

Measuring Visibility Bias in Digital Urban Experience: A Multimodal Spatial Statistical Analysis of Sina Weibo Data in Guangzhou

HUIMIN QU

Guangzhou Huashang College

Gongxiang Huang

gongxianghuang64@gmail.com

Guangzhou Huashang College

Zhuqin Liang

Universiti Sains Malaysia

Pei Wen

Guangzhou Huashang College

Article

Keywords: digital urban experience, visibility bias, social media data, Sina Weibo, spatial statistics, null model, LISA, ResNet-18, Guangzhou

Posted Date: August 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10550099/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Measuring Visibility Bias in Digital Urban Experience: A Multimodal Spatial Statistical Analysis of Sina Weibo Data in Guangzhou

Authors:

Huimin Qu¹, Gongxiang Huang^{1*}, Zhuqin Liang^{2*}, Pei Wen¹

¹ School of creativity and design, Guangzhou Huashang College, 511328 Guangzhou, China

² School of Mathematical Sciences, Universiti Sains Malaysia, 11800, Penang, Malaysia

Corresponding Authors:

- Gongxiang Huang (Email: gongxianghuang64@gmail.com)
- Zhuqin Liang (Email: liangzhuqin@student.usm.my)

Abstract

Social media data are increasingly used in computational urban research, yet their spatial distributions are shaped by uneven participation, platform affordances and selective representation. This study develops a bias-aware spatial statistical framework for measuring digital urban visibility using 273,566 geotagged Sina Weibo posts from Guangzhou, China, in 2023. Rather than treating raw social media density as a proxy for urban importance, the study constructs an activity baseline and uses a Visibility Bias Index (VBI) to measure whether specific categories are overrepresented or underrepresented relative to overall posting activity. Text sentiment was classified into positive, neutral and negative categories, while user-generated images were analysed using a multi-label ResNet-18 classifier covering cultural and social life, natural landscape, urban architecture and public space. The image classifier achieved strong held-out validation performance (micro-F1 = 0.986), although this result is interpreted in relation to broad image categories and possible image-level similarity. Permutation-based null models, Z-scores, p-values, Global Moran's I and Local Indicators of Spatial Association were used to identify statistically significant and spatially clustered visibility bias. Results

show that visibility bias is not evenly distributed across Guangzhou. Emotional visibility is scale-dependent, while architectural visibility is consistently clustered across 500 m, 1000 m and hotspot-removal robustness tests. By linking statistical results to Guangzhou's modern CBD, historic districts, waterfront landscapes and mobility nodes, the study shows how platform-mediated urban representation can be quantified and geographically interpreted.

Keywords: digital urban experience; visibility bias; social media data; Sina Weibo; spatial statistics; null model; LISA; ResNet-18; Guangzhou

List of abbreviations: CNN, convolutional neural network; LISA, Local Indicators of Spatial Association; MAUP, modifiable areal unit problem; NLP, natural language processing; VBI, Visibility Bias Index; VGI, volunteered geographic information.

1. Introduction

Digital traces generated through social media platforms have become a major empirical resource for computational social science, data science and urban analytics. Geotagged posts, short texts, user-generated photographs and check-in records provide high-volume evidence of how urban places are encountered, described, photographed and circulated online [1–9]. These data are valuable because they record not only physical presence but also expressive acts: a post, image or sentiment label indicates that a user considered a place communicable within a platform environment. For this reason, social media data are increasingly used to study urban perception, place attachment, mobility, tourism, public space and everyday experience [10–21].

However, the same properties that make social media data analytically attractive also make them methodologically risky. Social media data do not provide a neutral census of urban life. They are shaped by demographic inequalities, uneven platform use, algorithmic affordances, location-sharing practices and the

self-selective nature of online expression [22–28]. Recent work on geospatial text data further shows that the presence of usable location information is itself topic-dependent, reinforcing the need to treat geotagged social media as a selective rather than comprehensive spatial signal [29]. As a result, high online visibility may reflect symbolic attractiveness, event intensity, photographic distinctiveness or platform norms rather than the objective importance of a place.

This distinction is particularly important for studies of urban imagery and urban perception. Many social media-based urban studies identify hotspots of posting activity, sentiment or imagery and then interpret these hotspots as meaningful urban clusters [5, 7, 10, 30–32]. Such maps are useful, but they cannot determine whether a category is disproportionately represented once overall posting activity is taken into account. A central commercial district may contain many positive posts simply because it contains many posts overall. A transport hub may appear negative because it is a high-activity mobility node, not because negative expression is statistically overrepresented.

This paper addresses that problem by reframing digitally mediated urban experience as a problem of visibility bias. Visibility bias refers to the statistical deviation between the observed digital visibility of a category and its expected visibility under a baseline of overall posting activity. The aim is therefore not to rank places by popularity or to claim that social media represents urban reality directly. Instead, the study asks which sentiments and visual categories become disproportionately visible after controlling for general activity volume. This framing moves the analysis from descriptive urban mapping towards a bias-aware measurement problem in computational social science.

The proposed framework responds to three limitations in existing research. First, many studies use social media density as an indicator of urban vitality, popularity or perception, but rarely distinguish activity volume from category-specific representational bias [5, 7, 10] Second, multimodal urban analytics often

combine text and images descriptively without an explicit baseline or null model [31, 33–41]. Third, spatial clusters are frequently interpreted visually without testing whether they are statistically meaningful or robust to spatial scale [42–47]. Recent EPJ Data Science work on large-scale georeferenced tweets has similarly emphasized that reliable interpretation of emotional and event-related patterns requires robust baselines and explicit attention to data biases [48].

Guangzhou, China, provides a particularly suitable empirical context for testing this framework because of its strong spatial heterogeneity. As both a historic commercial city and a contemporary first-tier metropolis, Guangzhou contains Lingnan heritage environments, traditional commercial streets, modern CBD landscapes, transport hubs, waterfront public spaces, university districts, exhibition areas and rapidly redeveloped urban zones. This coexistence of historic and modern, symbolic and infrastructural, recreational and mobility-oriented spaces makes Guangzhou an appropriate case for examining whether different forms of digital visibility are associated with distinct urban functions and symbolic environments, rather than treating social media activity as a uniform expression of urban experience.

