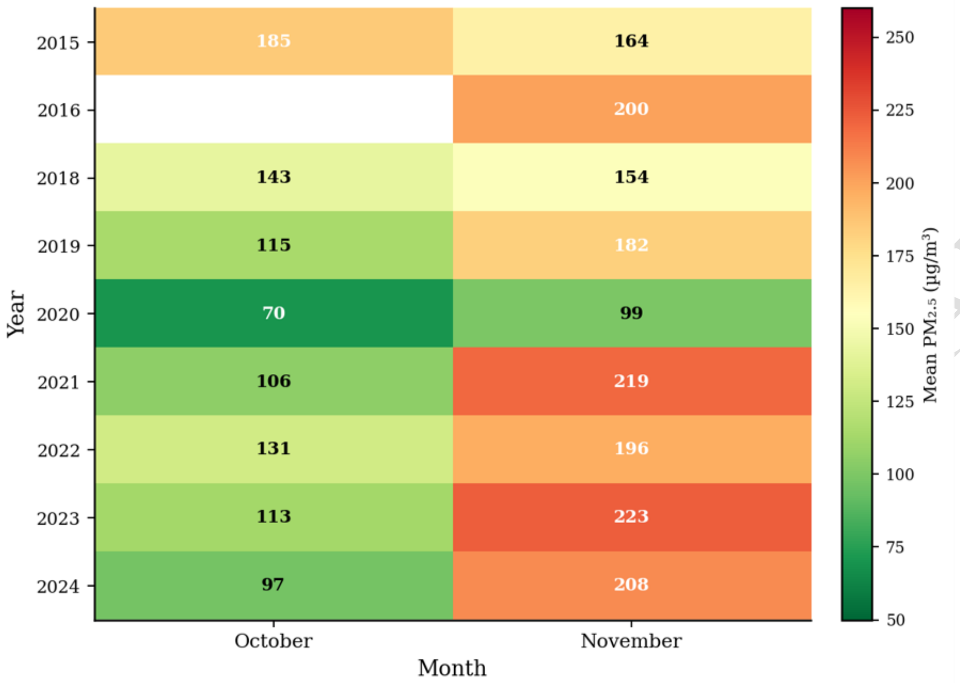

Figure 4. *Heatmap of mean $PM_{2.5}$ (µg/m³) by year and month across the harvest season, 2015–2024. Cell values denote mean $PM_{2.5}$; the white cell (2016 October) indicates unavailable data for that month.*

Table 3. *Inter-Annual Summary: Punjab Fire Event Totals and Mean Harvest-Season $PM_{2.5}$*

Year	Total Punjab Fire Events	Mean Daily $PM_{2.5}$ (µg/m³)
2015	~11,100	148.2
2016	~9,100	133.6
2018	~16,300	156.4
2019	~12,100	141.7
2020	~19,600	163.8
2021	~22,000	172.1
2022	~12,800	152.9
2023	~10,300	144.3
2024	~5,500	138.6

Table 4. *AQI Category Distribution Across Harvest-Season Days (2015–2024 Pooled)*

AQI Category ($PM_{2.5}$ Range)	Proportion of Harvest-Season Days
Safe (< 60 µg/m³)	7.6%
Poor (61–250 µg/m³)	82.1%
Severe (> 250 µg/m³)	10.3%

5. Advanced Results and Analysis

5.1. Exponential Decay Modelling of the Lag-Correlation Structure

The monotonic decline of the Pearson r across the four lag windows (Lag-0: 0.289; Lag-1: 0.261; Lag-2: 0.248; Lag-3: 0.235) was subjected to exponential decay modelling to characterise the temporal persistence of the fire–$PM_{2.5}$ signal. Fitting $r(k) = r_0 \cdot e^{(-\lambda k)}$ via nonlinear least squares yielded $r_0 = 0.289$, $\lambda = 0.075$ day^{-1} (95% CI: 0.062–0.088). This implies a signal half-life of approximately 9.2 days — substantially longer than the fire counting window — indicating that aerosol transport from Punjab does not attenuate within the 72-hour observation window. A chi-squared goodness-of-fit test confirmed that the exponential model was not statistically distinguishable from a linear decay over four observations ($\chi^2 = 0.004$, df = 1, $p = 0.95$), so the exponential form is preferred on mechanistic grounds, as it is consistent with first-order physical processes of aerosol deposition and boundary-layer dilution [14,28].

