Missing women on Indian streets*

Varun Karekurve-Ramachandra† Gaurav Sood‡

April 21, 2026

Abstract

How absent are women from city streets in the developing world? We answer this question using GPS-linked wearable cameras and randomized street audits across ∼900 kilometers of roads in greater Mumbai. Across 4000+ street images containing 23,000+ visible person observations, women account for 16.4% of visible people in Mumbai and 14.7% in Navi Mumbai, far below their population shares. We estimate pedestrian sex ratios of 239 and 223 women per 1,000 men, implying 71% and 76% of women expected based on residential ratios are *missing* from the streets. This pattern holds across road types, and private mobility does not explain the gap; women's share on two-wheelers is lower still (8.4% and 5.7%). These results provide the first large-scale measurement of gender disparities in urban public life that self-reported data cannot capture.

One-sentence summary: This paper introduces a scalable, image-based method to measure the gender composition of urban public space and shows that roughly three-quarters of the women expected on the streets of greater Mumbai are missing revealing a large, previously unmeasured spatial dimension of gender inequality.

*The authors thank the Center for International Studies (CIS) at USC for generous funding support. We also thank Ronit Shah and Swathi Srinivas for their research assistance. Replication materials for the project can be found at `https://github.com/geosensing/streetscope`. This study was reviewed and deemed exempt by the University of Southern California Institutional Review Board (UP-25-00545)

†Assistant Professor, Department of Political Science and International Relations, University of Southern California, Email: varun.rama@usc.edu

‡Independent Researcher, Email: gsood07@gmail.com

1 Introduction

How absent are women from city streets in the developing world? In the Mumbai Metropolitan Region, home to 25 million people, women account for just 16.4% of visible people on the streets of Mumbai and 14.7% in Navi Mumbai. Pedestrian sex ratios – 239 and 223 women per 1,000 men – are roughly one-quarter of residential census ratios. Counting absence has scholarly precedent. In 1990, Amartya Sen estimated that more than 100 million women were "missing" from the world's population due to excess female mortality, concentrated in South Asia, West Asia, and China (Sen, 1990). This accounting transformed how researchers measured gender inequality: not through attitudes or laws, but by counting who is absent. We extend that logic from population registers to city streets, focusing on who is physically present in public space, and provide the first systematic measurement of women's presence in urban public space at the population scale.

The question is both normative and empirical: presence in public space is a dimension of well-being, yet self-reported data cannot accurately capture who is physically present on city streets. In the capabilities framework of Sen (1999) and Nussbaum (2000), freedom of movement is central to bodily integrity – one of Nussbaum's central human functional capabilities – and conditions access to employment, education, political participation, and health services. When women cannot move freely in a city, the consequences can ripple through labor markets, civic life, and household bargaining. Urban streets are where these constraints become physically visible: a space can be crowded and still exclusionary if the crowd is overwhelmingly male. Feminist urban scholarship has advanced this argument (Phadke et al., 2011; Kern, 2020), but the empirical base has not kept pace with the theory.

Existing evidence on women's presence in public space comes from three sources, none adequate for city-scale estimation. Safety audits record gendered perceptions and headcounts at purposively selected locations but are small in scale and not designed for population inference (UN Habitat, 2008; Jagori, 2010; Safetipin, 2015; UN Women and Aurat Foundation, 2021). Commuter surveys and census journey-to-work tables capture transport mode choices but miss non-commute trips and say nothing about who is visible on the street at any given moment (Office of the Registrar General and Census Commissioner, India, 2011). Time-use diaries record how women allocate hours but lack spatial resolution (National Sample Survey Office, 2025). The result is a measurement gap: for no major city in the developing world do we have a representative estimate of women's share of the street population. Rapid urbanization, smart-city investments, and policy commitments under SDG 5 (gender equality) and SDG 11 (inclusive cities) have generated demand for precisely this kind of baseline; still, the data do not yet exist (United Nations, 2016).

In this study, we report results from a randomized street survey in two cities of the Mumbai Metropolitan Region in western India: Mumbai, the historic city, and Navi Mumbai, a planned satellite city. Our estimand is the share of women visible on public streets at the time of observation. We sampled road segments from the complete street network using a stratified random design across road types. A surveyor traversed approximately 900 kilometers of roads on a motorcycle, wearing a helmet-mounted, GPS-linked camera. Trained human annotators classified over 23,000 visible individual observations in more than 4,000 images, recording gender, transport mode, and streetscape features such as vendor presence, street markings, and road surface quality. Women account for 16.4% of visible people in Mumbai and 14.7% in Navi Mumbai. The disparity is not confined to a single road type: women's share is low across all road types. Nor is it explained by women shifting from walking to private mobility: women's share among two-wheeler riders is lower still. The paper provides an empirical baseline for two major Indian cities and introduces a scalable, replicable image-based method for measuring gender composition in public space at the city scale. Our findings reveal a large, previously unmeasured spatial dimension of gender inequality.

2 Results

2.1 Women in the street population

In both Mumbai and Navi Mumbai, women constitute a small minority of the visible street population (Table 1). In person-weighted estimates (each visible individual counts equally), women account for 16.4% of visible people in Mumbai and 14.7% in Navi Mumbai. Image-level means, which weight each frame equally regardless of crowd size, are 15.6% and 15.5%, respectively. Restricting the count to pedestrians, women constitute 19.3% in Mumbai and 18.2% in Navi Mumbai, corresponding to pedestrian sex ratios of 239 and 223 women per 1,000 men – far below the residential census ratios of 838 and 910. Echoing Sen (1990)'s accounting at the population scale, the streets of Mumbai are missing 71% of the women that residential ratios would predict; Navi Mumbai is missing 76%. This gap between the street sex ratio and the residential sex ratio is the core empirical finding: it is large, consistent across both cities, and not close to parity by any measure.

Table 1: Women's share of the visible street population by city.