Using 273,566 geotagged Sina Weibo posts collected between 1 January and 31 December 2023, this study integrates text-based sentiment classification, multi-label ResNet-18 image classification, 500 m grid aggregation, a Visibility Bias Index, permutation-based null modelling, Moran's I, LISA and robustness checks using 1000 m grids and removal of the top 1% high-posting grids. The contributions are threefold. Conceptually, the paper advances selective digital visibility as a measurable dimension of urban experience. Methodologically, it proposes a bias-aware spatial statistical workflow for distinguishing raw activity from disproportionate representation. Empirically, it reveals how emotional and visual visibility are structured across Guangzhou and how these structures vary by category and spatial scale.

The study is guided by three research questions. RQ1: How can visibility bias in social media representations of urban space be quantified through multimodal data? RQ2: Which emotional and visual categories are significantly overrepresented or underrepresented relative to baseline social media activity? RQ3: Do visibility-bias patterns exhibit statistically significant spatial clustering, and are these patterns robust across spatial specifications?

2. Related work: from digital urban sensing to visibility bias

This section critically reviews three connected bodies of work. The first concerns the use of social media as digital urban sensing. The second concerns multimodal measurement of urban perception through text and imagery. The third concerns bias, selective visibility and spatial inference. Together, these literatures show why a data-science approach to social media urban analysis requires not only content extraction, but also baseline modelling, uncertainty-aware inference and geographically meaningful interpretation.

2.1 Social media as digital urban sensing

The use of social media data in urban research emerged from broader debates around volunteered geographic information, citizen sensing, social sensing and urban computing. Geotagged posts and images provide spatially referenced traces of human activity and have been used to identify areas of interest, tourism hotspots, urban mobility patterns, land-use functions and event dynamics [1, 2, 5, 7–10, 13–18, 49]. These studies show that platform data can supplement traditional urban datasets by capturing aspects of everyday urban behaviour and representation that are not always visible in administrative statistics.

Yet the metaphor of social media as a sensor is incomplete. A physical sensor is designed to measure a defined phenomenon under known conditions; social media platforms are designed for communication, social interaction, entertainment and self-presentation. The resulting data are therefore produced through

platform practices and user choices rather than through controlled sampling [22–28, 50, 51]. Treating them as direct observations of urban reality risks confusing digital expression with the underlying urban condition.

Existing studies have made important progress in using geotagged data to identify urban hotspots, infer functional areas and analyse event dynamics. However, many of these studies still rely on the geography of raw activity or density [5, 7–10, 13–18, 48, 52–54]. This creates a conceptual ambiguity: a dense social media cluster may indicate high pedestrian activity, strong symbolic attraction, tourism concentration, platform popularity or simply uneven user participation. The present study therefore treats social media not as a direct measure of urban importance, but as a platform-mediated trace of selective visibility.

2.2 Multimodal measurement of urban perception

Urban perception is not expressed through a single modality. Textual posts can reveal affective reactions, opinions, place narratives and event-related meanings, whereas images can reveal visually salient urban features such as landmarks, streetscapes, greenery, waterfronts, architectural forms and public-space scenes [31, 33–41, 55–58]. Research on sentiment analysis has provided tools for mapping emotional expression in cities, while computer vision and deep learning have expanded the possibility of analysing large image collections [33–41, 55–59]. Recent reviews also stress that user-perspective research in urban public spaces increasingly requires combined computer vision, natural language processing and higher-quality user data [38].

Despite these advances, multimodal urban analytics faces two conceptual challenges. First, text and images do not represent the same dimension of urban experience. A positive text may be attached to a visually ordinary scene, and an architecturally distinctive image may carry neutral or negative text. Second, multimodal features are often mapped as observed distributions without testing whether they are overrepresented relative to baseline activity [12, 22, 24, 32, 60].

This study uses sentiment classification and multi-label image classification not as direct measures of urban quality, but as indicators of digital emotional and visual representation. The four visual categories are intentionally broad and interpretable, aligning with urban imagery theory and prior work on visual urban analytics [33, 37, 39–41, 61–63]. This design supports grid-level spatial analysis while avoiding the over-interpretation of fine-grained object recognition.

2.3 Bias, selective visibility and spatial inference

The most important limitation of social media-based urban analysis is not simply sample size or classification accuracy, but representational bias. Social media data are affected by demographic bias, platform bias, spatial bias, temporal bias and reporting bias [18, 22–28, 32, 64, 65]. In urban contexts, these biases produce selective visibility: some spaces become repeatedly photographed, tagged and emotionally narrated, whereas other spaces remain digitally muted despite their everyday importance.

In spatial analysis, such bias becomes particularly consequential. Hotspot maps, kernel density surfaces and choropleth maps can reveal concentration, but they cannot determine whether concentration is statistically disproportionate. Spatial patterns are also sensitive to zoning and scale, a problem widely discussed as the modifiable areal unit problem [46, 47]. Therefore, social media-based urban analysis requires not only mapping, but also baseline comparison, spatial inference and robustness checks [42–45, 47].

A further limitation concerns statistical inference. Many studies identify visually apparent clusters but do not test whether the patterns differ from what would be expected under a null model. Moran's I and Local Indicators of Spatial Association provide established tools for testing global and local spatial autocorrelation [42–44]. In this study, these tools are combined with a permutation-based null model to distinguish meaningful visibility bias from random spatial fluctuation.

3. Data and methods

3.1 Study area, dataset and preprocessing

The empirical analysis focuses on Guangzhou, China, a large metropolitan region characterized by historic districts, commercial centres, waterfront landscapes, transport hubs, university areas, public spaces and rapidly transforming redevelopment zones. These diverse spatial contexts make Guangzhou suitable for analysing how digitally mediated visibility varies across different urban functions and symbolic environments.

The dataset consists of 273,566 geotagged Sina Weibo posts collected from 1 January to 31 December 2023. Each post contains spatial coordinates and textual content, and a subset of posts is associated with one or more user-generated images. Records without valid coordinates were removed, and posts were projected to EPSG:4547 for metric spatial analysis. After spatial assignment, 273,564 posts were linked to active grid cells. The main analytical unit is a 500 m by 500 m grid cell, producing 1,147 active grid cells. A 1000 m grid specification was used for scale robustness testing (see Table 1).

For the image component, 3,161 labelled images were used to train and validate a multi-label image classifier. The four categories were Cultural and Social Life, Natural Landscape, Urban Architecture and Public Space. Because a single image may contain more than one category, image counts are interpreted at the image-label level rather than as mutually exclusive image classes.