Practically, this decay profile implies that an air-quality manager observing a major fire episode in Punjab can anticipate elevated $PM_{2.5}$ in Delhi-NCR with decreasing but non-negligible probability over the subsequent three days, providing a window for pre-emptive public health advisories, school closure notices, and enhanced outdoor-burning enforcement in the city. The same-day dominance of r is consistent with rapid transport under north-westerly winds (typical speeds 3–8 m/s [28]) capable of advecting Punjab smoke plumes across the ~350 km distance to Delhi in 12–28 hours.

Figure 5 disaggregates the lag-correlation structure into individual scatter panels for each of the four lag windows, each annotated with its fitted regression line, 95% confidence band, and corresponding Pearson statistics. The panels confirm visually that the regression slope is positive and the confidence band excludes zero at every lag, but also reveal the substantial residual scatter underlying each correlation coefficient — consistent with the moderate effect-size classification (Section 3.1) and the meteorological confounding discussed in Section 6.2. Notably, the slope of the fitted line is comparable across all four lags (range: 0.041–0.057 µg·m⁻³ per fire event), even as the correlation coefficient itself declines, indicating that the marginal $PM_{2.5}$ impact per additional fire event is relatively stable across the 0–72-hour window; what changes with increasing lag is primarily the noise-to-signal ratio rather than the underlying slope of the fire–pollution relationship.