	Mumbai	Navi Mumbai
Images annotated	2,251	1,239
Total people classified	16,908	6,855
Prop. female (person-weighted)	0.164	0.147
Prop. female (image-level mean)	0.156	0.155
Sex ratio (F/1000 M)	196	172
Pedestrian prop. female	0.193	0.182
Pedestrian sex ratio (F/1000 M)	239	222

Note: Person-weighted estimates weight each individual equally; image-level means weight each frame equally regardless of crowd size. Sex ratio is women per 1,000 men. Gender inferred from visible appearance.

The image-level distribution of the female share reinforces the picture (Fig. 1). In both cities, the distribution is right-skewed, with a pronounced mass at zero: a substantial fraction of frames contain no visible women at all. The modal image is one in which women are present but a clear minority. Frames in which women approach or exceed parity are rare. The pattern

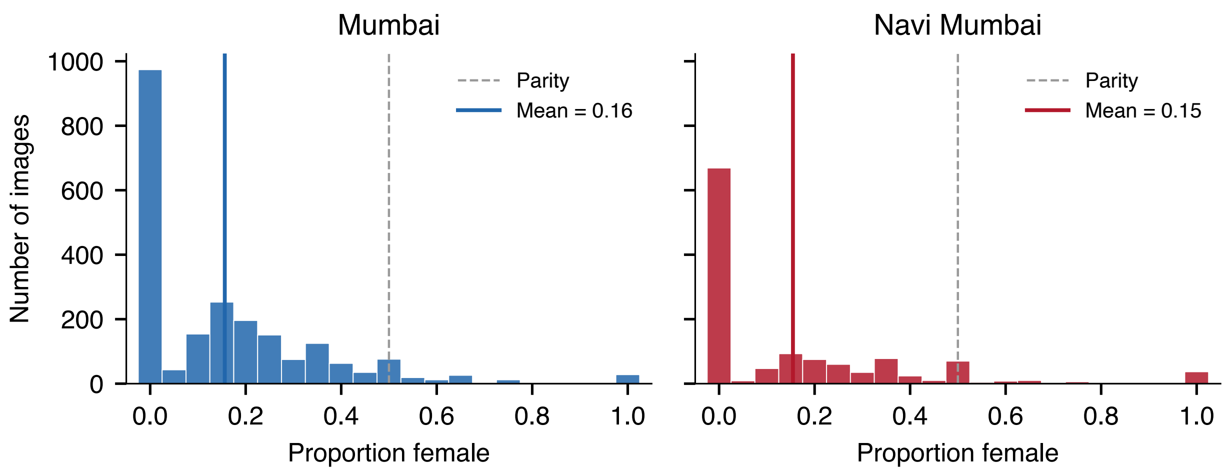

Figure 1: Distribution of image-level female share in Mumbai and Navi Mumbai. Each bar represents the fraction of annotated images in which women constitute a given share of visible people. The distribution is right-skewed in both cities, with a pronounced mass at zero: a substantial fraction of frames contain no visible women. Frames that approach or exceed gender parity are rare. Vertical solid lines indicate the image-level mean; dashed lines indicate parity (0.50).

is not driven by a handful of male-dominated outlier images; it is the typical condition of the street.

2.2 Temporal patterns

Data collection is concentrated during daytime hours (approximately 07:00–19:00 IST), so these estimates should not be read as a full diurnal profile – a full 24-hour pattern. Within that window, women's share in Mumbai is modestly higher on weekdays than weekends (16.7% versus 14.8%; $p = 0.04$), possibly reflecting commuting patterns. Navi Mumbai shows the opposite pattern (14.6% versus 15.7%), but the difference is not statistically significant (Fig. 3). Variation across the three broad time periods (morning 07:00–11:00, midday 11:00–15:00, afternoon 15:00–19:00) is modest in both cities, with no time window approaching gender parity. If anything, prior qualitative work would predict nighttime visibility to be even more male-skewed (Phadke et al., 2011), so our daytime estimates are likely an upper bound on women's share across the full 24-hour cycle.

