Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

How does urban form shape spatial and demographic inequalities in compound heat–ozone exposure? — Evidence from Guangzhou and Shenzhen, China

Lin Zhou
Beijing Normal University
Ran Chen
Northwest A&F University
Shidong Ge
Henan Agricultural University
Chenyu Fang
11112026043@bnu.edu.cn

Beijing Normal University

Article

Keywords: Compound heat–ozone pollution, Spatial and demographic inequalities, Explainable machine learning, Megacity comparison

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11441053/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

How does urban form shape spatial and demographic inequalities in compound heat–ozone exposure? — Evidence from Guangzhou and Shenzhen, China

Lin Zhou (First Author)[1,2]

1. School of Government, Beijing Normal University, Beijing 100875, China
2. School of Public Policy, University of Massachusetts Amherst, Amherst, MA 01003, USA

E-mail: linzhou@umass.edu
ORCID: 0000-0002-2812-1096

Ran Chen (Second Author)
College of Landscape Architecture & Arts, Northwest A&F University, Yangling, 712100, China.
E-mail: chenran@nwafu.edu.cn
ORCID: 0009-0003-1897-3396

Shidong Ge (Third Author)
College of Landscape Architecture and Art, Henan Agricultural University, Zhengzhou 450002, China.
Email:shidongge@henau.edu.cn
ORCID: 0000-0003-0828-7379

Chenyu Fang (Corresponding Author)
School of Government, Beijing Normal University, Beijing 100875, China
E-mail: 11112026043@bnu.edu.cn
ORCID: 0000-0002-5280-847X

Abstract

Compound heatwave and ozone pollution (CHOP) events have emerged as a critical environmental health threat in rapidly urbanizing regions, yet how their fine-scale spatial heterogeneity differs between megacities sharing a common climatic background has received limited attention. Taking Guangzhou and Shenzhen, two megacities of China's Greater Bay Area, this study maps CHOP exposure on a unified 100 m grid, compares its spatial pattern and population burden between the two cities, and identifies the urban-form factors that explain their divergence with SHapley Additive exPlanations (SHAP) interpretation. The two cities show distinct exposure intensity, spatial distribution, and inequality patterns. Guangzhou has a higher mean exposure of 17.7 days, with a continuous high-exposure belt extending from the traditional core to the southern built-up districts. Shenzhen has a mean exposure of 11.8 days, with a localized central-western cluster against a low-exposure eastern ecological belt. Spatial inequality is also higher in Guangzhou (Gini 0.30 versus 0.21). The population burden differs in both scale and demographic composition. About 90.3% of Guangzhou's residents experience more than 15 CHOP days per year, compared with only 17.1% in Shenzhen, and high-exposure areas in Shenzhen also over-represent lower-education and lower-income residents, along with households at the pregnancy stage. The dominant influencing factors also diverge. Exposure in Guangzhou is shaped mainly by horizontal development intensity, including population density and road density, while exposure in Shenzhen is more closely tied to three-dimensional building morphology, including building height and coverage. These results indicate that compound exposure in dense megacities is shaped by city-specific urban form, supporting spatially differentiated governance.

Keywords

Compound heat–ozone pollution; Spatial and demographic inequalities; Explainable machine learning; Megacity comparison

1. Introduction

Global warming has substantially altered the pattern of extreme heat events, with heatwaves becoming more frequent, more intense, and more spatially extensive (Domeisen et al., 2023; Ghanbari et al., 2023). Rapid urbanization further compounds these pressures by reshaping surface properties, modifying built-environment morphology, and concentrating populations in vulnerable settings. Recent estimates indicate that heat-related deaths accounted for more than 11% of all disaster-related fatalities globally between 2000 and 2019 (Cheval et al., 2024; X. Zhou et al., 2025) . Beyond direct thermal stress, elevated temperatures also accelerate atmospheric photochemical reactions and promote the formation of ground-level ozone (Liang et al., 2017), which in turn can intensify heatwave conditions through long-wave radiative trapping (Su et al., 2025). As a result, heatwaves and ozone pollution frequently co-occur as compound heat–ozone pollution (CHOP) events (Xu et al., 2026) . Compared with single-hazard exposures, CHOP events impose substantially greater environmental stress, exhibit more complex spatial signatures, and translate into stronger public-health burdens (Zhou et al., 2025), and have become one of the most pressing concerns at the intersection of urban systems research, environmental science, and public-health risk assessment (Schnell and Prather, 2017).

The synergistic relationship between heat and ozone is now well established: high temperatures accelerate photochemistry and reshape boundary-layer dynamics, while dense built environments and concentrated emissions amplify local accumulation. This coupling translates into substantial health consequences. Compared with exposure to heat or ozone alone, CHOP is associated with substantially higher cardiovascular mortality risks, with reported effect sizes 1.6 to 3.2 times larger (J. Su et al., 2025); the combined action of heat and air pollution further elevates non-accidental and respiratory mortality, and the attributable fraction of CHOP-related deaths in the Middle East and North Africa rose from 0.99% in 2000 to 2.01% in 2021 (Xu et al., 2024). Yet despite this clearly documented public-health relevance, the spatial distribution of CHOP exposure within cities and the populations bearing the resulting burden have received much less attention, leaving a critical gap for translating compound-risk evidence into actionable urban interventions. This gap is a matter of analytical resolution. Intra-urban CHOP analysis has been constrained by the spatial resolution of available datasets: regional air-quality grids and population datasets have historically been distributed at kilometer-to-county scales, which are too coarse to resolve neighborhood-level compound exposure or the urban-form factors that shape it.

What this coarse resolution masks is substantial within-city variation in both hazards. Heat exposure varies substantially with functional zoning, built-environment typology, and population settlement patterns, jointly shaped by urban form, land use, road networks, building density, green-space configuration, and the spatial distribution of residents (Namaiti et al., 2026). Ozone, although often treated as a regional-scale phenomenon, also exhibits strong intra-urban variability driven by local emissions, ventilation conditions, precursor sensitivity, and functional organization (Huang et al., 2023). Most CHOP research, however, continues to operate at macro scales (Ban et al., 2022), focusing on event frequency, expanding population exposure, and aggregated health attribution, with little attention to whether heat and ozone form a coherent intra-urban geography. Even where finer-scale CHOP analyses are attempted, the built-environment dimension is typically captured through aggregate or two-dimensional measures, and three-dimensional building morphology has rarely been incorporated, despite its role in shaping both surface heat budgets and pollutant accumulation in dense vertical environments. As a result, the fine-scale geography of compound risk in dense megacities, and the specific urban-form factors that produce it, remain poorly understood.

Beyond resolution, intra-urban CHOP patterns and their underlying drivers may differ even between cities sharing the same climate. Megacities can develop along very different physical templates despite a comparable climate, varying in built-up continuity, three-dimensional form, ecological-space configuration, and functional organization. Guangzhou and Shenzhen, the two core megacities of China's Greater Bay Area, offer an informative case for examining this. Guangzhou is a long-established city dominated by horizontal sprawl, with basin-like topography that constrains pollutant dispersion. Shenzhen is younger, more compact, and vertically intensified, with a substantial preserved ecological belt in its eastern districts. Whether such urban-form differences translate into contrasting fine-scale CHOP geographies and contrasting sets of dominant drivers is an open empirical question. A direct comparison of the two cities can therefore help separate the potential roles in shaping fine-scale CHOP exposure.

To address these gaps, this study compares fine-scale CHOP exposure between Guangzhou and Shenzhen using a unified 100 m grid framework, and applies explainable machine learning to identify city-specific influencing factors. Three research questions are addressed: (1) how the intra-urban spatial distribution of CHOP exposure compares between the two megacities; (2) how the population burden of CHOP exposure differs between them, including the demographic profile of

high- and low-exposure neighborhoods; and (3) what socio-demographic and built-environment factors are associated with CHOP exposure within and between the two megacities.

This study contributes to the literature in three main ways. First, it develops a fine-scale (100 m) analytical framework that integrates socio-demographic attributes with both two-dimensional urban form and three-dimensional building morphology, enabling consistent intra-urban comparison of CHOP exposure across cities and providing a more complete characterization of urban-form influences than prior CHOP research, which has typically relied on aggregate or two-dimensional measures. Second, by directly comparing two megacities within the same climate region, the study shows empirically that fine-scale CHOP exposure patterns and population burdens diverge between cities even when regional climate is held constant, providing evidence that compound risk cannot be predicted from climate alone and must be examined at the intra-urban scale. Third, by applying explainable machine learning to identify the urban-form and socio-demographic factors associated with CHOP exposure in each city, the study reveals city-specific influencing factors of CHOP exposure and provides an empirical basis for spatially targeted, differentiated CHOP governance in dense megacities of the Greater Bay Area and beyond.

The remainder of the paper is organized as follows. Section 2 introduces the study area, datasets, and methods. Section 3 presents the spatial differentiation of CHOP exposure, the population exposure burden, and the dominant influencing factors in Guangzhou and Shenzhen. Section 4 discusses the findings and their policy implications, and Section 5 concludes.

2. Data and methods

2.1. Methodological framework

An analytical framework was developed to characterize and compare CHOP exposure between Guangzhou and Shenzhen at fine spatial resolution (Fig. 1). The framework comprises three steps. First, daily LST and MDA8 ozone (environmental exposure), together with socio-demographic and built-environment datasets, the latter spanning both two-dimensional urban form and three-dimensional building morphology, were harmonized into a unified 100 m gridded database. Second, CHOP events were identified at the grid scale and aggregated to annual exposure frequency, after which the spatial signatures of compound exposure in the two cities were compared through spatial maps, latitudinal–longitudinal profiles, descriptive statistics, inequality indices, and population-burden estimates. Third, an ensemble machine-learning model coupled

with SHAP interpretation was used to identify the socio-demographic and built-environment factors associated with the divergence in CHOP exposure.