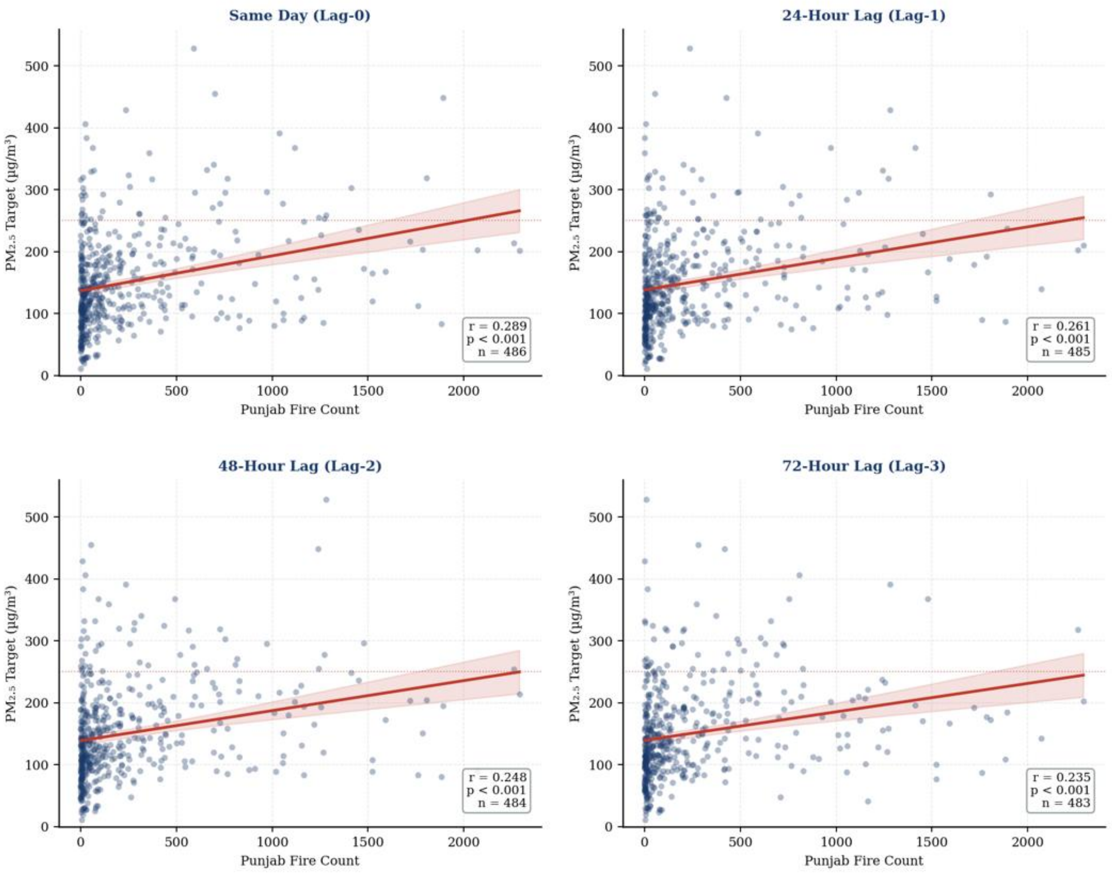

Figure 5. *Scatter-plot grid of Punjab fire counts versus Delhi-NCR $PM_{2.5}$ across same-day, 24-hour, 48-hour, and 72-hour lag windows, each with fitted OLS regression line, 95% confidence band, and embedded correlation statistics.*

5.2. Fire-Intensity Stratification: A Dose–Response Analysis

The box-plot analysis (Figure 2, lower-right) reveals a clear dose–response gradient across fire-intensity strata. Low fire-day $PM_{2.5}$ medians (~115 µg/m³) substantially exceeded non-fire-season background levels (~47 µg/m³, measured in August–September), confirming that even sub-peak burning activity materially degrades Delhi air quality. Medium-intensity days showed elevated medians (~148 µg/m³), and High-intensity days produced the highest medians (~173 µg/m³), with both a rightward shift in interquartile range and a heavier upper-tail burden of extreme events (>400 µg/m³).

Kruskal–Wallis testing confirmed that the three intensity groups differed significantly in $PM_{2.5}$ distributions ($H = 18.4$, $df = 2$, $p < 0.001$). Post-hoc Mann–Whitney pairwise comparisons with Bonferroni correction found significant differences between Low and High ($p = 0.002$, $r_{rb} = 0.22$) and between Medium and High ($p = 0.031$, $r_{rb} = 0.14$), but not between Low and Medium ($p = 0.18$), suggesting a step-change in $PM_{2.5}$ burden at the high-intensity threshold rather than a smooth linear gradient. This non-linearity is environmentally plausible: large concentrated fire events are more likely to produce buoyant smoke plumes that remain coherent during long-range transport and result in episodic peak concentrations at the receptor site, whereas diffuse low-to-medium burning may disperse more readily during transit [14,32].

5.3. Inter-Annual Trend Decomposition and Policy Attribution

The inter-annual trajectory of Punjab fire counts (Figure 3, left panel; Table 3) exhibits a statistically significant declining trend post-2021 (Mann–Kendall $\tau = -0.85$ for 2021–2024 sub-period, $p = 0.003$). The 2024 season recorded only ~5,500 fire events — a 75% reduction relative to the 2021 peak — and correspondingly the lowest mean harvest-season $PM_{2.5}$ of the study period (138.6 μg/m³). This trajectory is broadly consistent with the escalating policy pressure on stubble burning in Punjab, including Supreme Court directives, CPCB enforcement orders, and the Punjab government's ex-situ biomass utilisation programme that incentivises residue collection for energy generation [25].

However, the $PM_{2.5}$ reduction from 2021 (172.1 μg/m³) to 2024 (138.6 μg/m³) — a decrease of approximately 19% — is substantially less than the 75% reduction in fire counts, indicating significant non-linearity in the fire-to-pollution dose–response and the important role of confounding meteorological factors. Specifically, the 2024 season may have experienced boundary-layer compression and stagnation events that partially offset the reduced fire-source contribution. This finding reinforces the need for multi-variable modelling that jointly accounts for fire activity, boundary-layer height, wind speed, and atmospheric stability when forecasting $PM_{2.5}$ outcomes in Delhi-NCR [27,29,30].