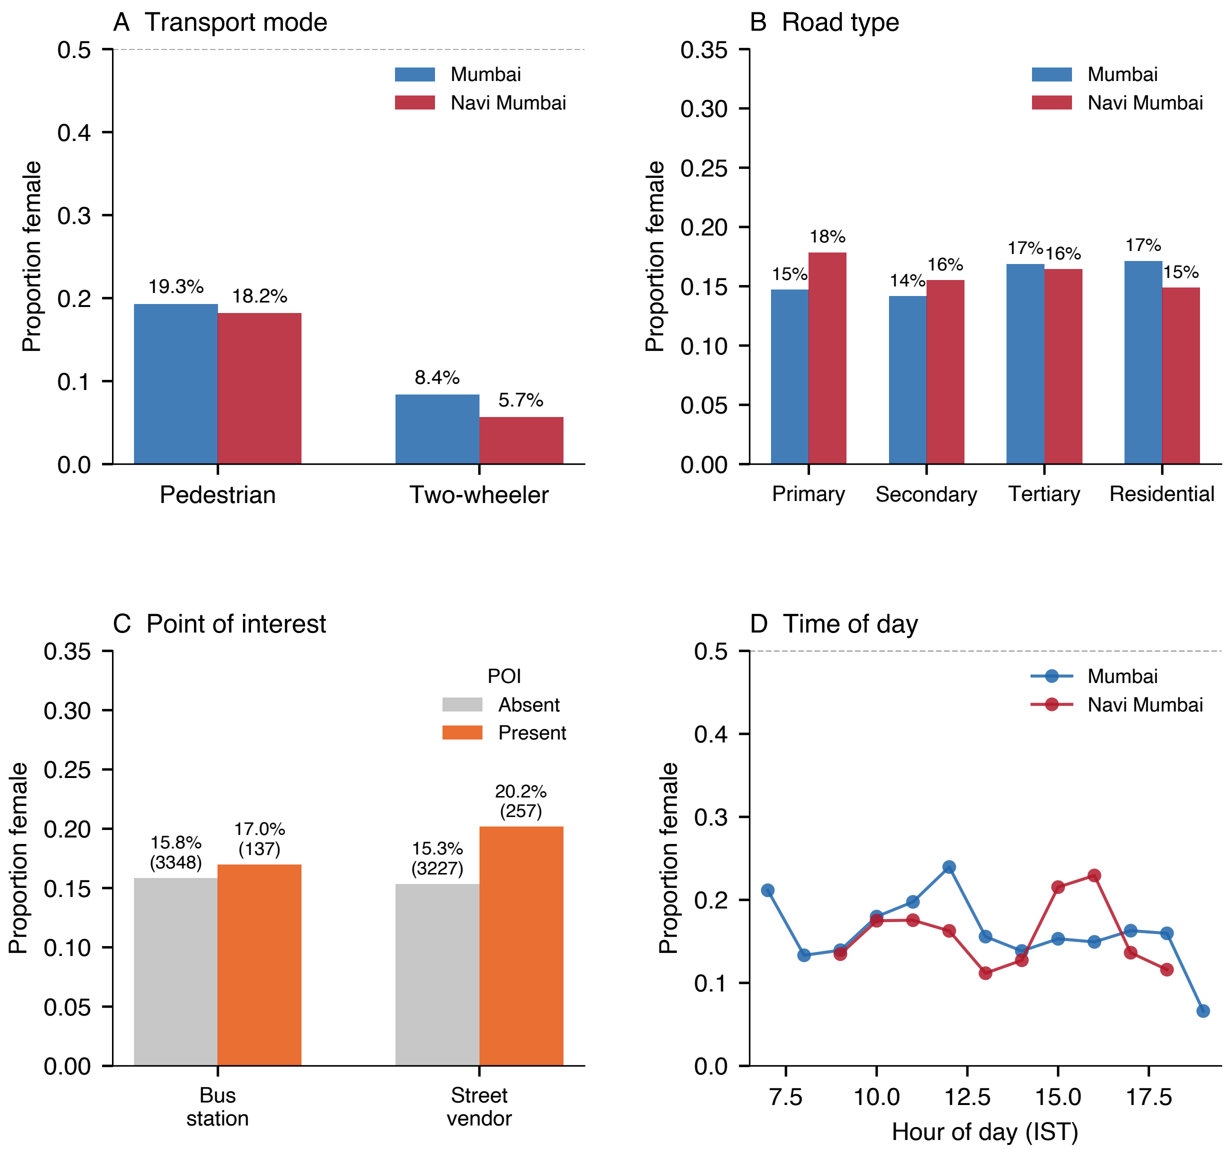

Figure 2: Women's share of visible people by transport mode, road type, point of interest, and time of day.

Notes: (**Panel A**) Women's share is lower among two-wheeler riders than among pedestrians in both cities, inconsistent with modal substitution as an explanation for low female street presence. (**Panel B**) In Navi Mumbai, women's share is highest on primary roads (18%) and declines toward residential streets (15
(**Panel C**) Locations tagged with a street vendor or bus station show modestly higher female shares than locations without; sample sizes in parentheses. Both cities pooled. (**Panel D**) Women's share by hour of day (IST). No time window approaches parity. Coverage is concentrated between 07:00 and 19:00; points with fewer than five images excluded.

2.3 Modal substitution

A natural objection is that women may be present in public spaces but traveling in enclosed vehicles rather than on foot. We cannot observe car occupants, but cars account for a small share of Mumbai's vehicle fleet, and the two-wheeler (motorbikes, scooters, and bicycles) is by far the dominant private mode. Two-wheelers are fully visible in the imagery. If modal substitution explained the pedestrian gap, women's share among two-wheeler users should be higher than among pedestrians. It is not. Women constitute 8.4% of two-wheeler users in Mumbai and 5.7% in Navi Mumbai, substantially lower than their pedestrian shares (Fig. 2A). The most visible private transport mode is even more male-dominated than walking. We cannot rule out substitution into enclosed public transport like buses and Mumbai's suburban rail, whose interiors are not visible from the street; we address this in ongoing work. Whatever forces suppress women's street presence, they are not offset by a shift into two-wheeled private mobility.

2.4 Points of interest

Do women appear more frequently in locations with specific functions? Bus-adjacent scenes show a modestly higher female share than non-bus scenes (17.0% versus 15.8%; $p = 0.048$), consistent with the idea that transit access points concentrate purposeful female trips. Street-vendor locations show a more pronounced difference (20.2% versus 15.3%; $p < 0.001$), suggesting that informal-commerce clusters draw somewhat more gender-mixed foot traffic (Fig. 2 Panel C). Railway-station estimates are based on only eight tagged images and are too sparse to interpret. The POI patterns are suggestive rather than causal, but they point toward a plausible mechanism: locations that concentrate practical errands or mixed-use activity may produce somewhat less male-skewed street populations than purely through-moving road segments.

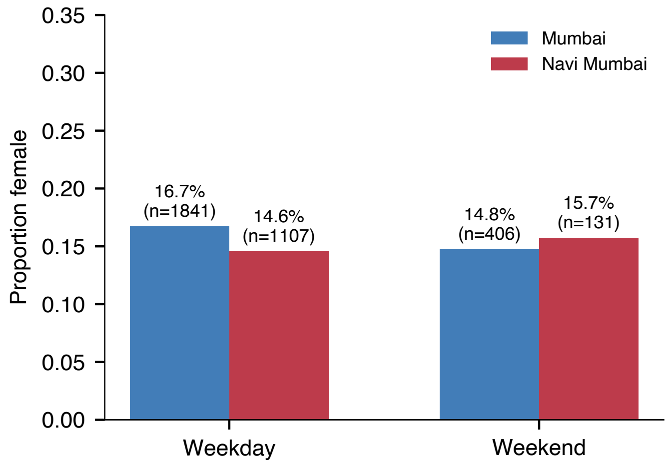

Figure 3: Women's share of visible people on weekdays versus weekends

Note: Mumbai shows a modestly higher female share on weekdays (person-weighted; $p = 0.048$), possibly reflecting gendered commuting patterns. Navi Mumbai shows the reverse but without statistical significance. Sample sizes in parentheses.