Ethical considerations and privacy

The source material was human-generated online content from Sina Weibo. The posts were publicly accessible at the time of collection and were retrieved programmatically for research; however, the resulting study dataset is not itself a publicly archived dataset and has no stable direct URL. This was a non-interventional observational study: no users were recruited, contacted, surveyed or subjected to an experiment. No prospective institutional ethics approval or individual informed consent was obtained. The authors treated

the work as secondary research using public online content rather than research involving recruited human participants. Account identifiers, usernames and profile fields were excluded from the analytical dataset before analysis. Images were used only to derive broad category labels and are not reproduced or shared, and all published spatial results are aggregated to 500 m or 1000 m grid cells. No individual user or post is identifiable in the manuscript or supporting outputs. All methods were carried out in accordance with applicable regulations and institutional guidance concerning privacy and the research use of publicly accessible online data.

Table 1. Dataset and analytical units.

Item	Value
Study area	Guangzhou, China
Platform	Sina Weibo
Observation period	1 January-31 December 2023
Original geotagged posts	273,566
Grid-assigned posts	273,564
Main spatial unit	500 m x 500 m grid
Active 500 m grids	1,147
Scale robustness unit	1000 m x 1000 m grid
Labelled image dataset	3,161 images
Coordinate reference system	EPSG:4547

3.2 Analytical framework

The analytical framework is designed to separate three concepts that are often conflated in social media-based urban analysis: total activity, category-specific representation and statistically significant visibility bias. Total activity refers to the overall number of posts in each grid cell. Category-specific representation refers to the distribution of positive sentiment, negative sentiment or visual categories across

the city. Visibility bias refers to the deviation between category-specific representation and the amount expected from total activity.

The workflow therefore proceeds from data cleaning and multimodal content extraction to grid aggregation, baseline construction, VBI calculation, null-model testing, spatial autocorrelation analysis and robustness checks. This structure allows the analysis to move beyond descriptive maps towards inferential claims about over-visibility, under-visibility and spatial clustering (see Figure 1).




Figure 1. Analytical framework for measuring multimodal visibility bias in digital urban experience.

3.3 Multimodal content classification

Textual content was analysed using SnowNLP to obtain a sentiment score between 0 and 1. Posts with scores below 0.4 were classified as negative, scores from 0.4 to 0.6 as neutral and scores above 0.6 as positive.

These thresholds produced 169,824 positive posts, 27,255 neutral posts and 76,485 negative posts after grid

assignment. A manually checked subset of 1,000 posts achieved an overall validation accuracy of 92.7%, providing a practical quality check for large-scale sentiment classification. Nevertheless, SnowNLP is a comparatively traditional Chinese sentiment tool and may be less effective than transformer-based or LLM-assisted approaches in detecting irony, sarcasm, emerging internet slang, and highly context-dependent emotional expression [58, 66]. Recent studies of Chinese and cross-lingual sentiment analysis show both the growing relevance of large language models and the continuing need to interpret model performance in relation to language, domain, and task setting [58, 66]. Given that the core of this study lies in the grid-level quantification of relative spatial bias rather than the diagnosis of fine-grained post-level emotions, large-scale sample aggregation over more than 270,000 posts helps reduce the influence of random classification noise associated with traditional lexicon-based or machine-learning classifiers. For this reason, the sentiment results are interpreted as broad affective indicators aggregated to grid level, rather than as definitive post-level psychological states.

For visual content, a ResNet-18 convolutional neural network pretrained on ImageNet was fine-tuned for multi-label classification of user-generated urban images. The final fully connected layer was replaced with a four-unit output layer corresponding to Cultural and Social Life, Natural Landscape, Urban Architecture and Public Space. A sigmoid activation function was used because the categories were not mutually exclusive, allowing a single image to be associated with multiple visual categories [33–37, 39–41, 67, 68].

The validation results indicate strong performance for the broad visual categories used in this study. The micro-averaged F1-score was 0.986, and category-level F1-scores ranged from 0.978 for Urban Architecture to 0.988 for Public Space. This high performance should be interpreted cautiously. The four categories are broad urban imagery categories rather than fine-grained object classes, and the held-out validation set contained 317 images. Because the validation split was performed at the image level, the

reported F1-scores should be interpreted as image-level validation performance rather than as a strict test of post-level, user-level, or spatial generalization. Visually similar images from the same post, user, or nearby scene may therefore have appeared in both the training and validation subsets, potentially inflating model performance. The visual labels are consequently used for grid-level aggregation and spatial comparison, while future work should apply grouped post-level, user-level, or spatially separated validation strategies to further reduce concerns about data leakage and overfitting (see Table 2).

Table 2. Performance of the ResNet-18 multi-label image classifier.

Category / average	Precision	Recall	F1-score	Support
Cultural and Social Life	0.989	0.984	0.987	187
Natural Landscape	1.000	0.975	0.987	160
Urban Architecture	0.991	0.966	0.978	116
Public Space	0.988	0.988	0.988	253
Micro average	0.992	0.980	0.986	716
Macro average	0.992	0.978	0.985	716
Weighted average	0.992	0.980	0.986	716
Samples average	0.988	0.979	0.981	716

3.4 Visibility bias measurement and spatial statistical validation

The baseline represents the expected visibility of each grid cell under the assumption that the spatial distribution of a given sentiment or image category follows the distribution of overall Weibo activity. Let $i = 1, \dots, N$ index the active grid cells, and let c denote a sentiment or image category. Let T_i denote the total number of posts in grid i . The baseline activity share of grid i is defined as:

$$B_i = \frac{T_i}{\sum_{j=1}^N T_j} \quad (1)$$

Let C_{ic} denote the number of observations belonging to category c in grid i . The observed spatial share of category c in grid i is defined as:

$$O_{ic} = \frac{C_{ic}}{\sum_{j=1}^N C_{jc}} \quad (2)$$

The Visibility Bias Index for category c in grid i is calculated as:

$$VBI_{ic} = \frac{O_{ic} + \epsilon}{B_i + \epsilon} \quad (3)$$

Where ϵ is a small positive constant introduced to ensure numerical stability for grid-category combinations with zero observations. A value of $VBI_{ic} > 1$ indicates that category c is overrepresented in grid i relative to the activity baseline, whereas $VBI_{ic} < 1$ indicates underrepresentation. A value equal to 1 indicates that the observed category share is consistent with the baseline expectation.

For mapping and spatial statistical analysis, the index was transformed using the natural logarithm:

$$\log VBI_{ic} = \ln(VBI_{ic}) \quad (4)$$

Under this transformation, $\log VBI_{ic} > 0$ indicates over-visibility, $\log VBI_{ic} < 0$ indicates under-visibility, and values close to zero indicate approximate consistency with the activity baseline. The logarithmic transformation produces a measure that is symmetric around zero and facilitates comparison between positive and negative deviations.

