Towards Causal Understanding of Urban Air Pollution: Mechanistic Models under Sparse Sensing

Anonymous Author(s)
Affiliation
Address
email

Abstract

Understanding the causal drivers of urban air pollution remains a central challenge for environmental science and policy. While high-resolution source apportionment typically relies on dense monitoring networks or receptor-based chemical analysis, many cities must operate with sparse sensors and incomplete emission inventories. We frame air pollution source attribution as a causal inference problem, linking emissions to observed concentrations through mechanistic dispersion models. Using Gaussian plume formulations, we combine multiple emission categories—vehicular traffic, domestic emissions, brick kilns, industries, and power plants—with real-world sensor data from New Delhi. Our methodology estimates source-specific contributions under sparse observations via parameterized dispersion modeling, while also capturing the influence of missing or unobserved sources. By situating source apportionment within a causal modeling perspective, we emphasize both the opportunities and limitations of mechanistic approaches under real-world constraints, and propose a causal learning framework at the intersection of environmental science and machine learning.

1 Introduction

Urban air pollution is a major global health challenge, responsible for millions of premature deaths each year [WHO et al., 2021, Alpert et al., 2012]. Cities such as New Delhi frequently record $PM_{2.5}$ levels far above international guidelines, with severe consequences for respiratory and cardiovascular health [Goyal et al., 2021, Sharma and Dikshit, 2016, Sharma et al., 1998]. Effective interventions require not only measuring pollution but also identifying the causal drivers of elevated concentrations—whether from internal sources such as traffic, kilns, industries, and domestic combustion, or external contributors like crop burning and dust transport. Yet, attribution remains difficult under heterogeneous emissions, dynamic meteorology, and sparse monitoring networks [Sharma and Dikshit, 2016, Bikkina et al., 2019, Apte et al., 2017, Cusworth et al., 2018].

Existing source apportionment methods face well-known limitations. Receptor models such as PMF [Paatero and Tapper, 1994, Reff et al., 2007] and CMB [Watson et al., 1991, 2001] require detailed chemical speciation, which is expensive and difficult to scale. Statistical and machine learning models capture correlations but lack causal interpretability [Bhardwaj et al., 2024, Iyer et al., 2022]. Mechanistic dispersion models, including Gaussian plume formulations, encode the physics of transport but depend on detailed inventories and high-resolution meteorology, which are often unavailable. No single approach provides a satisfactory solution in data-constrained urban contexts.

We address this by framing source attribution as a causal inference problem, where emissions propagate through dispersion physics to produce observed concentrations. We develop a parameterized Gaussian Plume Dispersion Model (GPDM) tailored to sparse sensing, integrating multiple emission categories and allowing both fixed-source scaling and a flexible formulation that estimates an