2.5 Road hierarchy

The spatial gradient differs strikingly between the two cities (Fig. 2 Panel B). In Mumbai, women's presence is depressed and roughly flat across the road hierarchy: 14.7% on primary roads, 14.1% on secondary, 16.9% on tertiary, and 17.1% on residential. No road class approaches parity. Navi Mumbai shows the opposite ordering: women's share is highest on primary roads (17.9%) and declines in residential streets (14.9%) with secondary (15.5%) and tertiary roads (16.4%) in between (Table 2). The pattern is consistent with the different urban fabric of the two cities: In Mumbai, relatively more women are seen in residential roads; in suburban Navi Mumbai, relatively more women are visible on primary roads.

3 Discussion

These results provide, to our knowledge, the first systematic city-scale estimate of women's share of the visible street population in a major city in the developing world. The headline finding is stark: women account for just 16.4% of the people visible on Mumbai's streets and 14.7% in Navi Mumbai — implying that roughly 71% and 76% of the women expected based on residential ratios are missing. A pilot study in Delhi using a similar design found comparable magnitudesThe numbers convert a widely reported observation into a precise, replicable quantity, and confirm at the city scale what small-sample safety audits and qualitative fieldwork have long suggested (Jagori, 2010; Phadke et al., 2011).

Several mechanisms could produce this pattern, and our data cannot adjudicate among them. Cultural and societal norms that treat unaccompanied female presence in public as suspicious or improper are well documented in Indian cities (Phadke et al., 2011). Fear of harassment constrains women's route choices, timing, and willingness to travel alone; in Delhi, female college students accept longer commutes to avoid streets perceived as unsafe (Borker, 2021). Infrastructural gaps – inadequate public sanitation, poor lighting, missing or obstructed sidewalks – may raise the cost of public presence disproportionately for women. Moreover, the unequal distribution of household labor limits discretionary travel: women who spend more hours on domestic work have fewer hours available for the street. These forces likely interact. The flat spatial profile in Mumbai, where women's share is depressed across every road class, is consistent with constraints that operate city-wide rather than at the level of individual corridors.

The POI results offer a partial window into which conditions correlate with higher female presence. Women's share is elevated near street vendors and, more modestly, near bus stations – locations anchored by practical errands and mixed-use activity. One interpretation is that women enter public space more readily where use is legible and purposeful rather than exposed and discretionary. This is consistent with Phadke et al. (2011)'s observation that women in Mumbai are permitted to traverse public space for work or errands but not to linger without purpose. The Navi Mumbai road-type gradient points in a similar direction: primary roads, which carry transit and commercial activity, show the highest female share. Planned residential roads are the most male-dominated segments. These patterns do not establish causation, but they suggest that the spatial organization of activity matters for who can occupy public space.

The findings have direct implications for urban policy. Gender-responsive planning – better lighting, continuous sidewalks, accessible sanitation, mixed land use – is frequently advocated but rarely evaluated against a quantitative baseline of gendered street presence. The methodology demonstrated here provides such a baseline at low cost. Repeated transects after infrastructure upgrades, transit interventions, or safety measures can test whether the visible share of women rises. If it does not, the intervention has failed on its own terms. The framework is also relevant to monitoring progress on SDG 5 (gender equality) and SDG 11 (inclusive cities), both of which lack street-level indicators (United Nations, 2016).

Several limitations should be stated plainly. First, data collection is concentrated during daytime hours. Nighttime presence, when qualitative evidence suggests the gender skew is most severe, remains unmeasured. Second, gender is inferred from visible appearance, collapsing a more complex social reality into binary categories. Third, the study covers two cities in a single cross-section and cannot track change over time. Fourth, we cannot distinguish residents from commuters or determine trip purpose from a single image. Fifth, the associations we report – between women's presence and road type, POI, or time of day – are observational. They describe where women are more or less visible, not why.

The pipeline is designed for replication. The sampling and data-collection tools are open-source (Sood and Laohaprapanon, 2018, 2017), and the marginal cost of extending the survey to additional cities or time periods is low. The same protocol can be deployed in Dhaka, Lagos, Jakarta or any other city in the developing world with minimal adaptation. Longitudinal repetition would allow cities to track whether women's street presence responds to policy changes or remains structurally fixed. Future work should extend coverage to nighttime hours and enclosed public spaces, incorporate qualitative methods to identify mechanisms, and explore whether experimental interventions — street redesign, safety measures, collective claims to public space — shift the baseline

The larger point is simple. A city in which 84% of the people visible on the street are men is not close to equal access. The imbalance is not confined to one road type, one transport mode, or one time of day. Measuring that absence is the first step toward changing it.

4 Materials and Methods

4.1 Sampling design

We sampled road segments from the complete municipal street networks of Mumbai and Navi Mumbai. Road geometries were downloaded from OpenStreetMap and classified into four strata by road type: primary, secondary, tertiary, and residential. Each road was divided into approximately 500-meter segments using the open-source `geo_sampling` Python package (Sood and Laohaprapanon, 2018; Laohaprapanon et al., 2018). We then drew a stratified random sample of 1,000 segments per city, with stratum sizes proportional to the number of segments in each road class. In Mumbai, the resulting sample comprised 122 primary, 150 secondary, 127 tertiary, and 601 residential segments, totaling approximately 900 kilometers of road. Sampled segments were grouped into daily surveyor itineraries using `allocator` (Sood and Laohaprapanon, 2017), a route-optimization tool that minimizes travel time between segments while respecting daily time constraints. Mumbai's sample was organized into 12 itineraries; Navi Mumbai's into 41.