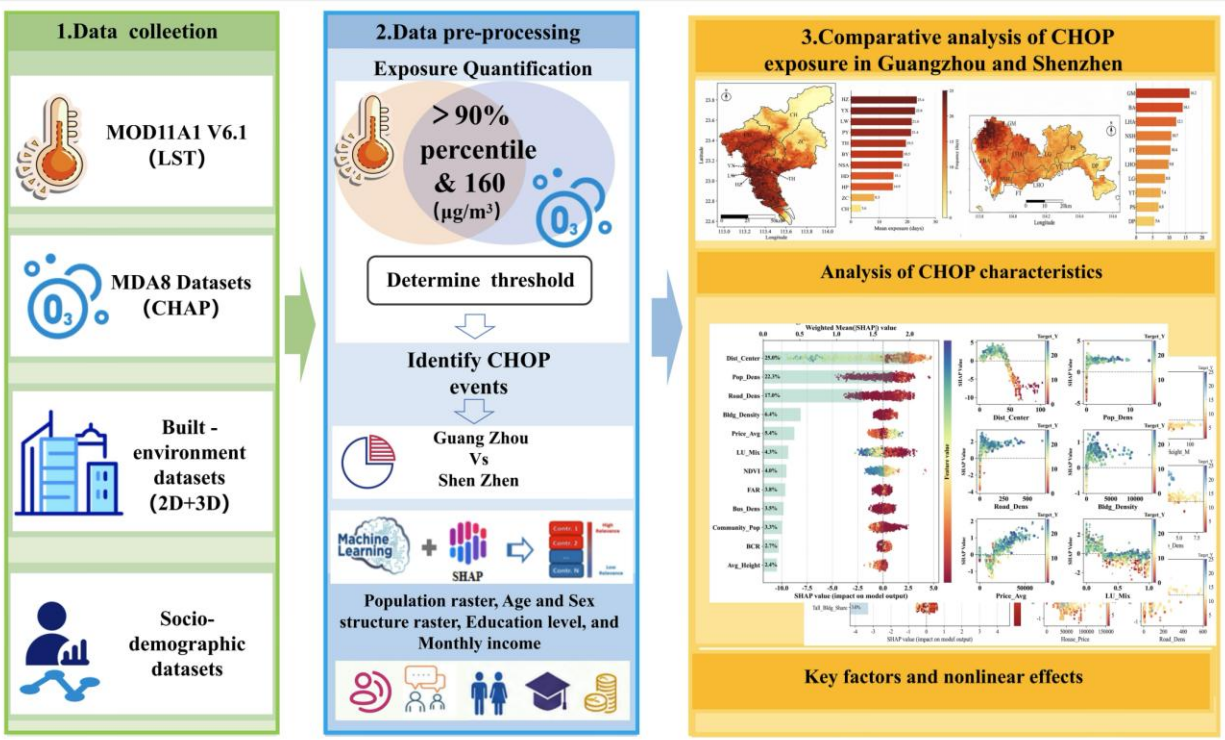


Fig. 1. The methodological framework

2.2. Study area

Guangzhou (GZ) and Shenzhen (SZ) are the two largest megacities of China's Greater Bay Area and lie within a single subtropical-monsoon climate domain, where intense summer solar radiation, high humidity, and frequent atmospheric stagnation render surface ozone highly sensitive to extreme heat (Fig. 2). The two cities share this regional environmental background yet differ in spatial scale, urban form, and ecological context, providing a natural setting for fine-scale comparison. Guangzhou occupies approximately 7,434 km² in the northern Pearl River Delta and is organized around eleven administrative districts, including Liwan (LW), Yuexiu (YX), Tianhe (TH), Haizhu (HZ), Panyu (PY), Nansha (NSA), Zengcheng (ZC), and Conghua (CH); its terrain transitions from the alluvial plain in the south to low hills in the north and northeast, forming a basin-like configuration that limits horizontal pollutant dispersion and concentrates summertime ozone in the urban core. Shenzhen, in contrast, spans only about 1,997 km^2on the eastern bank of the Pearl River Estuary but supports a permanent population of roughly 17.99 million, comparable

to Guangzhou's, resulting in a much more compact urban fabric; its territory rises from the densely built northwestern plain toward a mountain–forest–coast complex in the southeast that retains substantial ecological space. The contrast between Guangzhou's contiguous, horizontally extensive urbanization and Shenzhen's compact, vertically intensified development against a preserved eastern ecological belt forms the basis for the analyses below.

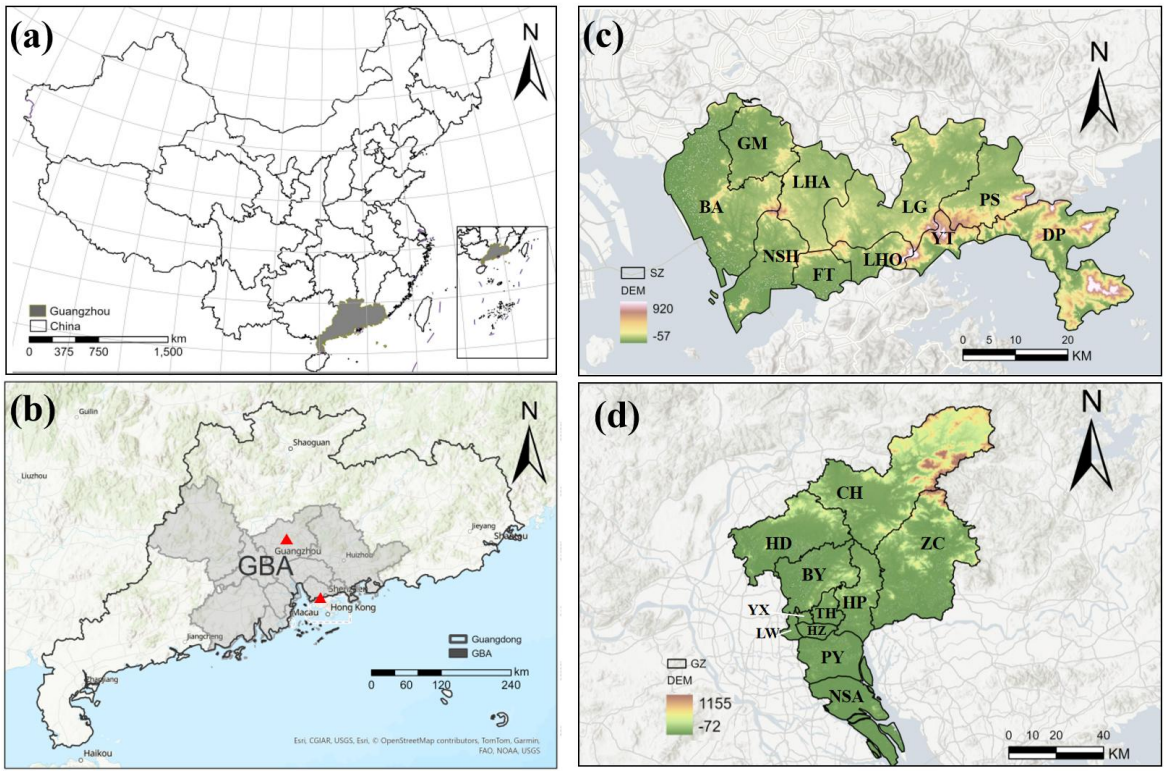


Fig. 2. Study area: (a) Location of the Greater Bay Area in China; (b) location of the GBA, Guangzhou, and Shenzhen in Guangdong Province; (c) DEM and district boundaries of Shenzhen; (d) DEM and district boundaries of Guangzhou.

2.3. Datasets

This study integrates multi-source datasets that span three domains (Table 1): socio-demographic attributes, built environment, and environmental exposure. The first two domains serve as the independent variables, while environmental exposure (daily LST and MDA8 ozone) is used to construct the dependent variable, the annual CHOP frequency. All datasets[1] were harmonized through unified projection, spatial alignment, and resampling, and were registered to

[1] Dataset and code could be found in https://doi.org/10.6084/m9.figshare.32396964

a 100 m population raster as the basic analytical unit. The predictor set includes 46 input features spanning socio-demographic attributes, two-dimensional urban form, and three-dimensional building morphology.

Table 1. Summary of variables and data sources.

Domain	Variable	Description	Abbreviation	Year	Source	Unit
Socio-demographic attributes	*Population density*	Total population per unit area	Pop_Dens	2019	Baidu Huiyan	persons/km²
	Gender	Male and female population counts	Gender_Male, Gender_Female			
	Age	Seven age groups from under-18 to 65+	Age_Under18, Age_18_24, Age_25_34, Age_35_44, Age_45_54, Age_55_64, Age_65Plus			
	Education	Low-, middle-, and high-education groups	Edu_Low, Edu_Mid, Edu_High			
	Income	Five monthly-income groups from low to very high	Inc_Low, Inc_LMid, Inc_Mid, Inc_High, Inc_VeryHigh			
	Vehicle ownership	Car-ownership status of residents	Asset_Car, Asset_NoCar			
	Life stage	Six population life-stage groups, from primary-school students to parenting adults	Stage_JrStudent, Stage_SrStudent, Stage_UnivStudent, Stage_GradStudent, Stage_Pregnant, Stage_Parenting			
	Family	Seven household stages defined by the youngest child, from pregnancy to high school	Fam_Pregnant, Fam_Baby_0_1, Fam_Toddler_1_3, Fam_Child_3_6, Fam_Primary, Fam_JrHigh, Fam_SrHigh			
	Housing price	Average housing price per unit area	Price_Avg	2020	Housing listings	RMB/m²

Domain	Variable	Description	Abbreviation	Year	Source	Unit
Built environment	*Road network*	Road-intersection density	Road_Dens	2020	OpenStreetMap	intersections/km²
	Locational accessibility	Distance to urban center	Dist_Center	2020	OpenStreetMap	km
	Land-use composition	Land-use diversity (entropy index)	LU_Mix	2020	Land-use dataset	dimensionless
	Public transport accessibility	Metro-station and bus-stop density	Metro_Dens, Bus_Dens	2020	OpenStreetMap	stations/km², stops/km²
	Ecological condition	Vegetation greenness (satellite-derived)	NDVI	2020	Landsat-8	dimensionless
	Building height	Mean and 90th-percentile building height	Avg_Height, P90_Height	2019	Building footprint and height dataset	m
	Building intensity	Building coverage ratio and floor-area ratio	BCR, FAR			dimensionless
	Building density	Number of buildings per unit area	Bldg_Density			buildings/km²
	Tall-building share	Proportion of tall buildings within each grid cell	Tall_Bldg_Share			%
Environmental exposure	*Thermal environment*	Grid-specific heat threshold (90th-percentile daily maximum LST in warm seasons)	LST_P90	2013–2022	MOD11A1 V6.1	°C
	Air pollution	Daily maximum 8-hour average ozone concentration	O3_MDA8	2019 (daily)	CHAP dataset	μg/m³

Notes: Cells in the "Data dimension" column are merged across all modules belonging to the same dimension. Variable codes are written in monospaced type to distinguish them from descriptive text. For socio-demographic modules sharing the same temporal coverage, source and unit, the values are nevertheless reported on each row to keep the table readable when extracted in isolation.