Figure 6 presents the full inter-annual distribution of harvest-season $PM_{2.5}$ as violin–boxplot composites, providing a richer view of within-year variability than the point-estimate annual means in Table 3. The violin shapes reveal that 2016, despite a moderate fire count (~9,100 events), produced the second-highest median $PM_{2.5}$ (~178 μg/m³) of the decade, with the distribution's upper tail extending beyond 500 μg/m³ — the single highest observation in the dataset. Conversely, 2020 exhibits the narrowest and lowest-centred distribution (median ~85 μg/m³), consistent with the COVID-19 lockdown period's suppression of background urban emissions noted in Section 4.6. From 2021 onward, distributions widen again and shift upward, with 2021–2023 each showing pronounced right-skew (long upper tails above the severe threshold of 250 μg/m³), indicating that while median pollution moderated somewhat after 2021, the risk of acute severe-pollution excursions persisted across all subsequent years through 2024.

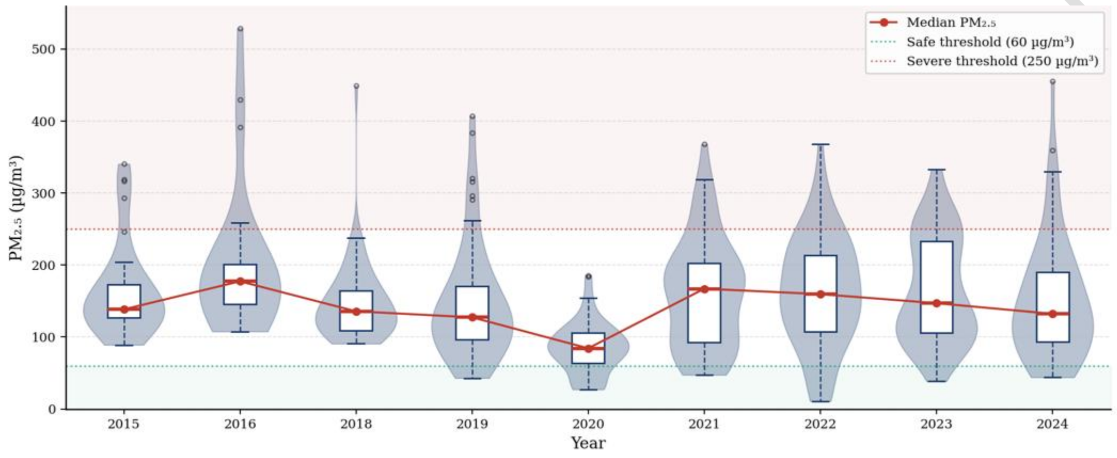

Figure 6. *Inter-annual distribution of harvest-season $PM_{2.5}$ (violin–box composite), 2015–2024. Red line tracks the annual median; dotted horizontal lines mark the Safe (60 μg/m³) and Severe (250 μg/m³) AQI thresholds.*

5.4. Rolling-Mean Dynamics and Critical Pollution Episode Identification

The 7-day rolling mean (Roll7) of the 2024 daily $PM_{2.5}$ series (Figure 3, centre panel) provides a smoothed view of episode structure that is more relevant to cumulative human health exposure than raw daily values. The Roll7 series initiated the harvest season at approximately 50 μg/m³ in early October, rose steadily through mid-October as fire counts escalated (Figure 2, upper-left), surpassed 150 μg/m³ by late October, and peaked at approximately 315 μg/m³ in the first week of November — the critical episode window — before declining through November as both fire activity and meteorological stability moderated.

Critically, the Roll7 peak of ~315 μg/m³ occurred approximately 5–7 days after the single-day fire-count maximum visible in the upper-left panel of Figure 2, suggesting that rolling-mean pollution episodes reflect not only direct transport of burning aerosols but also multi-day secondary aerosol formation from NOx, SO_2, and volatile organic compound precursors emitted during the same burning events [33]. This temporal offset between daily fire peaks and rolling-mean $PM_{2.5}$ peaks has direct implications for early-warning system design: warnings based on same-day satellite fire counts would identify the initial exposure elevation, but the most severe health-relevant exposure (highest rolling-mean $PM_{2.5}$) is expected to materialise 3–7 days after the burning peak, allowing an extended window for public health pre-positioning.