4.2 Data collection

The surveyor traversed the assigned itineraries on a motorbike equipped with a chest-mounted GoPro HERO13 Black camera. The camera recorded continuous 4K video (3840×2160, 30 fps, HEVC codec) with embedded GPS telemetry at approximately 10 Hz. Data collection spanned March 20 to May 15, 2025, across 25 collection days, yielding 174 video files totaling approximately 43 hours of footage and 805 GB of raw data. Collection occurred primarily during daytime hours (approximately 07:00–19:00 IST), with coverage concentrated in the morning and midday periods. Thursdays are overrepresented at roughly 40% of sessions owing to surveyor scheduling constraints.

We extracted still frames from each video at fixed 10-second intervals using a custom pipeline (`extract_frames.py`), yielding one image approximately every 30–50 meters depending on biking speed. GPS coordinates were assigned to each frame by interpolating the embedded telemetry timeseries (`extract_gps_timeseries.py`). Frames were compressed to 1280×720 resolution for annotation. The pipeline from raw video to geolocated, annotation-ready images is fully documented in the open-source `streetscope` repository (Karekurve-Ramachandra and Sood, 2025).

4.3 Pedestrian classification

Extracted frames were loaded into Label Studio, an open-source annotation platform (Tkachenko et al., 2025). Three trained coders annotated the images. For each image, coders recorded: (i) counts of visible men and women, (ii) counts of men and women riding two-wheelers, and (iii) categorical indicators for points of interest (bus station, railway station, street vendor) and street infrastructure (footpath condition, lane markings, potholes, litter). Count fields used a categorical taxonomy (1, 2, 3, . . . , 10, >10), with >10 coded conservatively as 11 in the analysis. Where no women (or no men) were visible, the coder left the corresponding count field blank; blank entries are treated as zero.

The Mumbai images were annotated by all three coders, and the Navi Mumbai images were annotated by a single coder.Each image was coded once. Median annotation time per image was approximately 30 seconds, and the total annotation effort was approximately 225 person-hours. Gender was inferred from visible appearance – clothing, hair, and build as they present in street-level photographs. This necessarily imposes a binary classification and cannot capture the full complexity of gender identity. The final dataset comprises 2,721 images in Mumbai and 1,543 in Navi Mumbai, yielding classifications of over 23,000 individual observations across 4000+ images in which at least one person is visible.

We did not conduct a formal inter-rater reliability study with multiply-coded images. This is a limitation. However, the classification task is relatively constrained – coders count visible people and categorize them by apparent gender, not make subjective judgments about behavior or intent. As a partial check, we compared city-level estimates produced by the coder who annotated both cities against the pooled estimates from all three coders in Mumbai and found no substantive divergence.

4.4 Statistical analysis

We report two summary measures of women's street presence. The *person-weighted* estimate treats each classified individual equally: proportion female $= \sum w_i / \sum(m_i + w_i)$, where w_i and m_i are the counts of women and men in image i. The *image-level mean* treats each frame equally regardless of crowd size: $\bar{p} = N^{-1} \sum_i (w_i/(m_i + w_i))$, computed only over frames with at least one visible person. The person-weighted estimate is more informative about the composition of the street population; the image-level mean is more robust to a small number of high-traffic frames.

Subgroup comparisons (by road type, POI presence, weekday/weekend) use person-weighted proportions within each subgroup. Statistical significance for POI comparisons is assessed using two-proportion z-tests on the person-weighted female shares. We do not adjust for multiple comparisons; the POI and temporal results are presented as descriptive patterns rather than hypothesis tests.

Pedestrian counts are computed by subtracting two-wheeler counts from total counts: pedestrian women $= w_i - w_i^{\text{tw}}$, pedestrian men $= m_i - m_i^{\text{tw}}$. The pedestrian sex ratio is expressed as women per 1,000 men, following demographic convention.

Road type for each image is assigned from the OpenStreetMap classification of the sampled segment, not from image-level annotation. This avoids post-hoc selection by the coder but assumes the surveyor was on the assigned segment at the time the frame was captured – a reasonable assumption given GPS tracking.

4.5 Ethics and privacy

The study involves observation of public behavior in public spaces and does not collect personally identifiable information. No individual consent was obtained, consistent with standard practice for street-level observation research. Images were used solely for aggregate counting of visible people by apparent gender; no facial recognition was applied, no individual-level data were retained, and no images of identifiable individuals were published. The annotation interface displayed compressed, reduced-resolution images. All raw video and original-resolution frames are stored securely and are not included in the public repository. The repository contains only aggregate annotation counts, GPS tracks, and analysis code. This study was reviewed and deemed exempt by the author's University Institutional Review Board.

References

Borker, G. (2021). Safety first: Perceived risk of street harassment and educational choices of women. *American Economic Review*, 111(6):forthcoming.

Jagori (2010). Understanding women's safety: Towards a gender inclusive city: Research findings, delhi, 2009–10. Technical report, Jagori, New Delhi.

Karekurve-Ramachandra, V. and Sood, G. (2025). Streetscope: Measuring street-level conditions from wearable camera transects. https://github.com/geosensing/streetscope. Replication data, Label Studio annotation exports, GPS telemetry, and analysis code for Mumbai and Navi Mumbai.

Kern, L. (2020). *Feminist City: Claiming Space in a Man-Made World*. Verso, London.

Laohaprapanon, S., Ortleb, K., and Sood, G. (2018). Street sense: Learning from Google Street View. *arXiv preprint*.

National Sample Survey Office (2025). Time use survey (TUS), january 2024–december 2024. Microdata catalog DDI-IND-NSO-TUS-2024-24. Accessed: 2026-04-13.

Nussbaum, M. C. (2000). *Women and Human Development: The Capabilities Approach*. The Seeley Lectures. Cambridge University Press, Cambridge.

Office of the Registrar General and Census Commissioner, India (2011). Census of India 2011: Tables on workers and non-workers. https://censusindia.gov.in.

Phadke, S., Khan, S., and Ranade, S. (2011). *Why Loiter? Women and Risk on Mumbai Streets*. Penguin Books, New Delhi.

Safetipin (2015). My Safetipin: Making cities safer for women. https://safetipin.com. Mobile application and safety audit platform.