A permutation-based null model was used to determine whether observed VBI values were statistically meaningful. The null model preserves the spatial distribution of overall posting activity while randomizing category allocation under the expectation that category-specific visibility should follow the activity baseline. The process was repeated 999 times. For each grid and category, observed VBI values were compared with the simulated null distribution to obtain null means, Z-scores and p-values. Grid cells were classified as significantly over-visible or significantly under-visible at $p < 0.05$. This inferential step follows

the logic of spatial statistical testing, where observed patterns are interpreted against explicit null expectations [42–44].

Spatial clustering was assessed using Global Moran's I and LISA with queen contiguity weights. Global Moran's I evaluates whether logVBI values are spatially clustered citywide, while LISA identifies local High-High, Low-Low, High-Low and Low-High configurations [42–44]. Robustness was evaluated in two ways: repeating the workflow using 1000 m grids and recalculating Moran's I after removing the top 1% of 500 m grids with the highest overall posting intensity. These tests address scale sensitivity and the possibility that results are driven by extreme hotspots, issues closely related to the modifiable areal unit problem [46, 47].

4. Results

The results are organized around the inferential sequence of the framework. The analysis first establishes the baseline geography of overall Weibo activity, then evaluates emotional visibility bias, visual visibility bias, spatial clustering and robustness. The purpose is not to describe every map feature, but to determine which patterns remain meaningful after controlling for baseline activity and testing against null expectations.

4.1 Baseline activity and the need for bias measurement

Overall Weibo activity is strongly uneven across Guangzhou. Central and symbolically prominent areas generate substantially more posts than peripheral or less digitally active spaces. This pattern is expected in a large metropolitan region where tourism, commercial activity, transport intensity and public events are spatially concentrated. However, the baseline map also shows why raw density cannot be interpreted directly as emotional or visual importance (see Figure 2). Areas with many posts have greater opportunity to generate every type of content.

The baseline therefore functions as a control for exposure. It asks what level of category-specific visibility would be expected if sentiment or image labels followed overall posting activity. Subsequent VBI results should be interpreted as deviations from this baseline rather than as raw counts. This distinction is crucial because the same high-activity district may be over-visible for one category, under-visible for another and approximately baseline for a third.




Figure 2. Baseline activity distribution of geotagged Weibo posts across 500 m grid cells in Guangzhou. Posting intensity is shown as $\log(1 + \text{total posts})$ to reduce skewness caused by extremely high-activity grids. Landmark labels are provided for interpretive reference only.

4.2 Emotional visibility bias: significant deviations from baseline

Positive and negative sentiment show different forms of visibility bias. The logVBI maps indicate where emotional expression is disproportionately visible after controlling for total activity. Positive logVBI identifies grids where positive expression exceeds baseline expectation, while negative logVBI identifies grids where negative expression is overrepresented. These patterns cannot be inferred from raw sentiment counts alone, because high-activity grids may contain many positive and negative posts simply by volume.

The null-model results demonstrate that emotional visibility is not simply random variation around total activity. For positive sentiment, 208 grid cells are significantly over-visible and 179 are significantly under-visible. For negative sentiment, 222 grid cells are significantly over-visible and 255 are significantly under-visible. The presence of both over-visible and under-visible emotional grids confirms that the emotional geography of Weibo differs from the baseline geography of posting activity (see Figure 3 and Table 3).




Figure 3. Emotional visibility bias and null-model significance. Panels (a) and (b) show logVBI values; diverging colour scales are centred at zero. Panels (c) and (d) show significant over-visible and under-visible grids based on permutation testing.

Table 3. Null-model significance classification of selected visibility-bias indicators.

Category	Significant over-visible	Significant under-visible	Not significant
Positive sentiment	208	179	760
Negative sentiment	222	255	670
Urban architecture image	149	576	422
Public space image	260	575	312

4.3 Spatial clustering and scale-dependent emotional visibility

LISA analysis shows that emotional visibility contains meaningful local clusters even when global autocorrelation is weak at the 500 m scale. Positive logVBI contains local High-High and Low-Low configurations, but also many High-Low outliers, indicating that positive emotional over-visibility is often spatially localized rather than forming a strong citywide clustering structure. Negative logVBI shows a stronger local structure, with more High-High and Low-Low clusters than positive logVBI (see Figure 4).

This scale-dependence has an important substantive interpretation. Emotional expression in social media may be attached less to an individual 500 m cell than to a broader functional environment. In Guangzhou, positive visibility is plausibly associated with larger waterfront, commercial and leisure environments such as the Pearl River corridor, Huacheng Square, Haixinsha, Beijing Road, Shangxiajiu and other pedestrian or consumption-oriented areas. Negative visibility, by contrast, may be more closely linked to mobility pressure, crowding and waiting experiences around transport-related environments such as Guangzhou South Railway Station and major interchange areas. The 1000 m results therefore suggest that emotional visibility often reflects area-based experience rather than isolated micro-location effects.




Figure 4. Emotional LISA clusters for positive and negative sentiment logVBI.

4.4 Visual visibility bias: architecture as the most robust dimension

The visual visibility analysis focuses on Urban Architecture and Public Space because they represent two contrasting dimensions of digital urban imagery: visually distinctive built form and collective everyday spatial environments. Architecture logVBI shows strong selectivity, with 149 significantly over-visible grids and 576 significantly under-visible grids. This indicates that architectural imagery is not evenly distributed across the city; it is concentrated in specific places where built form is more visually recognizable, photographable or symbolically communicable (see Figure 5).

The robustness of architectural visibility is geographically meaningful in the Guangzhou case. Modern landmark environments such as Zhujiang New Town, Canton Tower, Huacheng Square, Haixinsha and Pazhou contain highly recognizable skylines, towers, exhibition buildings and waterfront views that are particularly likely to be photographed and shared. At the same time, historic districts such as Yuexiu and

Liwan, including heritage streets, Lingnan architectural forms, arcade streets, Shamian and Chen Clan Ancestral Hall-type cultural landscapes, provide another source of visually distinctive architectural representation. The stable architecture logVBI result therefore reflects Guangzhou's dual spatial identity as both a modern metropolitan centre and a historic commercial city.