2.3.1. Socio-demographic and housing data

Population data were obtained from the Baidu Huiyan platform, which provides 100 m gridded profiles for 2019. The dataset records grid-level population density, gender composition, age structure, educational attainment, monthly income brackets, vehicle ownership, life-stage groups, and family composition. Socioeconomic information was captured by average housing price, constructed at the same grid scale from second-hand housing listings on Lianjia and Fang.com (2020), which serves as a proxy for residential socioeconomic status, complementing the income and education indicators.

To avoid confounding from differences in absolute population size across grid cells, all demographic attributes were converted to within-grid proportions before being matched with CHOP exposure. The resulting variables play two roles in the subsequent analysis: they characterize the demographic profile of high- and low-exposure neighborhoods, and they serve as input features in the predictive model. Continuous variables were checked for outliers and missing values prior to modeling.

2.3.2. Built environment

Two-dimensional urban-form variables, the first subdimension of the built environment, describe horizontal development intensity, locational accessibility, land-use composition, public-transport accessibility, and vegetation cover. Road intersection density (Road_Dens), obtained from OpenStreetMap (2020), serves as a proxy for network density and horizontal development intensity. Distance to the urban center (Dist_Center), also derived from OpenStreetMap, captures the position of each grid cell relative to the city's main core. Land-use mix (LU_Mix), computed from the land-use dataset and Landsat-8 imagery, was quantified by an entropy index that captures the diversity of co-located functions. Metro station density (Metro_Dens) and bus stop density (Bus_Dens), drawn from OpenStreetMap, characterize transit service levels and node-level activity intensity. The Normalized Difference Vegetation Index (NDVI) was computed from the near-infrared and red bands of Landsat-8 to represent vegetation greenness and surface ecological condition; higher values indicate denser vegetation, and lower values typically correspond to highly developed or sparsely vegetated areas.

Three-dimensional building-morphology variables, the second subdimension of the built environment, describe the vertical structure of the urban fabric, which is particularly relevant in dense megacities where vertical form modifies surface heat budgets, ventilation, and pollutant

accumulation. From a building footprint and height database for 2020–2021, six indicators were derived: mean building height (Avg_Height), the 90th-percentile of building height (P90_Height), building coverage ratio (BCR), floor-area ratio (FAR), building density (Bldg_Density), and the share of tall buildings (Tall_Bldg_Share).

2.3.3. Environmental exposure data

Two daily environmental datasets are used to identify CHOP events: surface heat and ground-level ozone. Daily maximum land surface temperature was sourced from the MODIS MOD11A1 V6.1 product (1 km resolution, 2013–2022) and used to derive the grid-specific 90th-percentile heat threshold (LST_P90). Daily MDA8 ozone (O3_MDA8) for 2019 was obtained from the China High Air Pollutants (CHAP) dataset (1 km resolution), and combined with the 2019 daily LST to identify CHOP events. The CHAP product is generated through a space–time extreme randomized trees (STET) ensemble that combines surface solar radiation, air temperature, ground monitoring, atmospheric reanalysis, and emission inventories, and has demonstrated high accuracy in capturing the spatiotemporal variability of surface ozone over China (Wei et al., 2022).

2.4. Methods

2.4.1. Identification of CHOP events

A CHOP event is defined at the grid scale as a day on which the same grid cell simultaneously experiences anomalously high heat and elevated ozone. Specifically, heat anomalies are determined using a relative, locally referenced threshold: for each grid cell, the 90th percentile of daily maximum LST during the warm season (May–October) of 2013–2022 is computed and adopted as the local heat threshold. Ozone pollution, by contrast, is identified using an absolute, regulation-anchored threshold: a grid cell is flagged as polluted on a given day when its MDA8 ozone concentration meets or exceeds 160 μg/m³, corresponding to the Grade II limit of China's Ambient Air Quality Standard (GB 3095-2012). When both criteria are satisfied simultaneously on a given day in 2019, the day is classified as a CHOP event. The daily flags are then summed over all days of 2019 to obtain the annual CHOP frequency at each grid cell, which serves as the dependent variable for the modeling analyses:

$$CHOP_i = \sum_{d=1}^{D} I\left(LST_{i,d} > P90_LST_i\right) \times I\left(O_{3\,i,d} \geq 160\right)$$

where $CHOP_i$ denotes the annual CHOP frequency at grid i, $LST_{i,d}$ is the daily maximum LST at grid i on day d, $P90_LST_i$ is the warm-season 90th-percentile threshold derived from the 2013–2022 LST series, $O_{3,i,d}$ is the daily MDA8 ozone concentration in the study year, D is the total number of days in the study year, and $I(\cdot)$ is the indicator function that equals 1 when its condition is satisfied and 0 otherwise.

2.4.2. Spatial inequality and population burden

To quantify the spatial differentiation of CHOP exposure within each city, three complementary inequality indicators are used: the coefficient of variation (CV), the Gini coefficient, and the Theil index. CV captures the relative dispersion of CHOP frequency across grid cells; the Gini coefficient summarizes the overall level of inequality, with higher values indicating greater concentration of exposure in fewer cells; and the Theil index captures unevenness with stronger sensitivity to the upper tail of the distribution. Larger values across these indicators consistently signal stronger intra-urban differentiation. The CV and Gini coefficient are defined as:

$$CV = \frac{\sigma}{\mu}, \qquad G = \frac{\sum_{i=1}^{n}\sum_{j=1}^{n}|x_i - x_j|}{2n^2\mu}$$

where x_i and x_j denote the CHOP frequencies at spatial units i and j, μ is the mean exposure frequency, σ is the standard deviation, and n is the number of spatial units. To compare the CHOP frequency distributions of Guangzhou and Shenzhen, the Mann–Whitney U test is applied, as it does not require normality, and the effect size $r = Z/\sqrt{N}$ is reported to indicate the magnitude of the difference.

The population burden of CHOP exposure is estimated by summing the gridded population at successive exposure thresholds (≥ 15, ≥ 20, and ≥ 25 days per year), yielding both the absolute number and the percentage of residents falling within each band. High- and low-exposure neighborhoods within each city are defined by exposure quartiles: grid cells in the top quartile are treated as high-exposure neighborhoods, and those in the bottom quartile as low-exposure neighborhoods.

2.4.3. Model selection and spatial cross-validation

Four tree-based ensemble models, namely XGBoost (Chen and Guestrin, 2016), CatBoost

(Prokhorenkova et al., 2018), LightGBM (Ke et al., 2017), and Random Forest (Breiman, 2001), were compared to identify the model best suited for explaining the spatial variation of CHOP exposure in each city. XGBoost is an extreme gradient-boosting framework that combines regularization with second-order optimization and is widely adopted in geospatial regression; CatBoost is a gradient-boosting decision tree that natively handles categorical variables and is robust to overfitting on small samples; LightGBM grows trees in a leaf-wise rather than level-wise manner, achieving high computational efficiency on high-dimensional data; and Random Forest is a bagging-based ensemble that aggregates predictions from many randomized decision trees and provides a robust nonparametric baseline. Predictive performance was evaluated using the coefficient of determination (R^2), mean absolute error (MAE), and root mean square error (RMSE), with hyperparameters tuned via ten-fold cross-validation and grid search on the training set.

Final predictive performance was evaluated under a spatial-block cross-validation scheme rather than a conventional random split, to account for the spatial autocorrelation inherent in environmental-exposure data. The study area was partitioned into non-overlapping 2 km × 2 km regular blocks (block2000), and each block was assigned in its entirety to either the training or the test set. This design suppresses leakage from spatially adjacent cells, which would otherwise share local environmental conditions across the train–test boundary and yield optimistic generalization estimates. The best-performing learner under this scheme was retained for the SHAP-based interpretation.

2.4.4. SHAP-based interpretation of influencing factors

The best-performing model identified in Section 2.4.3 was coupled with the SHapley Additive exPlanations (SHAP) framework (Lundberg and Lee, 2017) to identify the factors associated with CHOP exposure in each city. SHAP is well suited to this purpose because it accommodates the multicollinearity, nonlinearity, and threshold effects typical of urban built-environment data, where conventional regression methods often fall short. SHAP values, derived from cooperative game theory, decompose each prediction into the marginal contributions of individual features. A positive SHAP value indicates that the feature increases the predicted CHOP frequency at that grid cell, while a negative value indicates the opposite. At the global level, feature importance is summarized using a population-weighted mean of |SHAP|, in which each grid cell's SHAP value is weighted by its population so that areas with more residents contribute more

strongly to the feature ranking. SHAP dependence plots are then used to visualize the direction, nonlinearity, and potential thresholds of feature effects. Feature rankings, SHAP distributions, and key dependence curves are then compared between the Guangzhou and Shenzhen models. All modeling and interpretation were conducted in Python 3.9.

2.4.5. Robustness checks across spatial scales

A central methodological consideration is that the CHOP exposure surface is observed at the 1,000 m support, while the predictors resolve to 100 m. The 100 m models therefore do not learn an independently observed 100 m exposure field; instead, they use fine-grained covariates to explain a response that remains anchored to the 1,000 m support, since each 100 m cell inherits the CHOP value of its parent 1,000 m grid. This support mismatch, compounded by the strong spatial autocorrelation of environmental-exposure data, would render conventional random splits highly leaky, as both nearest-neighbor cells and cells nested within the same parent unit could appear in training and test sets simultaneously, yielding overly optimistic performance estimates. To address this challenge, spatially blocked cross-validation strategies were designed in tandem with the modeling structure described below.

To make full use of the multi-resolution data while remaining faithful to the 1,000 m support of the response, two complementary modeling pipelines were estimated and compared. The 100 m pipeline constitutes the primary analytical framework and uses the full set of fine-grained socio-demographic and built-environment predictors to explain the spatial variation of CHOP exposure. The 1,000 m pipeline, in which both predictors and response are defined on the 1,000 m unit, serves as a scale-aligned reference whose role is not to replace the 100 m analysis but to verify, when the analytical scale is fully aligned with the response support, whether the dominant patterns identified at 100 m persist, thereby providing complementary, robustness-oriented evidence in support of the 100 m results.