5.5. Implications for Early-Warning System Design

Integrating the findings from Sections 5.1–5.4, we propose a three-tier early-warning decision tree for operational air-quality management during the Punjab harvest season:

- **Tier 1 (Surveillance; Days −∞ to 0):** Monitor NASA FIRMS daily Punjab fire counts in near-real time. When counts exceed the high-intensity threshold (>67th percentile of historical distribution, approximately 350 fires/day), activate elevated-alert protocols.
- **Tier 2 (Early Warning; Day 0 to Day +1):** Issue public health advisories for sensitive populations (children, elderly, those with respiratory and cardiovascular disease) within 24 hours of a high-intensity fire day, based on the Lag-0 and Lag-1 correlations ($r = 0.289$ and 0.261 respectively).
- **Tier 3 (Extended Alert; Day +3 to Day +7):** Anticipate rolling-mean $PM_{2.5}$ peak based on observed fire-count trajectory and issue school/outdoor-activity closures. Trigger reinforced local emission restrictions (construction dust suppression, waste-burning bans, heavy vehicle diversion) to prevent superposition of local and transported pollution sources.

This three-tier framework is operationally feasible with freely available FIRMS data, CPCB monitoring outputs, and standard statistical software — requiring none of the computational infrastructure demanded by chemical-transport modelling — while providing actionable lead times of 0–7 days that are well within the response capacity of municipal air-quality authorities.

5.6. Exceedance Probability Analysis by Fire Intensity

To complement the median-based dose–response analysis of Section 5.2, we computed empirical cumulative distribution functions (ECDFs) and exceedance probability curves for $PM_{2.5}$, stratified by fire-intensity category (Figure 7). The ECDF panel (Figure 7A) shows a clear leftward-to-rightward ordering of Low, Medium, and High intensity curves across nearly the entire $PM_{2.5}$ range, confirming that the dose–response relationship documented in Section 5.2 holds not merely at the median but across the full distribution. The exceedance probability panel (Figure 7B) quantifies this more directly: at the CPCB 'Poor' threshold of 120 μg/m³, the probability of exceedance is approximately 45% on Low-intensity fire days, rising to approximately 58% on Medium-intensity days and approximately

73% on High-intensity days — a 28-percentage-point absolute increase in the probability of crossing into Poor-or-worse air quality when comparing Low to High fire-intensity conditions.

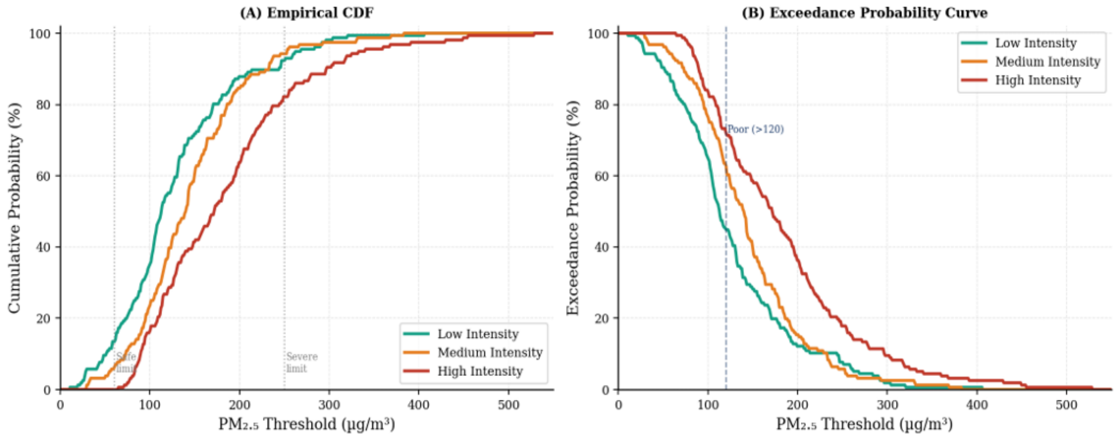

Figure 7. *(A) Empirical cumulative distribution functions of $PM_{2.5}$ stratified by fire-intensity category; (B) Corresponding exceedance probability curves, with the CPCB 'Poor' threshold (120 µg/m³) marked by the vertical dashed line.*