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

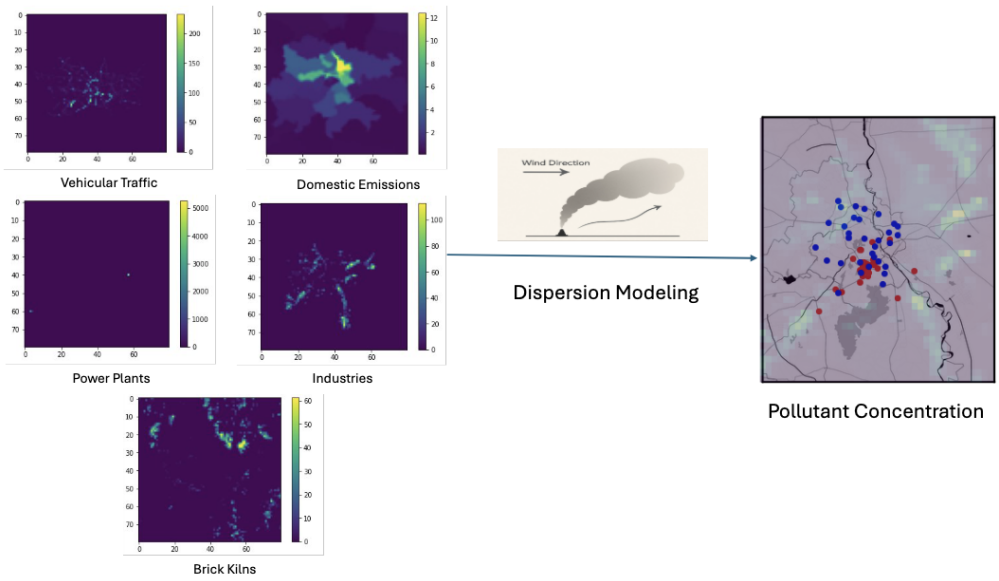

Figure 1: Conceptual framework and illustrative output for causal source attribution of urban air pollution. The left panels show emission inventories from multiple categories. The overlaid heatmap on the right illustrates the resulting concentration field for New Delhi

additional unknown source field. Parameters are learned via maximum likelihood against sensor observations from New Delhi, enabling interpretable attribution of source-specific contributions even under incomplete data.

Our contribution is methodological and conceptual: we present a causal framework that combines mechanistic dispersion physics with parameterized formulations, supporting reasoning about source effects and counterfactual interventions. This work is not a performance-driven empirical study but a position paper, highlighting both the opportunities and limitations of mechanistic causal modeling under sparse data. We propose air pollution source attribution as a benchmark causal challenge for the broader learning community.

2 Methodology

Problem Formulation We frame air pollution source attribution as a causal inference problem, where emission sources act as treatments that causally influence observed concentrations through atmospheric transport mechanisms. Formally, we represent the emissions inventory as a spatiotemporal tensor

$$E[i, j, t] = \{Q_b, Q_i, Q_p, Q_d, Q_v\}, \tag{1}$$

where (i, j) denotes the spatial grid cell, t denotes time, and Q_b, Q_i, Q_p, Q_d, Q_v represent the emission intensities from brick kilns, industries, power plants, domestic combustion, and vehicular traffic, respectively. The causal mechanism linking emissions to observed concentrations is governed by dispersion physics. Given meteorological conditions M_t (e.g., wind speed and direction, stability class) and emission fields E, the concentration at a sensor location s and time t is expressed as

$$C_{s,t} = f(E, M_t; \theta) + \epsilon, \tag{2}$$

where $f(\cdot)$ denotes a mechanistic dispersion model parameterized by θ, and ϵ captures noise and unobserved factors. Estimating θ under sparse sensing conditions allows us to quantify source-specific contributions, identify the presence of missing or unmeasured sources, and reason about counterfactual interventions (e.g., reducing traffic emissions by 30%).

Gaussian Plume Dispersion Model To instantiate $f(\cdot)$, we adopt the Gaussian Plume Dispersion Model (GPDM), a classical solution to the convection–diffusion equation under steady-state

conditions. The concentration at a downwind location (x, y, z) is given by

$$C(x, y, z) = \frac{Q}{2\pi U \sigma_y \sigma_z} \exp\left(-\frac{y^2}{2\sigma_y^2}\right) \left[\exp\left(-\frac{(z-H)^2}{2\sigma_z^2}\right) + \exp\left(-\frac{(z+H)^2}{2\sigma_z^2}\right)\right], \quad (3)$$

where Q is the source emission rate, U is the wind speed, H is the effective stack height, and σ_y, σ_z are lateral and vertical dispersion coefficients determined by Pasquill–Gifford stability parameters. The model naturally encodes the causal propagation of pollutants: emissions Q interact with meteorology $(U, H, \sigma_y, \sigma_z)$ to determine downwind concentrations C.

In practice, we extend this formulation to area and line sources by discretizing the urban region into a grid, assigning emissions $E[i, j, t]$ to each cell, and computing sensor concentrations as the additive superposition of contributions from all cells. To reduce computational overhead, we express the GPDM as a convolution operation between the emission inventory and pre-computed plume filters based on wind vectors, enabling efficient implementation on GPUs. This formulation retains the interpretability of first-principles physics while making parameter estimation tractable in large-scale, sparse-sensing settings.