The two pipelines were paired with pipeline-specific spatial validation designs that respond to their distinct sources of leakage. The 100 m pipeline was evaluated under two progressively stricter splits. The grouped 1,000 m split assigns all 100 m cells nested within the same parent 1,000 m grid to a single partition, so that no parent unit contributes samples to both training and test sets, directly eliminating the duplicated-label leakage produced by the 1,000 m support of the response. The 2,000 m spatial-block split goes further by partitioning the study area into non-

overlapping 2 km × 2 km regular blocks and assigning each block in its entirety to either the training or the test set, suppressing leakage from neighboring cells across the train–test boundary. The 1,000 m pipeline, in turn, was evaluated under both a conventional random split and the same 2,000 m spatial-block split, so that the stability of its results can be compared between the two patterns.

A parallel distinction was drawn at the evaluation stage. For the 100 m pipeline, two evaluation levels were reported. The row-level 100 m metric assesses model fit directly at the 100 m sample level and reflects how well the model captures fine-grained covariate patterns. The aggregated 1,000 m metric, by contrast, recombines the 100 m predictions within each parent grid and compares them against the observed 1,000 m exposure value; this metric is the more decisive of the two, as it aligns the assessment scale with the true spatial support of the response. The 1,000 m pipeline, by construction, is trained and evaluated only at the row-level 1,000 m. The best-performing learner in each pipeline was retained for the SHAP-based interpretation.

3. Results

3.1. Spatial distribution of CHOP exposure

The two megacities differ in both the overall level and the spatial pattern of CHOP exposure (Fig. 3). Guangzhou records a citywide mean annual frequency of 17.7 days, compared with 11.8 days in Shenzhen, with the standard deviation of district-level means also larger in Guangzhou (5.62 days versus 3.25 days). In Guangzhou (Fig. 3a), exposure is organized around a contiguous high-exposure belt centered on the traditional urban core of Yuexiu, Liwan, and Tianhe and extending southward across Haizhu, Panyu, and Nansha, with values attenuating gradually toward the peripheral mountainous and ecological districts of Conghua and Zengcheng in the north; the district-level ranking confirms a smoothly varying core-to-periphery gradient with limited fragmentation. Beyond the higher mean, the within-city distribution in Guangzhou is also more polarized, with a substantial share of grid cells deviating noticeably from the citywide average in both directions, so that very high and very low exposure values coexist within the same urban fabric. Shenzhen, by contrast, shows a more localized and fragmented structure (Fig. 3b): high values concentrate in the densely built central-western corridor, especially in Guangming and Bao'an, while substantial portions of the eastern districts (Yantian, Pingshan, and Dapeng) remain at relatively low exposure, owing to their preserved mountain-and-coast ecological setting. The

district-level ranking shows an east–west asymmetry rather than a smooth core-to-periphery decline. Despite the visible local concentration in the central-western corridor, the broad low-exposure ecological belt in Shenzhen's east keeps the city's overall distribution shape relatively flat compared with Guangzhou.

District-level statistics quantify these citywide patterns (Table 2). In Guangzhou, Haizhu records the highest mean exposure (23.34 ± 1.58 days, where ± denotes the within-district standard deviation across constituent grid cells), about 4.4 times the value in Conghua (5.29 ± 4.52 days). Yuexiu, Liwan, and Panyu also exceed 20 days (23.12 ± 1.79, 21.68 ± 1.48, and 21.38 ± 2.51 days), confirming that compound exposure in Guangzhou has spread well beyond the historical core into the southern built-up districts. Nansha records 17.95 days on average but with a markedly larger within-district standard deviation of 6.67 days, reflecting the coexistence of intensively developed industrial zones with coastal and peripheral pockets. In Shenzhen, the largest exposure burdens fall not on the historical urban core but on the rapidly expanding districts of Guangming (16.56 ± 3.34 days) and Bao’an (14.34 ± 4.17 days); Futian (10.47 ± 1.04 days) and Luohu (10.10 ± 1.76 days) record only moderate means, while Dapeng (6.46 ± 3.16 days) anchors the eastern low-exposure belt.

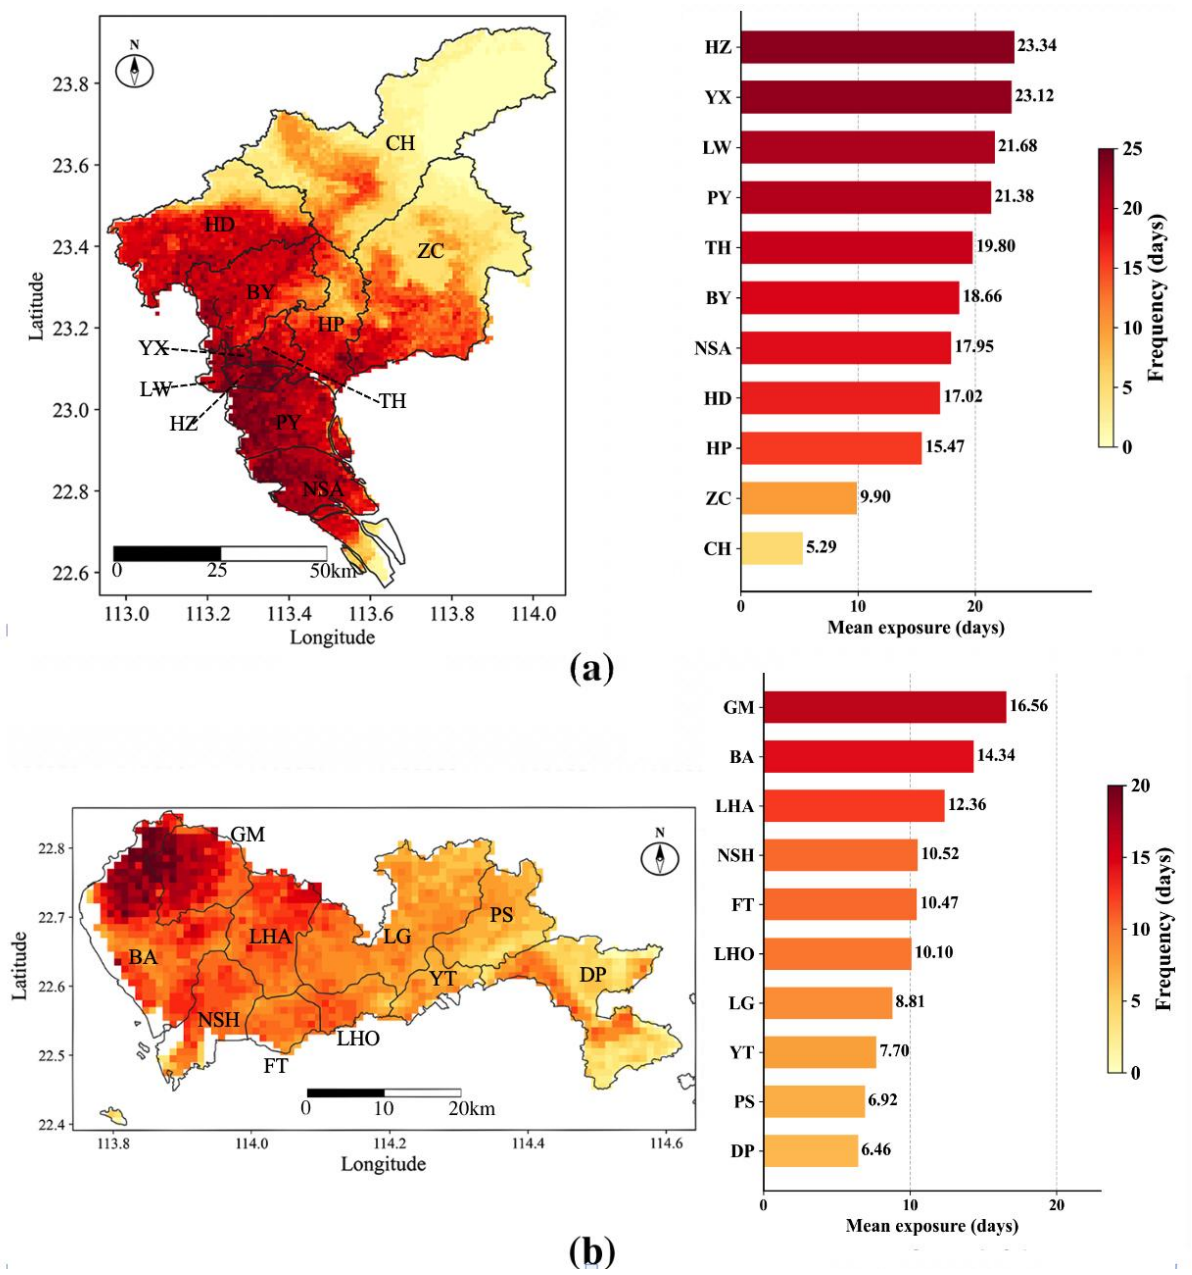


Fig. 3. Spatial distribution and district-level mean of CHOP exposure in Guangzhou and Shenzhen: (a) Guangzhou; (b) Shenzhen.

Table 2. District-level descriptive statistics of annual CHOP exposure days in Guangzhou and Shenzhen.

City	District	Abbreviation	Mean (days)	SD (days)
Guangzhou	Haizhu	HZ	23.34	1.58
	Yuexiu	YX	23.12	1.79
	Liwan	LW	21.68	1.48

	Panyu	PY	21.38	2.51
	Tianhe	TH	19.80	2.50
	Baiyun	BY	18.66	2.94
	Nansha	NSA	17.95	6.67
	Huadu	HD	17.02	4.53
	Huangpu	HP	15.47	4.39
	Zengcheng	ZC	9.90	5.45
	Conghua	CH	5.29	4.52
Shenzhen	Guangming	GM	16.56	3.34
	Bao'an	BA	14.34	4.17
	Longhua	LHA	12.36	1.82
	Nanshan	NSH	10.52	2.46
	Futian	FT	10.47	1.04
	Luohu	LHO	10.10	1.76
	Longgang	LG	8.81	1.54
	Yantian	YT	7.70	1.53
	Pingshan	PS	6.92	1.21
	Dapeng	DP	6.46	3.16

Beyond these mean and district-level differences, intra-urban inequality of CHOP exposure is also stronger in Guangzhou than in Shenzhen. The coefficient of variation (CV) is 0.53 in Guangzhou versus 0.39 in Shenzhen, indicating greater relative dispersion of grid-level exposure around the citywide mean. The Gini coefficient is 0.30 in Guangzhou versus 0.21 in Shenzhen, about 1.41 times higher, reflecting a more concentrated distribution in which a smaller share of cells accounts for a disproportionate share of total exposure days. The Theil index, which is more sensitive to the upper tail, is 0.13 in Guangzhou versus 0.07 in Shenzhen, about 1.81 times higher, suggesting that the highest-exposure cells in Guangzhou contribute disproportionately to overall inequality. The Mann–Whitney U test confirms that the two distributions differ at the $p < 0.001$ level, with an effect size r of 0.51, corresponding to a large difference.