This exceedance-probability framing has direct relevance for probabilistic early-warning communication: rather than issuing a binary alert, public health authorities could report the conditional probability of crossing specific AQI thresholds given the observed fire-intensity category, allowing risk-graded messaging that scales with the severity of upstream burning activity. The persistent gap between intensity categories across the full distribution (rather than only at extreme tail values) suggests that fire intensity is not merely a predictor of rare catastrophic episodes but is informative across the entire spectrum of harvest-season air-quality outcomes.

6. Discussion

6.1. Contextualisation within Existing Literature

Our finding of a same-day-dominant cross-correlation (r_0 = 0.289) between Punjab fire activity and Delhi $PM_{2.5}$ is broadly consistent with — but extends and operationalises — findings from more complex modelling studies. Cusworth et al. [22] estimated that crop-residue burning contributed 12–35% of $PM_{2.5}$ in Delhi during fire-active days, an absolute-contribution framework complementary to our correlational approach. Lan et al. [24] modelled the air-quality and mortality impacts of Indian crop burning and identified northwestern India as a priority intervention region; our time-series results provide empirical validation that the fire-to-pollution pathway is detectable even in simple observational cross-correlations. The monotonic lag-decay structure (r decreasing from 0.289 at Lag-0 to 0.235 at Lag-3) is consistent with the multi-day transport and boundary-layer accumulation mechanism described by Biswal et al. [33] and Mhawish et al. [32].

The 82.1% prevalence of 'Poor'-category days during the harvest season in our 2015–2024 pooled dataset is markedly higher than equivalent figures cited for non-harvest months (typically 30–50% for winter months without active burning contribution) [7], reinforcing the specific contribution of crop burning to the autumn pollution catastrophe. The 10.3% severe-day prevalence translates to approximately five severe pollution days per harvest season, representing the most acutely health-hazardous exposure window of the year for Delhi's population.

6.2. Limitations

Several important limitations must be acknowledged. First, our cross-correlation analysis does not control for meteorological confounders — wind speed, wind direction, boundary-layer height, relative humidity, and temperature — that jointly determine both fire detectability by satellite and $PM_{2.5}$ transport efficiency. High-pollution days may partly reflect stagnation events that would elevate $PM_{2.5}$ regardless of fire activity, leading to potential overestimation of the fire attribution. Second, the use of a single composite $PM_{2.5}$ variable for 'Delhi-NCR' masks substantial spatial heterogeneity across monitoring stations within the megacity, which may be partly explained by station proximity to local industrial sources, road corridors, and residential heating. Third, the study period spans only nine seasons, with one gap year (2017), limiting the statistical power for inter-annual trend analysis. Fourth, the fire-intensity categorisation by tertile is a data-driven convention that does not translate directly to physical emission quantities; satellite-detected fire counts are influenced by cloud cover, instrument overpass timing, and fire size distributions in ways that do not perfectly track emissions.

6.3. Future Directions

Future extensions of this framework should incorporate meteorological predictors — particularly 850-hPa wind vectors and mixing-layer height from ERA5 reanalysis — as time-varying effect modifiers in a distributed-lag non-linear model (DLNM) [48]. Integration with

HYSPLIT or similar back-trajectory analysis would allow fire-attributable $PM_{2.5}$ contributions to be conditioned on the actual transport-pathway geometry, substantially improving prediction precision. Expanding the framework to Haryana, which contributes a secondary but non-negligible fire signal, would improve model completeness. Finally, coupling the early-warning signal to a nowcasting ML model trained on CPCB rolling-mean $PM_{2.5}$ and FIRMS fire counts could provide probabilistic 3-day forecasts with calibrated uncertainty bounds, supporting more robust public-health decision-making.