Architecture is the most robust visual category. In the main 500 m model, architecture logVBI has a positive and significant Moran's I. It remains significant under the 1000 m grid specification and after removing the top 1% of high-posting grids. This robustness implies that architectural visibility bias is not merely a consequence of a few landmark hotspots or a particular grid resolution. Rather, it reflects a stable spatial structure in how built forms become digitally amplified.




Figure 5. Visual visibility bias and LISA clusters. Panels (a) and (b) show logVBI values; panels (c) and (d) show LISA clusters.

4.5 Robustness checks and synthesis

The robustness checks clarify which visibility-bias patterns are stable and which are scale-sensitive. Architecture logVBI is the most consistent result, remaining positive and significant in the 500 m main model, the 1000 m specification and the 500 m model excluding the top 1% high-posting grids. This finding supports the interpretation that architectural visibility bias represents a stable spatial pattern of visual amplification.

Emotional visibility is more scale-dependent. Positive and negative logVBI are not globally significant at 500 m, but both become significant at 1000 m. This does not mean that emotional visibility is absent. Instead, it suggests that affective expression is structured at broader spatial units, possibly associated with districts, functional corridors or larger place-based contexts. Cultural, natural and public-space logVBI show weaker global robustness, indicating that these dimensions of visibility may be more localized or spatially fragmented (see Table 4).

Taken together, the results demonstrate that digital urban experience is not simply a map of where people post. After controlling for baseline activity, testing against null expectations and examining spatial clustering, specific emotional and visual categories show different degrees of over-visibility, under-visibility, local clustering and scale sensitivity. The strongest and most robust evidence concerns the visual amplification of urban architecture, while emotional visibility requires multi-scale interpretation.

Table 4. Robustness checks of Global Moran's I under alternative spatial specifications.

Variable	500 m I (p)	1000 m I (p)	500 m without top 1% I (p)	Interpretation
Positive logVBI	-0.021 (0.699)	0.078 (0.022)	-0.021 (0.692)	Scale-sensitive; significant only at 1000 m
Negative logVBI	0.010 (0.360)	0.106 (0.006)	0.009 (0.374)	Stronger clustering at 1000 m
Cultural logVBI	0.034 (0.140)	-0.017 (0.650)	0.033 (0.139)	Not globally robust

Natural logVBI	0.028 (0.161)	0.042 (0.077)	0.026 (0.198)	Weak or marginal global clustering
Architecture logVBI	0.087 (0.006)	0.109 (0.002)	0.083 (0.006)	Robust across all tests
Public space logVBI	0.021 (0.218)	0.019 (0.216)	0.020 (0.248)	Not globally robust

Note. Global Moran's I was calculated using Queen contiguity weights and 999 permutations. Significant results are those with $p < 0.05$.

5. Discussion

5.1 From hotspots to visibility bias

The main contribution of this study is the shift from descriptive hotspot mapping to measurable visibility bias. In social media-based urban analysis, high-density maps are visually persuasive but inferentially limited. They show where content accumulates, but not whether a specific emotional or visual category is statistically overrepresented relative to overall activity [5, 7, 10, 32, 52–54]. VBI addresses this by introducing an explicit baseline and null-model comparison.

This distinction matters because it prevents raw social media volume from being mistaken for urban meaning. A place with high posting activity is not necessarily emotionally positive, visually architectural or publicly meaningful in a disproportionate sense. Conversely, a place with modest activity may be strongly over-visible for a particular category. By introducing a baseline and null model, the analysis makes the inferential target explicit and reduces the risk of over-interpreting descriptive concentrations.

For data science, this framing is valuable because it treats social media data as biased measurement rather than simply big data. The problem is not that social media data are unusable because they are biased; rather, the problem is that bias must be modelled, quantified and interpreted [22–28, 50, 51, 60]. Visibility bias offers one way to operationalize this principle in spatial urban analysis.

5.2 Urban-geographical mechanisms behind visibility bias in Guangzhou

The Guangzhou case shows that visibility bias is not only a statistical pattern but also an urban-geographical process. The most robust result is architectural visibility. This can be interpreted through the city's dual spatial identity: on the one hand, Guangzhou contains modern high-rise landscapes and recognizable metropolitan symbols around Zhujiang New Town, Canton Tower, Huacheng Square, Haixinsha and Pazhou; on the other hand, older districts such as Yuexiu, Liwan, Shamian and historic commercial streets contain Lingnan architectural heritage, arcade streets, temples, ancestral halls and culturally symbolic urban forms. These environments are more likely to be photographed, shared and circulated because they provide legible visual evidence of the city. The robustness of architecture logVBI across 500 m, 1000 m and hotspot-removal tests therefore suggests that built-form visibility is not driven only by a few extreme hotspots, but by a broader geography of visually communicable urban symbols.

Emotional visibility appears to operate differently. Positive and negative logVBI are weakly clustered at the 500 m scale but become significant at the 1000 m scale. This suggests that affective expression is often attached to larger functional areas, corridors or experiential zones rather than to individual grid cells. Positive visibility may emerge around waterfront leisure spaces, pedestrian commercial districts, scenic spaces, event-oriented plazas and cultural-tourism environments. Negative visibility may be associated with transport hubs, interchange areas, crowding, waiting, congestion and mobility pressure. Guangzhou South Railway Station, major metro-transfer environments and dense commercial mobility corridors are therefore important interpretive contexts for understanding why negative expression may become spatially meaningful only at a broader scale.

This difference between visual and emotional visibility is theoretically important. Architectural visibility is often object or scene-oriented: a tower, skyline, heritage street or landmark facade can be photographed within a specific grid and remain recognizable across platforms. Emotional visibility is more

relational and contextual: it may arise from the atmosphere of a district, the experience of moving through a transport node, the density of people in a commercial area, or the leisure setting of a waterfront. The results therefore suggest that different modalities of digital urban experience operate at different spatial scales and through different mechanisms of platform communication.

5.3 Implications for computational social science and urban interpretation

The findings have implications for computational social science because they show how platform-generated data can be treated as biased observational data rather than as a direct mirror of society. Social media visibility should not be interpreted as a simple proxy for public preference, spatial quality or inclusive urban experience [12, 22–28, 32, 60]. Instead, it should be understood as a platform-mediated signal shaped by user participation, communicability, visual distinctiveness, place function and the social norms of posting.