3.2. Population exposure burden and demographic profile

The two cities differ in the share of residents exposed to elevated CHOP frequencies. Using 15 CHOP days per year as a moderate-to-high exposure threshold, approximately 15.14 million

residents of Guangzhou (about 90.3% of the city's population) live in such areas, compared with 2.62 million residents in Shenzhen (about 17.1%); the ratio between the two cities reaches 5.78 to 1. Raising the threshold to 20 days widens the gap further: 8.34 million residents of Guangzhou (49.7%) remain in high-exposure areas, compared with only 0.73 million (4.8%) in Shenzhen, a ratio of approximately 11.4 to 1. At the 25-day extreme, around 0.37 million residents of Guangzhou are still affected, whereas no Shenzhen grid cells reach this level. Compound exposure has therefore become a near-citywide condition for Guangzhou's residents but remains spatially confined in Shenzhen.

The demographic composition of high- and low-exposure neighborhoods is broadly similar in Guangzhou (Table 3). Across gender, education, income, and vehicle ownership, within-area shares in the top exposure quartile differ from those in the bottom quartile by less than 2 percentage points. The largest deviations occur in age structure: residents aged 45–54 are 3.6 percentage points more concentrated in high-exposure areas (13.4% versus 9.8%), while those aged 18–24 are 5.9 percentage points less common (17.3% versus 23.2%). A consistent life-stage shift accompanies this pattern, with university students under-represented in high-exposure areas ($\Delta = -3.3$ pp) and graduate-stage residents slightly over-represented ($\Delta = +1.2$ pp). In family composition, households at the pregnancy stage are also marginally more prevalent in high-exposure areas (12.0% versus 10.2%, $\Delta = +1.8$ pp), while other family categories shift by less than 1.5 percentage points. These patterns reflect a more settled, mature resident structure in Guangzhou's contiguous urban core, in contrast with the younger, more mobile profile of peripheral districts. The relatively small income- and education-related differences indicate that compound exposure in Guangzhou is broadly shared across socioeconomic strata rather than being concentrated on any single disadvantaged group.

In Shenzhen, by contrast, high-exposure areas display pronounced socioeconomic and household-composition gradients. The clearest gradient appears in educational attainment: the share of low-education residents is 7.2 percentage points higher in high-exposure areas (69.7% versus 62.5%), while the share with higher education is 4.6 percentage points lower (9.0% versus 13.6%). A parallel income gradient emerges, with lower-middle-income residents over-represented in high-exposure areas (35.4% versus 30.6%, $\Delta = +4.8$ pp) and higher- and very-high-income residents jointly under-represented by 7.3 percentage points. Age structure shifts toward younger working-age groups, with 25–34 year-olds making up 56.5% of high-exposure

populations versus 51.6% in low-exposure ones (Δ = +4.9 pp), and a small but consistent male bias is also visible (Δ = +3.0 pp). Family composition mirrors these signals: households at the pregnancy stage are notably more prevalent in high-exposure areas (26.4% versus 21.8%, Δ = +4.6 pp), and households with infants (3.3% versus 3.2%), toddlers (5.1% versus 4.9%), young children aged 3–6 (13.5% versus 12.9%), and primary-school students (61.8% versus 60.3%) are slightly over-represented, while households with junior- or senior-high-school students are somewhat under-represented (11.1% versus 12.6% for junior-high; 4.3% versus 5.2% for senior-high). Together, these patterns indicate that the spatial concentration of compound exposure in Shenzhen overlaps with the residential geography of younger workers, lower-income and lower-education residents, and households at the most health-sensitive family stages, including pregnancy and very early childhood.

Table 3. Within-area shares (%) of socio-demographic groups in high- and low-exposure neighborhoods of Guangzhou and Shenzhen.

Group	Indicator	Guangzhou			Shenzhen		
		High (%)	Low (%)	Δ (pp)	High (%)	Low (%)	Δ (pp)
Gender	Male	62.8	64.2	−1.4	68.2	65.2	+3.0
	Female	37.2	35.8	+1.4	31.8	34.8	−3.0
Age	Under 18	1.0	1.2	−0.2	0.6	0.8	−0.2
	18–24	17.3	23.2	−5.9	14.1	13.0	+1.2
	25–34	45.2	46.5	−1.3	56.5	51.6	+4.9
	35–44	19.2	17.7	+1.5	19.8	21.7	−1.9
	45–54	13.4	9.8	+3.6	8.0	11.1	−3.1
	55–64	3.0	1.3	+1.7	0.8	1.5	−0.7
	65+	0.9	0.3	+0.6	0.1	0.3	−0.2
Education	Low	66.2	67.3	−1.1	69.7	62.5	+7.2
	Middle	21.4	22.4	−1.0	21.3	23.9	−2.6
	High	12.4	10.3	+2.0	9.0	13.6	−4.6
Income	Low	8.2	8.8	−0.5	2.4	2.3	+0.1
	Lower-middle	32.7	34.5	−1.8	35.4	30.6	+4.8
	Middle	34.0	32.8	+1.2	32.9	30.7	+2.3

	High	17.4	17.5	−0.1	20.2	24.6	−4.4
	Very high	7.7	6.5	+1.3	9.0	11.9	−2.9
Vehicle ownership	Has car	8.2	7.8	+0.4	9.7	11.6	−1.9
	No car	91.8	92.2	−0.4	90.3	88.4	+1.9
Life stage	Junior-high student	4.2	3.2	+1.0	3.6	4.8	−1.2
	Senior-high student	21.4	22.1	−0.7	29.5	26.8	+2.7
	University student	57.4	60.7	−3.3	36.1	40.6	−4.5
	Graduate student	4.9	3.7	+1.2	4.2	5.8	−1.6
	Pregnancy	12.0	10.2	+1.8	26.4	21.8	+4.6
	Parenting	0.1	0.1	0.0	0.2	0.2	0.0
Family	Has pregnant member	0.8	0.9	−0.1	0.9	0.8	+0.1
	Baby (0–1 yr)	3.0	3.4	−0.4	3.3	3.2	+0.1
	Toddler (1–3 yr)	4.5	5.0	−0.5	5.1	4.9	+0.2
	Child (3–6 yr)	11.8	12.9	−1.1	13.5	12.9	+0.6
	Primary student	61.8	61.6	+0.3	61.8	60.3	+1.5
	Junior-high student	12.4	11.2	+1.2	11.1	12.6	−1.6
	Senior-high student	5.7	5.0	+0.7	4.3	5.2	−0.9

Note: Values are within-area shares (%) of the resident population. High-exposure and low-exposure areas refer to grid cells in the top and bottom exposure quartiles within each city, respectively. Δ (pp) is the percentage-point difference between the two areas (high minus low).

3.3. Influencing factors of CHOP exposure

Among the four candidate ensemble learners, all delivered high accuracy across both cities and scales of analysis, with R^2 above 0.87 in every configuration (Table 4). Within Guangzhou, XGBoost achieved the lowest RMSE and MAE at the 1,000 m scale, while LightGBM marginally led at the 100 m scale; within Shenzhen, LightGBM led at both scales. As the focus of the attribution analysis lies in identifying robust, city-specific influencing factors rather than in optimizing predictive accuracy alone, and as XGBoost combined consistent top-tier performance with stable error metrics across scales and well-documented compatibility with SHAP-based interpretation, XGBoost was selected as the primary model for SHAP-based interpretation in both cities.

Table 4. Predictive performance of four ensemble models for CHOP exposure across spatial scales in

Guangzhou and Shenzhen.

Model	Scale	Guangzhou			Shenzhen		
		R^2	RMSE	MAE	R^2	RMSE	MAE
XGBoost	100 m	0.908	1.666	1.225	0.901	1.412	1.011
	1,000 m	0.965	1.366	0.962	0.918	1.270	0.875
CatBoost	100 m	0.902	1.715	1.277	0.899	1.425	1.037
	1,000 m	0.964	1.401	1.014	0.918	1.270	0.883
LightGBM	100 m	0.906	1.682	1.247	0.904	1.387	0.992
	1,000 m	0.964	1.388	0.990	0.918	1.269	0.860
Random Forest	100 m	0.877	1.929	1.465	0.881	1.551	1.128
	1,000 m	0.950	1.650	1.200	0.899	1.408	0.996

As a robustness check across spatial scales, we examined the cross-validated performance of the selected XGBoost model across the spatial validation patterns and evaluation scales (Table 5). The 100 m and 1,000 m pipelines are consistent in their overall pattern, and the predictive performance of the 100 m pipeline is preserved under both grouped_1000m and the stricter block2000 splits. Under grouped_1000m, the aggregated_1000m R^2 reaches 0.965 in Guangzhou and 0.918 in Shenzhen; under block2000, it remains 0.950 in Guangzhou and 0.880 in Shenzhen, indicating that the spatial patterns identified by the 100 m pipeline are not driven by within-parent label duplication or short-range spatial dependence between adjacent units. The 1,000 m pipeline, evaluated under random and block2000 splits, follows a consistent overall pattern at a lower magnitude. These results support the SHAP-based interpretation drawn from the 100 m pipeline as the basis for the city-specific interpretation that follows.

Table 5. Cross-validated predictive performance of XGBoost under alternative spatial cross-validation schemes and evaluation scales (100 m and 1,000 m pipelines).