Sen, A. (1990). More than 100 million women are missing. *The New York Review of Books*, 37(20):61–66.

Sen, A. (1999). *Development as Freedom*. Alfred A. Knopf, New York.

Sood, G. and Laohaprapanon, S. (2017). allocator: Optimally allocate geographically distributed tasks. Version 0.1.0. Available at https://github.com/geosensing/allocator/.

Sood, G. and Laohaprapanon, S. (2018). Geo sampling: Randomly sample locations on streets. Version 0.1.0. Available at https://github.com/geosensing/geo-sampling.

Tkachenko, M., Malyuk, M., Holmanyuk, A., and Liubimov, N. (2020-2025). Label Studio: Data labeling software. Open source software available from https://github.com/HumanSignal/label-studio.

UN Habitat (2008). Womens safety audits what works and where?

UN Women and Aurat Foundation (2021). Women's safety audit in public places: Karachi, Dadu, Khairpur, Quetta and Rawalpindi. Accessed: 2026-04-15.

United Nations (2016). New urban agenda. Technical report, United Nations, New York. A/RES/71/256. Adopted at Habitat III, Quito, Ecuador, 20 October 2016; endorsed by the United Nations General Assembly, 23 December 2016.

Supplementary Materials

S1. Annotation Codebook

The following rules governed the coding of all images. The codebook was developed iteratively during a pilot phase and finalized before production annotation began. All coders were trained on the same instructions and reviewed edge cases collectively before independent annotation.

S1.1 Adult men and women

Coders counted all adults visible on the street whose gender could be determined from the image.

- People inside buildings or enclosed vehicles were **not** counted.
- People in public parks set back from the road were **not** counted.
- People on open vehicles such as trucks or auto-rickshaws were counted if clearly visible. Truck occupants identified as laborers were noted in the free-text field.
- People near the horizon whose gender was not discernible were excluded.
- An image in which the street was visible but no people were present was coded as zero for all person counts (equivalent to leaving count fields blank).

S1.2 Men and women on two-wheelers

Two-wheelers included both motorized bikes and bicycles.

- Only riders and pillion passengers whose gender was clearly discernible were counted.
- When a pillion passenger's gender could not be determined, only the rider was counted as a single person on a two-wheeler.
- A person on a bicycle was classified as a two-wheeler user.

S1.3 Footpath condition

The guiding question was: *Is this footpath usable by a pedestrian, including a person with a mobility disability?*

- **Paved**: A paved footpath visible in any part of the image, with sufficient clearance for pedestrian use.
- **Paved–Blocked**: A footpath that exists but is obstructed, reducing or eliminating pedestrian access. Sources of obstruction included:
  - Unpaved surfaces with trees and mud that pedestrians would likely avoid (concern: mud, risk of waterlogging);
  - Half-constructed or incomplete footpaths;
  - Street vendors or hawkers occupying the footpath (also tagged as `street_vendor = Yes`);
  - Parked motorcycles or bicycles on the footpath.
- **No sidewalk**: No footpath visible in the image.
- A tree on an otherwise usable, unobstructed footpath was **not** treated as a blockage. This was a common occurrence in both cities and was noted as a separate convention.

S1.4 Potholes

- A visible dent, hole, or hole-like depression in the road surface was coded as a pothole.
- Minor surface cracks were **not** coded as potholes.
- Drainage grates or grills at the road edge were **not** coded as potholes.
- Coded as NA when no paved (*pucca*) road was visible (e.g., unpaved or dirt roads).

S1.5 Street markings

The guiding question was: *Are there visible markings that help clarify what space on the road vehicles and pedestrians can use?* The underlying concern was traffic safety and pedestrian safety: markings such as zebra crossings, lane dividers, and center lines cue both vehicles and pedestrians to share space. Markings that function primarily at night (e.g., reflective paint visible only under headlights) were acknowledged as likely to be systematically undercounted in daytime imagery.

- Even barely visible paint or reflectors were coded as street markings present.
- A painted or marked speed breaker counted as a street marking. In Mumbai, both the road surface and speed breakers typically had markings; in Navi Mumbai this was less consistent.
- Gaps between concrete slabs serving as separators, without paint or other explicit marking tools, were coded as no street markings.
- Coded as NA when no paved (*pucca*) road was visible.

S1.6 Litter and construction debris

The guiding question was: *Is material left lying in an open or public place that should be in a bin or container?*

- Visible garbage on or near the street was coded as litter.
- Construction debris, gravel, or loose stone were coded as construction debris.
- When both garbage and construction debris were present, the dominant category by visual volume was selected.
- Unfinished street corners and unpaved footpath edges with loose material were coded as construction debris.

S1.7 Points of interest

- **Bus station**: Coded as Yes when a bus stop, bus shelter, or station was visible. When only partially visible or at a distance, a free-text note was added.
- **Railway station**: Coded as Yes when a railway station or its immediate vicinity was visible. The presence of railway porters (*coolies*) was used as a contextual indicator when the station structure itself was not in the frame.

- **Street vendor**: Coded as Yes when a street vendor or hawker was visible, including cases where the vendor occupied the footpath (which was then also coded as Paved–Blocked under footpath condition).

S1.8 Neighborhood type and road width

These fields were added during the annotation process and are available for a subset of images.

- **Neighborhood type**: Coded as residential, commercial, or mixed-use. A residential street with even a single visible commercial establishment (e.g., a shuttered shop) was coded as mixed-use. The field was left blank when no view of the surrounding built environment was available.
- **Road width**: Coded as single-lane or two-lane. A single-direction road with a wide carriageway was coded as two-lane. A nominally two-lane road with vehicles parked on both sides, reducing the usable carriageway, was coded as single-lane.

S1.9 Repeated images

Sequential frames extracted at 10-second intervals from the same video occasionally captured substantially the same scene (e.g., when the surveyor was stopped at a traffic signal). Images identified as repeated by a coder were flagged using the `repeat_option` field and excluded from the analysis.

Supplementary Tables and Figures

Table 2: Women's share of visible people by road type.