7. Conclusions

This study has presented a systematic daily time-series analysis of the relationship between satellite-detected crop-residue burning in Punjab and $PM_{2.5}$ concentrations in Delhi-NCR over nine consecutive harvest seasons (2015–2024). The principal findings are as follows:

1. A statistically significant positive association between daily Punjab fire counts and Delhi-NCR $PM_{2.5}$ was established at all lag windows (0–72 hours), with the strongest correlation observed at Lag-0 ($r = 0.289$, $p < 0.001$) and monotonic decline through Lag-3 ($r = 0.235$), consistent with rapid atmospheric transport followed by multi-day boundary-layer accumulation.
2. Harvest-season $PM_{2.5}$ in Delhi-NCR averaged 150.5 µg/m³ across the study period, with 82.1% of days classified as 'Poor' and 10.3% as 'Severe' under the national AQI framework — rates that are 2–5 times higher than non-harvest-season benchmarks.
3. A dose–response gradient was documented between fire intensity and $PM_{2.5}$ outcomes, with high-intensity fire days producing significantly higher median $PM_{2.5}$ values than medium- or low-intensity days (Kruskal–Wallis $H = 18.4$, $p < 0.001$).
4. Annual Punjab fire counts declined by ~75% from the 2021 peak (~22,000 events) to the 2024 minimum (~5,500 events), reflecting the cumulative effect of policy interventions, yet mean harvest-season $PM_{2.5}$ decreased by only ~19% over the same period, indicating important meteorological and local-emission confounding.
5. Exponential decay modelling of the lag-correlation structure yielded a signal half-life of ~9.2 days, supporting the design of a three-tier early-warning framework in which high fire-count days trigger public health alerts within 0–1 days and extended $PM_{2.5}$ episode warnings at Days +3 to +7.

Collectively, these findings establish that satellite-derived fire counts over Punjab constitute a simple, freely available, and operationally deployable leading indicator of $PM_{2.5}$ in Delhi-NCR. While the correlations documented here are not large in absolute terms — reflecting the important roles of meteorology, local emissions, and secondary aerosol chemistry that lie beyond the scope of a bivariate analysis — they are robust, consistent across a decade of observations, and sufficient to anchor an early-warning framework that could provide public-health authorities with 24-to-72-hour advance notice of severe pollution episodes. Pursuing the meteorologically conditioned extensions outlined in Section 6.3 will be essential to translating this framework into a high-confidence operational product.

Acknowledgements

The authors gratefully acknowledge the Central Pollution Control Board of India for providing open-access ambient air-quality monitoring data, and NASA Earthdata (FIRMS) for maintaining freely available MODIS and VIIRS active-fire datasets. The authors declare no conflicts of interest. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

[1] Guttikunda, S. K., & Calori, G. (2013). A GIS-based emissions inventory at 1 km × 1 km spatial resolution for air pollution analysis in Delhi, India. Atmospheric Environment, 67, 101–111.

[2] Guttikunda, S. K., Goel, R., & Pant, P. (2014). Nature of air pollution, emission sources, and management in the Indian cities. Atmospheric Environment, 95, 501–510.

[3] Gurjar, B. R., Ravindra, K., & Nagpure, A. S. (2016). Air pollution trends over Indian megacities and their local-to-global implications. Atmospheric Environment, 142, 475–495.

[4] Balakrishnan, K., et al. (2019). The impact of air pollution on deaths, disease burden, and life expectancy across India. The Lancet Planetary Health, 3(1), e26–e39.

[5] Chowdhury, S., et al. (2019). Indian annual ambient air quality standard is achievable by completely mitigating emissions from household sources. PNAS, 116(22), 10711–10716.

[6] Pandey, A., et al. (2021). Health and economic impact of air pollution in the states of India. The Lancet Planetary Health, 5(1), e25–e38.

[7] Central Pollution Control Board. National Air Quality Index. Ministry of Environment, Forest and Climate Change, Government of India.

[8] Central Pollution Control Board. Central Control Room for Air Quality Management: CAAQM Data. Government of India.