For urban studies and planning, this does not reduce the value of social media data; rather, it clarifies their appropriate use. VBI can identify which spaces are digitally amplified and which are comparatively muted after controlling for total activity. Over-visible spaces may reveal the symbolic anchors of the city, such as landmark CBD landscapes, heritage districts, waterfront spaces or event-oriented public areas. Under-visible spaces may indicate everyday, residential, infrastructural or less photogenic environments that remain central to urban life but are less likely to become part of platform-mediated narratives. Such information can complement, but not replace, surveys, interviews, fieldwork, land-use data and street-level observation [12, 33, 38, 39, 60].

The Guangzhou findings therefore should be read neither as a ranking of urban places nor as a direct measure of satisfaction. They show how different urban functions and symbolic environments become unevenly expressed through digital platforms. This distinction is particularly important for cities with high

internal heterogeneity, where modern CBDs, historic districts, transport hubs, waterfront landscapes and everyday residential spaces coexist but are not equally visible online.

5.4 Methodological implications and limitations

Methodologically, the study demonstrates the value of combining baseline comparison, null-model inference, spatial statistics and robustness checks. The framework does not assume that every mapped pattern is meaningful. Instead, it asks whether the pattern is disproportionate, statistically significant, spatially structured and robust to alternative specifications. This approach builds on established spatial statistical tools while adapting them to the problem of platform-mediated urban representation [29, 42–45, 47, 48].

Several limitations remain. First, the analysis is based on Sina Weibo and therefore reflects the users, affordances and posting norms of a single platform. The results should be interpreted as platform-mediated expressed visibility rather than as a representative account of all urban residents' perceptions [18, 22–28, 32, 65]. Cross-platform comparison would be valuable for testing whether similar visibility-bias patterns appear on image-centred, short-video or review-based platforms.

Second, the text and image classifiers have different validation profiles. SnowNLP offers a practical and reproducible approach for large-scale Chinese sentiment scoring, and the manually checked subset indicated acceptable overall accuracy. However, because SnowNLP is not a transformer-based contextual model, it may underperform on sarcasm, irony, ambiguous internet slang and highly context-dependent expressions. Future work should compare SnowNLP with BERT-based, RoBERTa-based, or LLM-based Chinese sentiment models and report category-level precision, recall, and F1 scores for the text classifier.

Third, the ResNet-18 image classifier achieved high validation performance, but the micro-F1 score of 0.986 should not be interpreted as evidence of fine-grained visual understanding. The classification scheme uses broad categories, and image-level random splitting may allow visually similar images from the same

post, user or location to appear in both training and validation subsets. Future work should apply grouped validation by post, user or spatial cluster, and should test whether model performance remains stable under stricter splitting strategies. These steps would further reduce concerns about overfitting and data leakage.

Fourth, the analysis is affected by spatial scale. The robustness results show that some forms of visibility are scale-sensitive. This is not merely a technical weakness; it is substantively important. Digital visibility may operate at different scales depending on whether the content concerns affect, architecture, public space or landscape. Future work should therefore treat scale as an object of analysis rather than only as a modelling parameter [46, 47].

6. Conclusion

This study developed a multimodal spatial statistical framework for measuring visibility bias in digital urban experience. Using 273,566 geotagged Sina Weibo posts from Guangzhou in 2023, the analysis combined sentiment classification, ResNet-18-based multi-label image classification, grid-level aggregation, baseline modelling, a Visibility Bias Index, permutation-based null testing, Global Moran's I, LISA and robustness checks.

The results show that social media representations of Guangzhou are spatially selective. Emotional visibility does not simply mirror total activity and becomes more evident at the 1000 m scale, suggesting broader affective geographies associated with functional areas, mobility environments and leisure districts. Visual visibility, especially urban architecture, shows the strongest and most stable spatial clustering across all robustness tests, reflecting Guangzhou's combination of modern metropolitan landmarks and historic Lingnan urban forms. These findings demonstrate that digital urban imagery is structured not only by where people post, but by which emotions and visual categories become disproportionately visible in particular urban spaces.

The study contributes conceptually by defining selective digital visibility as a measurable dimension of urban experience. It contributes methodologically by showing how baseline modelling, null models, spatial statistics and robustness checks can be integrated into social media-based urban analysis. It contributes empirically by revealing how emotional and visual visibility are differentially structured in Guangzhou. More broadly, the study argues that social media data should be interpreted not as a neutral proxy for urban reality, but as a biased and analyzable layer of platform-mediated urban representation.

References

1. Goodchild MF (2007) Citizens as sensors: The world of volunteered geography. *GeoJournal* 69:211–221. <https://doi.org/10.1007/s10708-007-9111-y>
2. Zheng Y, Capra L, Wolfson O, Yang H (2014) Urban computing: Concepts, methodologies, and applications. *ACM Trans Intell Syst Technol* 5:. <https://doi.org/10.1145/2629592>
3. Batty M (2013) *The New Science of Cities*. The MIT Press
4. Kitchin R (2014) *The Data Revolution: Big Data, Open Data, Data Infrastructures & Their Consequences*. SAGE Publications Ltd
5. Hu Y, Gao S, Janowicz K, et al (2015) Extracting and understanding urban areas of interest using geotagged photos. *Comput Environ Urban Syst* 54:240–254. <https://doi.org/10.1016/j.compenvurbsys.2015.09.001>
6. Resch B, Summa A, Sagl G, et al (2015) Urban Emotions—Geo-Semantic Emotion Extraction from Technical Sensors, Human Sensors and Crowdsourced Data. In: *Lecture Notes in Geoinformation and Cartography*
7. García-Palomares JC, Gutiérrez J, Mínguez C (2015) Identification of tourist hot spots based on social networks: A comparative analysis of European metropolises using photo-sharing services and GIS. *Applied Geography* 63:408–417. <https://doi.org/10.1016/j.apgeog.2015.08.002>
8. Soliman A, Soltani K, Yin J, et al (2017) Social sensing of urban land use based on analysis of Twitter users' mobility patterns. *PLoS One* 12:. <https://doi.org/10.1371/journal.pone.0181657>