Pipeline	Spatial CV scheme	Evaluation scale	GZ			SZ		
			R^2	RMSE	MAE	R^2	RMSE	MAE
100 m	grouped_1000m	row_100m	0.908	1.666	1.225	0.901	1.412	1.011
		aggregated_1000m	0.965	1.366	0.962	0.918	1.271	0.875
	block2000	row_100m	0.878	1.930	1.411	0.848	1.622	1.156

Pipeline	Spatial CV scheme	Evaluation scale	GZ			SZ		
			R^2	RMSE	MAE	R^2	RMSE	MAE
1,000 m	random	aggregated_1000m	0.950	1.632	1.154	0.880	1.422	0.939
		row_1000m	0.907	2.226	1.613	0.819	1.883	1.362
	block2000	row_1000m	0.874	2.593	1.912	0.770	1.969	1.424

SHAP-based attribution of the selected XGBoost model reveals different combinations of dominant influencing factors in Guangzhou and Shenzhen (Figs. 4 and 5). In Guangzhou, distance to the urban center (Dist_Center) is the most influencing factor, with a population-weighted mean |SHAP| of 24.8% (Fig. 4). Its dependence relationship is monotonic and gradual: SHAP values are around +3 to +5 within roughly 10–20 km of the center, decline through the 20–50 km range, and reach values between −5 and −10 beyond 50 km, indicating that central locations contribute additional CHOP days while peripheral districts such as Conghua suppress them. Population density (Pop_Dens) ranks second at 21.5% and shows a clear threshold pattern, with SHAP values fluctuating around zero at low densities and turning strongly positive at higher densities. Road intersection density (Road_Dens) follows at 15.8%, with negligible contribution below 200 intersections/km² and SHAP values rising into the +2 to +3 range above 200–300 intersections/km². The combined contribution of population and road densities reaches 37.3%, indicating that horizontal development intensity, expressed through population concentration and road network density, jointly shapes Guangzhou's compound exposure on top of the underlying core-to-periphery gradient. Building density (Bldg_Density, 7.6%) and the building coverage ratio (BCR, 2.3%) add further positive contributions, reinforcing the role of horizontal high-density development. Average housing price (Price_Avg, 5.6%) shows a non-monotonic response, with positive SHAP values at low to moderate prices and a turning point in the higher-price range. Land-use mix (LU_Mix, 4.3%) and community-level population (Community_Pop, 4.0%) contribute modestly, while the Normalized Difference Vegetation Index (NDVI, 3.8%) consistently contributes negatively, with greener cells associated with lower predicted CHOP exposure.

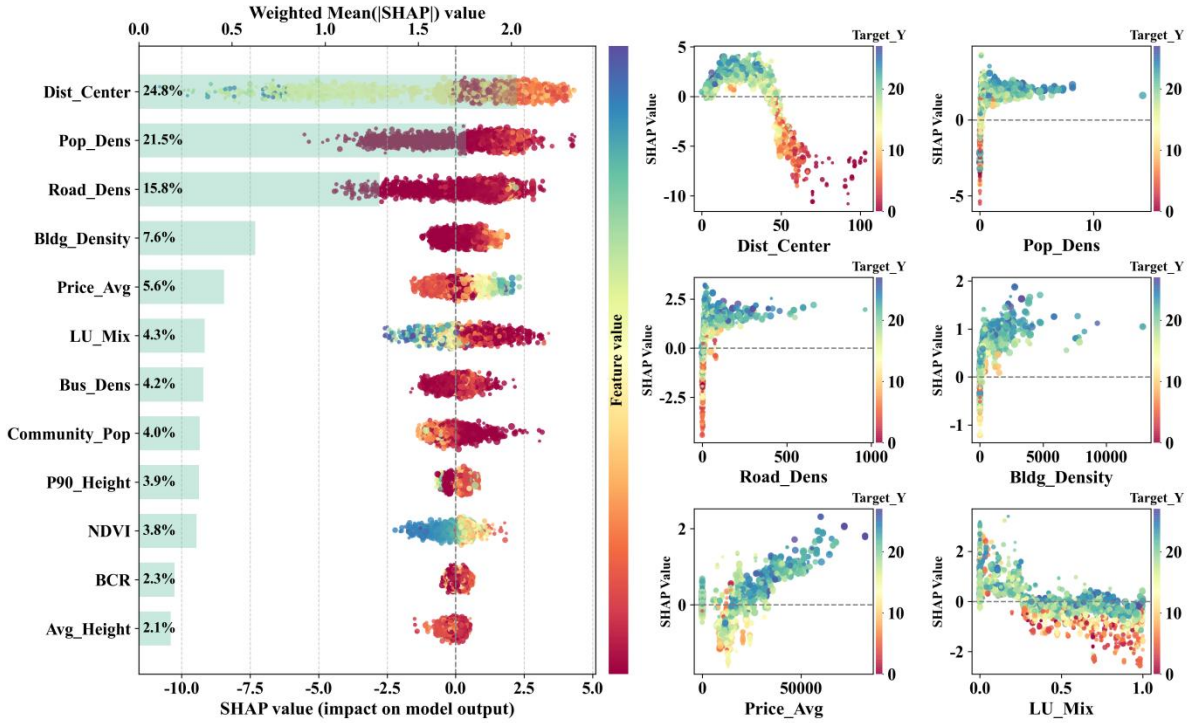


Fig. 4. Population-weighted SHAP feature importance and dependence plots for Guangzhou.

In Shenzhen, distance to the urban center (Dist_Center) remains the most influencing factor but operates over a more compact effective range, with a population-weighted mean |SHAP| of 22.2% (Fig. 5). The dependence plot shows that positive SHAP values are concentrated within roughly 20 km of the urban core and turn negative at intermediate distances, indicating a stronger coupling between high CHOP exposure and the urban core and a more distinct suppressing role of the eastern ecological belt. The most striking divergence from Guangzhou lies in the relative weight of 3D building morphology among the secondary factors. Mean building height (Avg_Height) ranks second at 17.6%, with SHAP values rising sharply from near zero at low heights into the +2 to +4 range above roughly 10–20 m and stabilizing thereafter. The building coverage ratio (BCR) ranks third at 9.6%, contributing close to zero below 0.20 and turning positive above 0.25–0.30. Population density (Pop_Dens) ranks fourth at 8.5%, substantially lower than its 21.5% in Guangzhou but still a non-trivial contributor. Together with FAR (3.1%), building density (Bldg_Density, 4.9%), and the share of tall buildings (Tall_Bldg_Share, 3.0%), three-dimensional building morphology accounts for a combined 38.2% of feature importance in Shenzhen, compared with roughly 16% in Guangzhou; horizontal development intensity (Pop_Dens + Road_Dens) shows the opposite pattern, contributing 15.6% in Shenzhen versus 37.3%

in Guangzhou. Average housing price (Price_Avg, 7.6%) again exhibits a nonlinear, threshold-like effect, though the precise turning point should be verified against the dependence plot. Community-level population (Community_Pop, 6.4%) and bus stop density (Bus_Dens, 6.2%) play moderate roles, with Bus_Dens contributing positively where transit-node density is highest. The Normalized Difference Vegetation Index (NDVI) is consistently negative in both cities (3.8% in Guangzhou and a similar magnitude in Shenzhen), and among socio-demographic predictors, the share of low-income and elderly residents shows positive SHAP values in both cities, indicating spatial co-occurrence of these groups with higher predicted CHOP exposure.

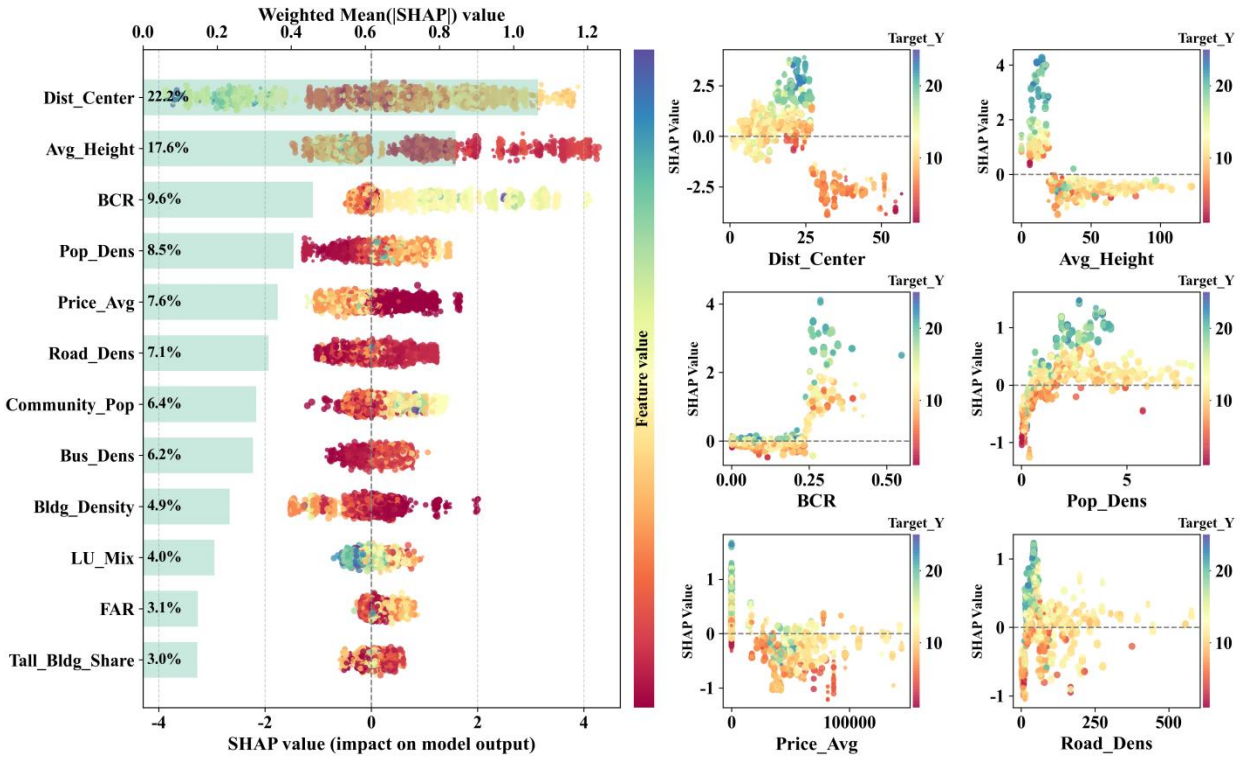


Fig. 5. Population-weighted SHAP feature importance and dependence plots for Shenzhen

4. Discussion

4.1. Two contrasting intra-urban exposure patterns

Even under a shared regional climate and a common ozone-pollution background, Guangzhou and Shenzhen develop sharply different intra-urban exposure patterns, a contrast that earlier macro-scale assessments of compound exposure (Cui et al., 2025; J. Su et al., 2025) cannot directly

capture. Guangzhou exhibits a continuous, smoothly varying core-to-periphery exposure gradient, while Shenzhen displays a localized, fragmented pattern in which a central-western high-exposure cluster coexists with a sizeable low-exposure ecological belt. This contrast suggests that compound exposure cannot be inferred from regional climate or aggregate emission inventories alone; the continuity of the built-up area, the configuration of ecological space, and the dimensionality of urban form jointly shape intra-urban exposure.