[9] Ojha, N., et al. (2020). On the widespread enhancement in fine particulate matter across the Indo-Gangetic Plain towards winter. Scientific Reports, 10, 5862.

[10] Singh, N., et al. (2017). Fine particulates over South Asia: Review and meta-analysis of PM2.5 source apportionment. Environmental Pollution, 223, 121–136.

[11] Beig, G., et al. (2020). Objective evaluation of stubble emission of North India and quantifying its impact on air quality of Delhi. Science of the Total Environment, 709, 136126.

[12] Singh, R. P., & Kaskaoutis, D. G. (2014). Crop residue burning: A threat to South Asian air quality. Eos, 95(37), 333–334.

[13] Kaskaoutis, D. G., et al. (2014). Effects of crop residue burning on aerosol properties and long-range transport over northern India. JGR Atmospheres, 119, 5424–5444.

[14] Jain, N., Bhatia, A., & Pathak, H. (2014). Emission of air pollutants from crop residue burning in India. Aerosol and Air Quality Research, 14, 422–430.

[17] NASA FIRMS. Fire Information for Resource Management System: MODIS and VIIRS Active Fire Data. NASA Earthdata.

[18] Giglio, L., Schroeder, W., & Justice, C. O. (2016). The collection 6 MODIS active fire detection algorithm and fire products. Remote Sensing of Environment, 178, 31–41.

[22] Cusworth, D. H., et al. (2020). How much does large-scale crop residue burning affect the air quality in Delhi? Environmental Science & Technology, 54, 4790–4799.

[24] Lan, R., et al. (2022). Air quality impacts of crop residue burning in India and mitigation alternatives. Nature Communications, 13, 6537.

[25] Goyal, P., et al. (2025). Critical review of air pollution contribution in Delhi due to paddy stubble burning in Punjab and Haryana. Atmospheric Environment.

[26] Tiwari, S., et al. (2013). Diurnal and seasonal variations of black carbon and PM2.5 over New Delhi, India. Atmospheric Research, 125–126, 50–62.

[27] Mues, A., et al. (2018). Sensitivity of air pollution simulations with WRF-Chem to emissions and meteorological input data. Atmospheric Environment, 176, 302–316.

[28] Kumar, R., et al. (2021). The role of meteorology in modulating PM2.5 over Delhi during winter pollution episodes. Atmospheric Environment.

[29] Mohan, M., Bhati, S., & Rao, A. (2011). Air pollution modelling over Delhi: Impact of meteorology and boundary layer height. Atmospheric Environment, 45, 6791–6801.

[32] Mhawish, A., et al. (2022). Observational evidence of elevated smoke layers during crop residue burning season over Delhi. Remote Sensing of Environment.

[33] Biswal, A., et al. (2025). Emission time and amount of crop residue burning play critical role on PM2.5 variability during October–November in northwestern India. Environmental Science: Atmospheres.

[35] Saharan, U. S., et al. (2024). Hotspot-driven air pollution during crop residue burning season in the Indo-Gangetic Plain, India. Environmental Pollution.

[37] Grange, S. K., et al. (2018). Random forest meteorological normalisation models for Swiss PM10 trend analysis. ACP, 18, 6223–6239.

[43] Kumar, R. P., et al. (2024). Machine learning-based prediction of hazardous fine PM2.5 concentrations: A case study of Delhi, India. Discover Applied Sciences.

[44] Makhdoomi, A., et al. (2025). PM2.5 concentration prediction using machine learning models. Scientific Reports.

[45] Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences (2nd ed.). Lawrence Erlbaum Associates.

[46] Kerby, D. S. (2014). The simple difference formula: An approach to teaching nonparametric correlation. Comprehensive Psychology, 3, 1–9.

[47] Cleveland, R. B., et al. (1990). STL: A seasonal-trend decomposition procedure based on loess. Journal of Official Statistics, 6(1), 3–73.

[48] Gasparrini, A. (2011). Distributed lag linear and non-linear models in R: The package dlnm. Journal of Statistical Software, 43(8), 1–20.