9. Frias-Martinez V, Soto V, Hohwald H, Frias-Martinez E (2012) Characterizing urban landscapes using geolocated tweets. In: Proceedings - 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust and 2012 ASE/IEEE International Conference on Social Computing, SocialCom/PASSAT 2012
10. Cranshaw J, Schwartz R, Hong JI, Sadeh N (2012) The Livelihoods project: Utilizing social media to understand the dynamics of a city. In: ICWSM 2012 - Proceedings of the 6th International AAAI Conference on Weblogs and Social Media
11. Cui Y, Xie X, Liu Y (2018) Social media and mobility landscape: Uncovering spatial patterns of urban human mobility with multi source data. *Front Environ Sci Eng* 12:.. <https://doi.org/10.1007/s11783-018-1068-1>
12. Ilieva RT, McPhearson T (2018) Social-media data for urban sustainability. *Nat. Sustain.* 1
13. Girardin F, Blat J, Calabrese F, et al (2008) Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Comput* 7:.. <https://doi.org/10.1109/MPRV.2008.71>
14. Hollenstein L, Purves RS (2010) Exploring place through user-generated content: Using Flickr tags to describe city cores. *Journal of Spatial Information Science* 1:21–48. <https://doi.org/10.5311/JOSIS.2010.1.3>
15. Crandall D, Backstrom L, Huttenlocher D, Kleinberg J (2009) Mapping the world's photos. In: WWW'09 - Proceedings of the 18th International World Wide Web Conference
16. Kennedy L, Naaman M, Ahern S, et al (2007) How flickr helps us make sense of the world: Context and content in community-contributed media collections. In: Proceedings of the ACM International Multimedia Conference and Exhibition
17. Cheng Z, Caverlee J, Lee K (2010) You are where you tweet: A content-based approach to geo-locating Twitter users. In: International Conference on Information and Knowledge Management, Proceedings
18. Noulas A, Scellato S, Lambiotte R, et al (2012) A tale of many cities: Universal patterns in human urban mobility. *PLoS One* 7:.. <https://doi.org/10.1371/journal.pone.0037027>
19. Ahas R, Aasa A, Silm S, Tiru M (2010) Daily rhythms of suburban commuters' movements in the Tallinn metropolitan area: Case study with mobile positioning data. *Transp Res Part C Emerg Technol* 18:.. <https://doi.org/10.1016/j.trc.2009.04.011>
20. Kang C, Ma X, Tong D, Liu Y (2012) Intra-urban human mobility patterns: An urban morphology perspective. *Physica A: Statistical Mechanics and its Applications* 391:.. <https://doi.org/10.1016/j.physa.2011.11.005>
21. Longley PA, Adnan M, Lansley G (2015) The geotemporal demographics of twitter usage. *Environ Plan A* 47:.. <https://doi.org/10.1068/a130122p>
22. Boyd D, Crawford K (2012) Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Inf Commun Soc* 15:.. <https://doi.org/10.1080/1369118X.2012.678878>
23. Lazer D, Kennedy R, King G, Vespignani A (2014) The parable of google flu: Traps in big data analysis. *Science* (1979). 343
24. Olteanu A, Castillo C, Diaz F, Kiciman E (2019) Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *Front. Big Data* 2

25. Tufekci Z (2014) Big Questions for social media big data: Representativeness, validity and other methodological pitfalls. In: Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014
26. Ruths D, Pfeffer J (2014) Social media for large studies of behavior. *Science* (1979) 346:..
<https://doi.org/10.1126/science.346.6213.1063>
27. Morstatter F, Pfeffer J, Liu H, Carley KM (2013) Is the sample good enough? Comparing data from twitter's streaming API with Twitter's firehose. In: Proceedings of the 7th International Conference on Weblogs and Social Media, ICWSM 2013
28. Salganik MJ (2017) Bit by bit: Social research in the digital age. Princeton University Press
29. Mast J, Lemoine-Rodríguez R, Rittlinger V, et al (2026) Geospatiality: the effect of topics on the presence of geolocation in English text data. *International Journal of Geographical Information Science* 40:..
<https://doi.org/10.1080/13658816.2025.2460051>
30. Garcia-Palomares JC, Maria Henar S-O, Borja M-G, et al (2017) The pulse of the city through Twitter: relationships between land use and spatiotemporal demographics. *arXiv preprint arXiv:170507956*
31. Quercia D, O'Hare NJ, Cramer HSM (2014) Aesthetic capital: What makes london look beautiful, quiet, and happy? In: Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing
32. Shelton T, Poorthuis A, Zook M (2015) Social media and the city: Rethinking urban socio-spatial inequality using user-generated geographic information. *Landsc Urban Plan* 142:.. <https://doi.org/10.1016/j.landurbplan.2015.02.020>
33. Biljecki F, Ito K (2021) Street view imagery in urban analytics and GIS: A review. *Landsc Urban Plan* 215:..
<https://doi.org/10.1016/j.landurbplan.2021.104217>
34. Naik N, Philipoom J, Raskar R, Hidalgo C (2014) Streetscore-predicting the perceived safety of one million streetscapes. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops
35. Dubey A, Naik N, Parikh D, et al (2016) Deep learning the city: Quantifying urban perception at a global scale. In: *Lecture Notes in Computer Science*
36. Seresinhe CI, Preis T, Moat HS (2017) Using deep learning to quantify the beauty of outdoor places. *R Soc Open Sci* 4:..
<https://doi.org/10.1098/rsos.170170>
37. Zhang F, Zhou B, Ratti C, Liu Y (2019) Discovering place-informative scenes and objects using social media photos. *R Soc Open Sci* 6:.. <https://doi.org/10.1098/rsos.181375>
38. Zhu Y, Zhang Y, Biljecki F (2025) Understanding the user perspective on urban public spaces: A systematic review and opportunities for machine learning. *Cities* 156:.. <https://doi.org/10.1016/j.cities.2024.105535>
39. Chen W, Huang H, Dong J, et al (2018) Social functional mapping of urban green space using remote sensing and social sensing data. *ISPRS Journal of Photogrammetry and Remote Sensing* 146:.. <https://doi.org/10.1016/j.isprsjprs.2018.10.010>