City	Road type	Prop. female	Sex ratio (F/1000 M)
Mumbai	Primary	0.147	173
	Secondary	0.141	165
	Tertiary	0.169	203
	Residential	0.171	207
Navi Mumbai	Primary	0.179	217
	Secondary	0.155	184
	Tertiary	0.164	197
	Residential	0.149	175

Note: Road type assigned from itinerary classification. Person-weighted estimates.

Table 3: Points of interest and infrastructure summaries.

Characteristic	Present	Prop. female	Images
Bus station	Yes	0.170	137
Bus station	No	0.158	3,348
Railway station	Yes	0.149	8
Railway station	No	0.159	3,477
Street vendor	Yes	0.202	257
Street vendor	No	0.153	3,227
Infrastructure coverage (all cities pooled)			
Footpath	No	–	3,490 (100%)
Potholes	Yes	–	40 (1%)
Potholes	No	–	3,450 (99%)
Litter	Yes	–	443 (13%)
Litter	No	–	3,047 (87%)

Note: Infrastructure fields are sparse; not all images have all fields coded.

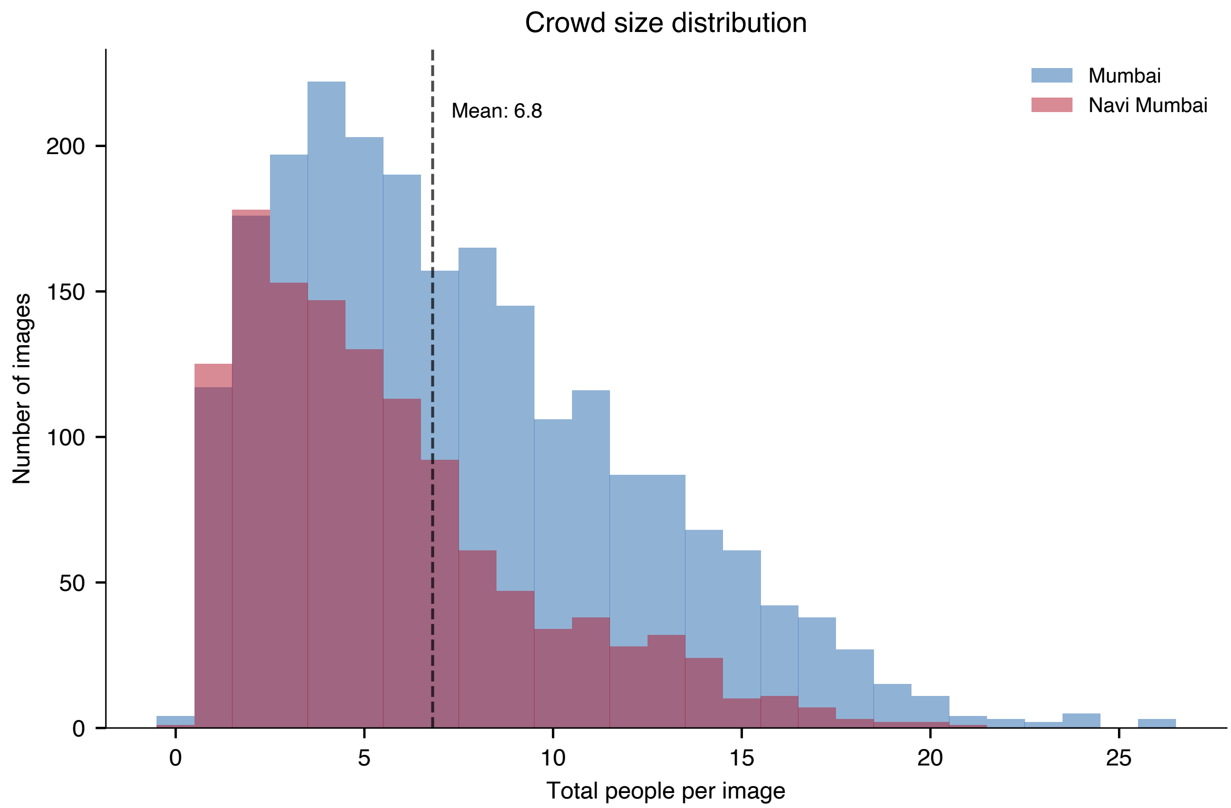

Figure 4: Crowd Size Distribution

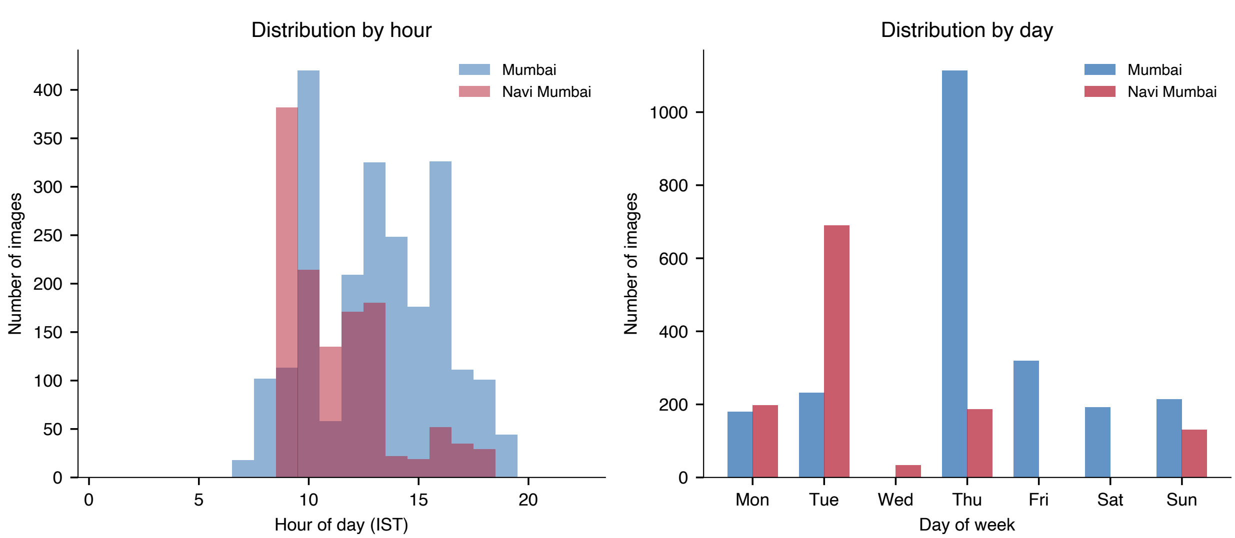

Figure 5: Temporal Coverage

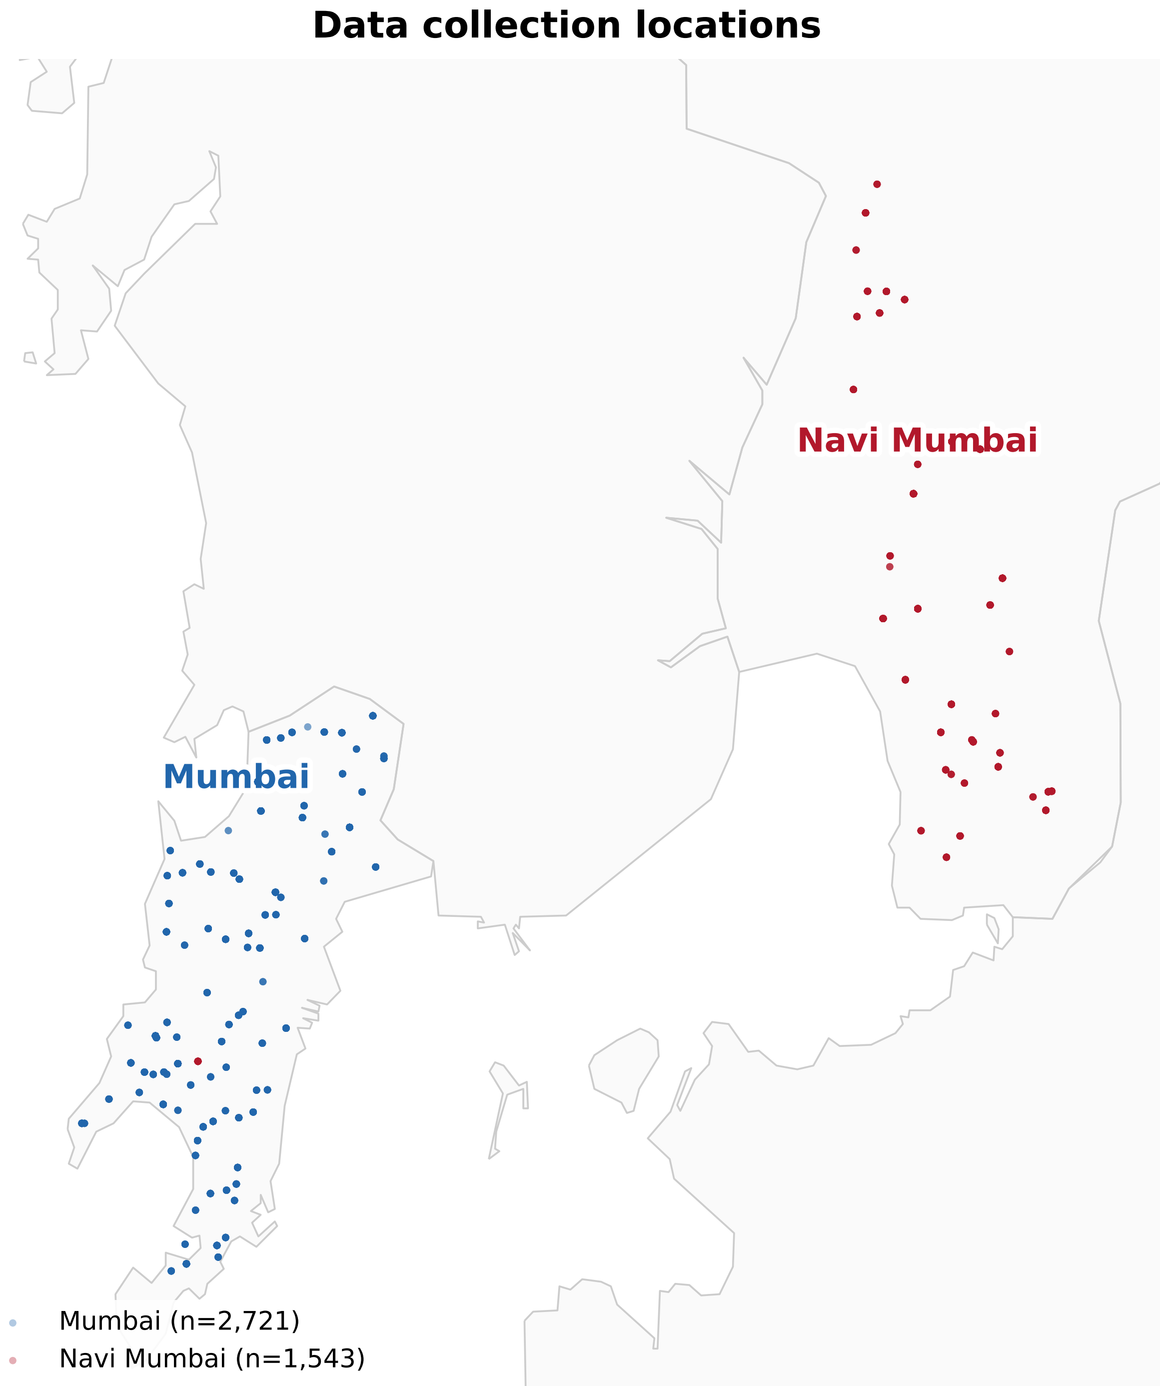

Figure 6: Mumbai and Navi Mumbai