The exposure pattern in Guangzhou reflects the cumulative effects of long-standing horizontal urbanization. Its historical core and southward extension share high impervious-surface ratios, dense road networks, and intense human activity, all of which amplify the urban heat island and provide both the precursor base and the photochemical conditions for ozone formation. National-scale studies have shown that compound heat–ozone exposure rose substantially in Chinese cities between 2003 and 2020, with central districts of some megacities contributing more than 80% of citywide compound exposure by 2020 (J. Su et al., 2025). The Guangzhou pattern is consistent with this national trend but with one difference: rather than concentrating exclusively in the historical core, Guangzhou's compound exposure has expanded southward into Haizhu, Panyu, and the broader contiguous built-up zone, where population and functional spillover from the core, accelerated urban development, and intensified transport connections act together to elevate heat exposure, ozone formation, and population activity within a single urban tissue.

Shenzhen's exposure pattern, by contrast, is more localized and confined to a dense central-western cluster. The central-western corridor concentrates industrial space, transport infrastructure, logistics activity, and high-density residential functions, generating intense local emissions and human-activity loads that promote heat–ozone co-occurrence within a relatively confined area. Pearl River Delta studies have shown that ozone in the region is sensitive to both volatile organic compound (VOC) and nitrogen oxide (NOx) controls, and that local emissions and inter-city transport jointly shape ozone formation in core cities such as Guangzhou and Shenzhen (Peng et al., 2025). Shenzhen’s central-western cluster is therefore unlikely to be governed by a single administrative or sectoral process; rather, it emerges from the interaction of industrial emissions, traffic, urban-edge expansion, and regional transport. The relatively low exposure in eastern Shenzhen further illustrates how preserved ecological space and lower development intensity can substantially attenuate compound risk, even within the same megacity climate, and supports treating ecological space as an active component of compound-risk management.

4.2. Population burden and environmental justice

From an environmental-justice perspective, the differences in CHOP exposure between Guangzhou and Shenzhen are not merely environmental gradients but signals of unequal public-health burdens (Peng et al., 2026). The two cities show two contrasting configurations of compound exposure inequality. Guangzhou represents a citywide-burden configuration: a large majority of residents are jointly exposed to moderate-to-high compound levels, with relatively small differences across income, education, and life-stage groups, so that the pressure on public health is broad-based rather than concentrated. Shenzhen represents a localized-mismatch configuration: although the absolute exposed population is smaller, high-exposure areas show pronounced over-representation of lower-education, lower-income, younger working-age, and pregnancy-stage residents, indicating that compound risk is selectively concentrated on socio-economically and biologically more vulnerable groups.

Two related processes can help explain why the burden in Guangzhou is so broadly distributed. The first is structural: the contiguous, longstanding nature of Guangzhou's built-up area means that grid cells with moderately high CHOP exposure dominate large parts of the city, mechanically encompassing residents of diverse income, education, and life-stage groups. The second is residential-sorting: in a mature, polycentric urban fabric, advantaged groups have not been able to systematically relocate away from the contiguous built-up zone, because exposure differentiation across that zone is comparatively muted. Consistent with this interpretation, Guangzhou's high- and low-exposure neighborhoods differ by less than 2 percentage points across most income and education categories. The resulting inequality is therefore primarily one of magnitude rather than of distribution: the public-health burden of compound exposure is substantial precisely because it is shared by such a large fraction of the resident population.

Shenzhen's pattern, by contrast, aligns more closely with the classic environmental-justice expectation that disadvantaged groups disproportionately settle in heat-exposed and pollution-burdened neighborhoods (Tian et al., 2025). Higher-income and higher-education households can leverage residential choice and adaptive resources to avoid central-western high-exposure areas, while lower-income and lower-education residents face stronger constraints from housing affordability, employment-related residential location, and commuting cost, making it harder for them to relocate to safer environments. The over-representation of younger workers and pregnancy-stage residents in Shenzhen's high-exposure areas extends this argument: in rapidly

urbanizing, mobile cities, young workers' residential choices are shaped by job opportunities, housing costs, commuting distances, working hours, and outdoor activity patterns, which can push them toward higher-exposure zones regardless of physiological vulnerability. The Guangzhou–Shenzhen comparison extends this evidence to compound heat–ozone exposure in highly urbanized Chinese megacities and shows that two contrasting inequality configurations, a broad-based public-health burden and a localized vulnerability mismatch, can coexist within a single urban agglomeration.

4.3. Urban-form influencing factors of the two exposure patterns

The SHAP-based interpretation makes the urban-spatial influencing factors behind the two patterns explicit. Distance to the urban center ranks first in both cities, indicating a clear core-to-periphery gradient in CHOP exposure. This variable is not a direct physical factor of heat or ozone formation; it functions as a composite proxy for development intensity, population activity, traffic accessibility, ecological-space share, and functional layout. Beyond this shared first-order driver, the secondary influencing factors diverge sharply between the two cities, and these differences map closely onto the two exposure patterns.

In Guangzhou, the prominence of population density and road density alongside central distance indicates that compound exposure is shaped predominantly by horizontal accumulation of population activity, road networks, and impervious surfaces in the contiguous built-up area. The strong connectivity of population, infrastructure, and function across the core and the southward built-up corridor allows thermal stress, ozone formation, and population exposure to coalesce within a single urban tissue. In Shenzhen, by contrast, mean building height, building coverage ratio, and other three-dimensional building features rank highest after central distance, indicating that compound exposure is more closely tied to three-dimensional building morphology. These variables capture not only physical form but also the joint effects of vertical density, compact street structure, and population clustering. Compound risk in Shenzhen is therefore more likely to emerge in vertically intensified, mixed-function blocks, giving the risk a stronger neighborhood-scale and morphology-based character than in Guangzhou.

Several variables operate consistently across the two cities and complement the city-specific contrast. NDVI is consistently negative, indicating that high-greenness areas suppress predicted CHOP exposure. In a compound heat–ozone context, the protective role of vegetation likely arises

from the bundle of conditions that high-greenness areas tend to exhibit, including lower development intensity, less imperviousness, lower population activity, and better ventilation. NDVI thus marks settings that contrast with high-intensity built environments, helping to explain the lower exposure observed in the eastern districts of Shenzhen and the peripheral mountainous areas of Guangzhou. Housing price and socio-demographic variables further suggest that CHOP exposure is shaped not only by physical processes but also by the social-spatial differentiation of urban areas. Lower housing prices, higher shares of low-income residents, and higher shares of older adults are associated with elevated exposure in both cities, consistent with the joint influence of housing-market segmentation, public-service allocation, and residential-choice constraints (Ma et al., 2025) .

4.4. Policy implications

The two contrasting exposure patterns call for differentiated rather than uniform policy responses. CHOP risk governance should move beyond single-hazard regulation toward spatially differentiated, equity-oriented compound-risk management, with the focus and instruments matched to the urban-form factors associated with exposure in each city (Tian et al., 2025).

For Guangzhou, the central policy priority is to reduce the long-term compound exposure shared by a large majority of residents in the contiguous built-up zone. As the burden is broad-based rather than concentrated on a small disadvantaged subgroup, policy responses should emphasize area-wide retrofits across the connected built-up zone, including more street-level shading, pocket parks, blue–green corridors, permeable pavement, and improved cooling around aging neighborhoods, schools, community health centers, and commuting corridors (Tian et al., 2025; Yin et al., 2023). District-scale integrated retrofits that combine road cooling, building-envelope efficiency, public cooling spaces, and ozone-precursor emission reduction are likely to deliver greater benefit than single-point greening, given the spatially extensive nature of the compound burden.

For Shenzhen, the central policy priority is twofold: to address the localized high-exposure areas in the central-western corridor, particularly emerging high-development zones such as Guangming, and to safeguard the relatively low-exposure status of the eastern ecological belt. As the exposure pattern is more localized and shaped by three-dimensional building morphology, CHOP risk management should be incorporated into planning, land development, transport

configuration, and housing supply at the build-out stage rather than retrofitted after high-density development is complete. Cooling, ventilation, shading, and precursor reduction should be embedded as design constraints in new corridors and high-rise clusters, with controls on building coverage and the protection of ventilation pathways particularly important in vertically intensified blocks. At the same time, the eastern ecological belt anchored by Dapeng should be treated as part of the city's compound-risk management system, by maintaining ecological continuity, controlling encroachment, and preserving land–sea breeze pathways.

Beyond city-specific responses, CHOP governance more generally requires cross-sector coordination among meteorological, environmental, planning, transport, housing, and public-health authorities. A joint heat–ozone early-warning system, integrating heat alerts, ozone alerts, gridded exposure maps, and sensitive-population data into street- or community-level response plans, can support the activation of outdoor-work adjustments, staggered commuting, community visits, maternal-health alerts, and the opening of cooling centers during high-risk periods. Ozone precursor management should be tightly coupled with heat-risk control: as compound exposure tends to coincide with elevated temperatures, intense human activity, and dense built environments, VOC and NOx control measures should be prioritized in compound high-exposure areas rather than allocated by district-average concentrations alone. Guangzhou should reinforce time-targeted management of traffic, residential, construction, and road-related emissions across the contiguous built-up zone, while Shenzhen should focus on mobile sources and VOCs around the central-western corridor, transit nodes, and emerging high-development zones, especially during seasons when ozone and heat extremes co-occur.