40. Wei J, Yue W, Li M, Gao J (2022) Mapping human perception of urban landscape from street-view images: A deep-learning approach. *International Journal of Applied Earth Observation and Geoinformation* 112:. <https://doi.org/10.1016/j.jag.2022.102886>
41. Zhou B, Lapedriza A, Khosla A, et al (2018) Places: A 10 Million Image Database for Scene Recognition. *IEEE Trans Pattern Anal Mach Intell* 40:. <https://doi.org/10.1109/TPAMI.2017.2723009>
42. Anselin L (1995) Local Indicators of Spatial Association—LISA. *Geogr Anal* 27:. <https://doi.org/10.1111/j.1538-4632.1995.tb00338.x>
43. Cliff AD (Andrew D, Ord JK, Cliff AD (Andrew D (1981) *Spatial processes : models & applications*. Taylor & Francis
44. MORAN PA (1950) Notes on continuous stochastic phenomena. *Biometrika* 37:. <https://doi.org/10.1093/biomet/37.1-2.17>
45. Tobler WR (1970) A Computer Movie Simulating Urban Growth in the Detroit Region. *Econ Geogr* 46:234–240. <https://doi.org/10.2307/143141>
46. Buzzelli M (2019) Modifiable Areal Unit Problem. In: *International Encyclopedia of Human Geography, Second Edition*
47. Fotheringham AS, Wong DWS (1991) The modifiable areal unit problem in multivariate statistical analysis. *Environ Plan A* 23:. <https://doi.org/10.1068/a231025>
48. Meyer AK, Brandt T (2026) Understanding the spatial and temporal impact of global events through large-scale social media data. *EPJ Data Sci* 15:. <https://doi.org/10.1140/epjds/s13688-026-00657-y>
49. Sui D, Goodchild M (2011) The convergence of GIS and social media: Challenges for GIScience. *International Journal of Geographical Information Science* 25
50. Kitchin R, McArdle G (2016) What makes Big Data, Big Data? Exploring the ontological characteristics of 26 datasets. *Big Data Soc* 3:. <https://doi.org/10.1177/2053951716631130>
51. Kitchin R (2014) Big Data, new epistemologies and paradigm shifts. *Big Data Soc* 1:. <https://doi.org/10.1177/2053951714528481>
52. Bernabeu-Bautista Á, Serrano-Estrada L, Perez-Sanchez VR, Martí P (2021) The geography of social media data in urban areas: Representativeness and complementarity. *ISPRS Int J Geoinf* 10:. <https://doi.org/10.3390/ijgi10110747>
53. Domènech A, Mohino I, Moya-Gómez B (2020) Using flickr geotagged photos to estimate visitor trajectories in world heritage cities. *ISPRS Int J Geoinf* 9:. <https://doi.org/10.3390/ijgi9110646>
54. Niu H, Silva EA (2023) Understanding temporal and spatial patterns of urban activities across demographic groups through geotagged social media data. *Comput Environ Urban Syst* 100:. <https://doi.org/10.1016/j.compenvurbsys.2022.101934>
55. Pang B, Lee L (2008) Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2:. <https://doi.org/10.1561/1500000011>

56. Hutto CJ, Gilbert E (2014) VADER: A parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014
57. Taboada M, Brooke J, Tofiloski M, et al (2011) Lexicon-Based Methods for Sentiment Analysis. Association for Computational Linguistics 37:267–307
58. Wang Z, Huang D, Cui J, et al (2025) A review of Chinese sentiment analysis: subjects, methods, and trends. *Artif Intell Rev* 58:. <https://doi.org/10.1007/s10462-024-10988-9>
59. Liu (2012) Sentiment analysis and opinion mining. Springer
60. Glaeser EL, Kominers SD, Luca M, Naik N (2018) BIG DATA AND BIG CITIES: THE PROMISES AND LIMITATIONS OF IMPROVED MEASURES OF URBAN LIFE. *Econ Inq* 56:. <https://doi.org/10.1111/ecin.12364>
61. Lynch K (1960) The image of the city. MIT Press
62. Lefebvre H (1991) The production of space. Wiley-Blackwell
63. Whitt JA, Zukin S (1996) The Cultures of Cities. *Contemp Sociol* 25:. <https://doi.org/10.2307/2077290>
64. Zhang G, Zhu AX (2018) The representativeness and spatial bias of volunteered geographic information: a review. *Ann. GIS* 24
65. Graham M, Hale SA, Gaffney D (2014) Where in the World Are You? Geolocation and Language Identification in Twitter. *Professional Geographer* 66:. <https://doi.org/10.1080/00330124.2014.907699>
66. Zhu X, Gardiner S, Roldán T, Rossouw D (2024) The Model Arena for Cross-lingual Sentiment Analysis: A Comparative Study in the Era of Large Language Models. In: WASSA 2024 - 14th Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media Analysis, Proceedings of the Workshop
67. He, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition
68. Deng J, Dong W, Socher R, et al (2009) ImageNet: A Large-Scale Hierarchical Image Database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009

Declarations

Ethics approval and consent to participate

The dataset contains human-generated social media content, but the study did not recruit or interact with human participants and involved no survey, intervention or experiment. The posts were publicly accessible at

the time of collection. No prospective institutional ethics approval or individual informed consent was obtained; the authors treated the study as secondary observational research using public online content. Direct account identifiers, usernames and profile fields were excluded from the analytical dataset before analysis. Images were used only to derive broad category labels and are not reproduced or publicly shared. All reported spatial results are aggregated to 500 m or 1000 m grid cells so that no individual user or post is identifiable. All methods were carried out in accordance with applicable regulations and institutional guidance concerning privacy and the research use of publicly accessible online data.

Consent for publication

Not applicable.

Availability of data and materials

The raw Sina Weibo posts are third-party user-generated content containing text, images and precise geolocation. They cannot be redistributed because of privacy considerations and restrictions on the redistribution of platform content. The posts were publicly accessible at the time of collection, but there is no stable public URL for the complete study dataset. A de-identified, aggregated grid-level dataset containing the variables necessary to reproduce the reported statistical analyses, together with the analysis scripts, is available from the corresponding authors upon reasonable request for research and verification purposes, subject to a data-use agreement prohibiting attempts to re-identify users and the redistribution of the original platform content.

Competing interests

The authors declare that they have no competing interests.

Funding

1. This research is funded by the Scientific Research Project of Guangzhou Huashang College (Project No.: 2023HSKT01).
2. Guangzhou Huashang College's School-level Project "Professional Investigation" Integrated Theory-Practice Course" (HSRHKC2024015).

Authors' contributions

Huimin Qu conceived the study, collected and preprocessed the Sina Weibo dataset, developed the multimodal analytical framework, implemented the sentiment analysis, image classification, visibility bias modelling, null-model simulations, robustness analysis, and spatial statistical analysis, interpreted the results, prepared all figures and tables, and drafted the manuscript.

Gongxiang Huang contributed to the conceptualization of the study, supervised the research, critically reviewed the methodology and discussion, and substantially revised the manuscript.

Zhuqin Liang contributed to the statistical modelling, methodological development, interpretation of the spatial statistical analyses, and critical revision of the manuscript.

Pei Wen contributed to data interpretation, visualization design, manuscript revision, and language editing.

All authors read and approved the final manuscript.