5. Conclusion

This study compares fine-scale compound heat–ozone pollution exposure between Guangzhou and Shenzhen, two megacities of the Greater Bay Area that share a common regional climate yet differ markedly in urban form. The analysis integrates daily LST and MDA8 ozone with socio-demographic and built-environment attributes on a unified 100 m grid, and compares four ensemble models under spatial-block cross-validation, and interprets the selected model using SHAP-based attribution. The main findings are as follows:

This study provides a fine-scale comparative assessment of compound heat–ozone pollution exposure between Guangzhou and Shenzhen, two megacities that share a common regional climate

yet differ markedly in urban form. By integrating 100 m population rasters with daily LST and MDA8 ozone as environmental-exposure data and with socio-demographic and built-environment attributes spanning both two-dimensional urban form and three-dimensional building morphology, and by coupling four ensemble learners with SHAP attribution under spatially blocked validation, the analysis yields the following findings:

(1) The two megacities exhibit different intra-urban CHOP exposure patterns. Guangzhou displays a continuous, smoothly varying core-to-periphery gradient that extends from the traditional core to the southern built-up corridor, with a citywide mean exposure of 17.7 days and stronger internal inequality (Gini 0.30). Shenzhen displays a more localized, fragmented pattern in which a central-western cluster (notably Guangming and Bao'an) coexists with a sizeable low-exposure eastern ecological belt, with a citywide mean exposure of 11.8 days and a comparatively flatter overall distribution (Gini 0.21).

(2) Population exposure burdens differ in both magnitude and demographic composition. In Guangzhou, about 90.3% of residents experience more than 15 CHOP days per year, generating a broad-based citywide public-health pressure with comparatively small socio-demographic disparities. In Shenzhen, the share of residents in the same exposure tier is only 17.1%, but high-exposure areas show systematic over-representation of lower-education, lower-income, younger working-age residents, and households at the pregnancy stage, indicating a pronounced spatial mismatch between localized high-risk environments and socially or biologically vulnerable groups.

(3) The two patterns are associated with different urban-form factors. Distance to the urban center is the most influencing factor in both cities, but the secondary influencing factors diverge: in Guangzhou, exposure is associated predominantly with horizontal development intensity (population density and road density), reflecting long-standing horizontal urbanization; in Shenzhen, the opposite pattern holds, with exposure more strongly tied to three-dimensional building morphology (building height and coverage), reflecting compact, vertically intensified development.

Several limitations should be acknowledged. First, the analysis identifies population-weighted exposure differences rather than directly observed health outcomes, and therefore cannot be interpreted as morbidity, mortality, or individual-level health disparities; future work should couple gridded CHOP exposure with epidemiological records to translate exposure inequality into

health-damage inequality. Second, the population exposure estimates are based on static residential distributions and do not fully capture dynamic exposure arising from commuting, work, and outdoor activity; integrating mobility data would help characterize true population-level compound exposure. Third, the study does not include emission-source apportionment, land–sea breeze diagnostics, or street-canyon ventilation simulations, so the interpretation of precursor sources, coastal meteorology, and local dispersion remains a mechanistic inference; future research should incorporate chemical transport modeling and microclimate simulation to characterize the physical and chemical pathways linking urbanization processes to compound exposure.

References

Ban, J., Lu, K., Wang, Q., Li, T., 2022. Climate change will amplify the inequitable exposure to compound heatwave and ozone pollution. One Earth 5, 677–686. https://doi.org/10.1016/j.oneear.2022.05.007

Breiman, L., 2001. Random forests. Machine learning 45, 5–32. https://doi.org/10.1023/A:1010933404324

Cheval, S., Amihăesei, V.-A., Chitu, Z., Dumitrescu, A., Falcescu, V., Irașoc, A., Micu, D.M., Mihulet, E., Ontel, I., Paraschiv, M.-G., Tudose, N.C., 2024. A systematic review of urban heat island and heat waves research (1991–2022). Clim. Risk Manage. 44, 100603. https://doi.org/10.1016/j.crm.2024.100603

Cui, M., Hong, C., Liu, W., He, K., Zhang, Q., 2025. Future warming exacerbates heatwave-ozone compound extremes in China. Npj Clim. Atmos. Sci. 8, 341. https://doi.org/10.1038/s41612-025-01218-2

Domeisen, D.I.V., Eltahir, E.A.B., Fischer, E.M., Knutti, R., Perkins-Kirkpatrick, S.E., Schär, C., Seneviratne, S.I., Weisheimer, A., Wernli, H., 2023. Prediction and projection of heatwaves. Nat. Rev. Earth Environ. 4, 36–50. https://doi.org/10.1038/s43017-022-00371-z

Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16). ACM, New York, pp. 785–794. https://doi.org/10.1145/2939672.2939785

Ghanbari, M., Arabi, M., Georgescu, M., Broadbent, A.M., 2023. The role of climate change and urban development on compound dry-hot extremes across US cities. Nat. Commun. 14, 3509. https://doi.org/10.1038/s41467-023-39205-x

Huang, Z., Luo, M., Gao, M., Ning, G., Ge, E., Chan, T.O., Wu, S., Zhang, H., Tang, Y., 2023. Different characteristics of independent and compound extreme heat and ozone pollution events in the greater bay area of China. Int. J. Appl. Earth Obs. Geoinf. 124, 103508. https://doi.org/10.1016/j.jag.2023.103508

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.-Y., 2017. Lightgbm: A highly efficient gradient boosting decision tree. Adv. neural inf. process. syst. 30.

Liang, L., Engling, G., Zhang, X., Sun, J., Zhang, Y., Xu, W., Liu, C., Zhang, G., Liu, X., Ma, Q., 2017. Chemical characteristics of PM 2.5 during summer at a background site of the Yangtze River Delta in China. Atmos. Res. 198, 163–172. https://doi.org/10.1016/j.atmosres.2017.08.012

Lundberg, S.M., Lee, S.-I., 2017. A unified approach to interpreting model predictions. Adv. neural inf. process. syst. 30.

Ma, K., Fang, F., Lin, Y., Tong, F., He, C., Yao, Y., Zhu, J., Wang, H., Xing, X., 2025. Spatial heterogeneity in the joint effect of heat wave and ozone pollution on mortality at the township-level in China. Sustainable Cities Soc. 106807. https://doi.org/10.1016/j.scs.2025.106807

Namaiti, A., Zeng, S., He, W., Liu, X., Zeng, J., 2026. Exploring diurnal spatiotemporal heterogeneity in urban heat exposure: A novel perspective from urban form-function coupling. Sustainable Cities Soc. 138, 107161. https://doi.org/10.1016/j.scs.2026.107161

Peng, J., Jang, C., Zhu, Y., Xing, J., Wang, S., Zhao, B., Long, S., Li, J., Wen, Q., Yan, X., 2025. Quantifying multi-year ozone precursor contributions over the pearl river delta region of China using a deep learning-based response surface model. Front. Environ. Sci. Eng. 19, 119. https://doi.org/10.1007/s11783-025-2039-y

Peng, J., Wang, Q., Yang, Z., Dong, J., Yu, X., Corcoran, J., 2026. Anthropogenic climate change drives rising global heat stress and its spatial inequality. Nat. Commun. 17, 2310. https://doi.org/10.1038/s41467-026-69164-y

Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A., 2018. CatBoost: Unbiased boosting with categorical features. Adv. neural inf. process. syst. 31.

Schnell, J.L., Prather, M.J., 2017. Co-occurrence of extremes in surface ozone, particulate matter, and temperature over eastern north America. Proc. Natl. Acad. Sci. 114, 2854–2859. https://doi.org/10.1073/pnas.1614453114

Su H., Qi Z., Wang Q., 2025. Impacts of land use characteristics on extreme heat events: Insights from

explainable machine learning model. Sustainable Cities Soc. 120, 106139. https://doi.org/10.1016/j.scs.2025.106139
Su, J., Jiao, L., Xu, G., 2025. Intensified exposure to compound extreme heat and ozone pollution in summer across Chinese cities. Npj Clim. Atmos. Sci. 8, 78. https://doi.org/10.1038/s41612-025-00966-5
Tian, P., Yan, Y., Li, J., Zhang, H., Liu, Y., Ying, C., 2025. Assessing compound exposure of vulnerable populations under dual inequality of urban heatwaves and green space. Sustainable Cities Soc. 134, 106951. https://doi.org/10.1016/j.scs.2025.106951
Wei, J., Li, Z., Li, K., Dickerson, R. R., Pinker, R. T., Wang, J., Liu, X., Sun, L., Xue, W., & Cribb, M. (2022). Full-coverage mapping and spatiotemporal variations of ground-level ozone (O_3) pollution from 2013 to 2020 across China. Remote Sens. Environ., 270, 112775.
Xu, R., Bi, Y., Jia, X., Zheng, Y., Feng, N., Wang, S., Luo, L., Lv, Z., Huang, S., Zheng, X.-Y., Sun, H., Chen, G., Wei, J., Liu, Y., 2026. Global spatiotemporal variation and death burden of compound heatwave and ozone pollution events. Environ. Sci. Technol. 60, 8404–8414. https://doi.org/10.1021/acs.est.5c17533
Xu, R., Sun, H., Zhong, Z., Zheng, Y., Liu, T., Li, Y., Liu, L., Luo, L., Wang, S., Lv, Z., Huang, S., Shi, C., Chen, W., Wei, J., Xia, W., Liu, Y., 2024. Ozone, heat wave, and cardiovascular disease mortality: A population-based case-crossover study. Environ. Sci. Technol. 58, 171–181. https://doi.org/10.1021/acs.est.3c06889
Yin, Y., He, L., Wennberg, P.O., Frankenberg, C., 2023. Unequal exposure to heatwaves in los angeles: Impact of uneven green spaces. Sci. Adv. 9, eade8501. https://doi.org/10.1126/sciadv.ade8501
Zhou, C., Wang, G., Jiang, H., Li, S., Shi, X., Hu, Y., Cabral, P., 2025. Spatio-temporal patterns of compound dry-hot extremes in China. Atmos. Res. 314, 107795–107807. https://doi.org/10.1016/j.atmosres.2024.107795
Zhou, X., Li, M., Huang, X., Liu, T., Zhang, H., Qi, X., Wang, Z., Qin, Y., Geng, G., Wang, J., 2025. Urban meteorology–chemistry coupling in compound heat–ozone extremes. Nat. Cities 2, 847–856. https://doi.org/10.1038/s44284-025-00302-1

Data Availability Statement: Data could be found in https://doi.org/10.6084/m9.figshare.32396964.

Acknowledgments: The authors would like to thank to the anonymous reviewers for their helpful comments.
This work was supported by the Research Start-up Fund for Young Faculty at Beijing Normal University (Grant No. 1263100020).

Competing Interests Statement: The authors declare that they have no financial and personal relationships with other people or organizations that can inappropriately influence their work.