Parameterized Formulations To adapt the Gaussian plume model for causal source attribution under sparse data conditions, we introduce parameterized formulations that distinguish between known and unknown emission sources.

Fixed-source GPDM. In the fixed-source setting, we assume that the spatial distribution of each source type is known from approximate emission inventories (e.g., brick kilns, industries, power plants, domestic combustion, vehicular traffic). We introduce non-negative scaling coefficients α_s for each source type s, which quantify their relative contributions:

$$C_{s,t}^{\text{pred}} = \sum_{s \in S} \alpha_s \cdot f(E_s, M_t), \quad (4)$$

where $f(\cdot)$ is the Gaussian plume operator applied to source map E_s. Estimating α_s thus provides interpretable source-specific attribution factors.

Flexible-source GPDM. In many urban settings, inventories are incomplete and important contributors (e.g., garbage burning, road dust, construction activities) may be missing. To account for this, we extend the fixed-source formulation by introducing an additional latent emission field E_u with its own scaling parameter α_u. The flexible formulation becomes:

$$C_{s,t}^{\text{pred}} = \sum_{s \in S} \alpha_s \cdot f(E_s, M_t) + \alpha_u \cdot f(E_u, M_t), \quad (5)$$

where E_u is estimated jointly with the scaling coefficients. This allows the model to capture unobserved or unmeasured sources while preserving mechanistic interpretability.

Parameter Estimation We estimate the scaling parameters $\{\alpha_s\}$ and, in the flexible-source case, the unknown field E_u using maximum likelihood estimation (MLE). Let $C_{s,t}^{\text{obs}}$ denote the observed concentration at sensor s and time t, and $C_{s,t}^{\text{pred}}$ the corresponding prediction from the parameterized GPDM. Assuming Gaussian measurement noise, the likelihood is

$$\mathcal{L}(\alpha; E_u) \propto \exp\left(-\frac{1}{2\sigma^2} \sum_{s,t} \left(C_{s,t}^{\text{obs}} - C_{s,t}^{\text{pred}}\right)^2\right). \quad (6)$$

We optimize this objective with respect to the parameters using stochastic gradient descent on GPU, leveraging the convolutional implementation of GPDM filters described in Section 2. Non-negativity constraints are imposed on all scaling coefficients α_s to ensure physical plausibility.

Causal Interpretation The estimated coefficients α_s can be interpreted as causal effect sizes of each emission source type on observed concentrations, conditional on meteorological factors. By embedding dispersion physics within the estimation process, our framework ensures that attribution is grounded in first-principles causal mechanisms rather than correlations alone. The latent emission field E_u represents the influence of unobserved or unmeasured sources, analogous to accounting for hidden confounders in causal inference. Finally, the parameterized GPDM enables counterfactual reasoning: by adjusting α_s for a given source (e.g., simulating a reduction in traffic emissions), we can generate counterfactual concentration maps that estimate the causal impact of interventions on urban air quality.

3 Illustration of Framework

We apply the parameterized Gaussian Plume Dispersion Model (GPDM) to New Delhi using sparse $PM_{2.5}$ sensor data [cpc], meteorological data [NCEP, NWS, NOAA, U.S. DoC, 2000], and approximate emission inventories (see Appendix §A). While absolute prediction errors are high (MAPE $\approx$ $104\% \pm 6\%$), this is consistent with the well-documented uncertainties in dispersion modeling caused by coarse wind data, incomplete emission inventories, and simplifying steady-state assumptions [Guttikunda and Calori, 2013, Namdeo et al., 2012]. Rather than focusing on precise prediction, our aim is to illustrate how mechanistic causal models provide insight into the contributions of different sources under real-world constraints.

Source Contributions in High and Low Pollution Areas The GPDM reveals several causal patterns when comparing relatively high- and low-pollution regions. Power plants, due to their tall stacks ($>$200 m), contribute primarily to regional background pollution rather than to localized peaks, dispersing emissions in the upper atmosphere. By contrast, industries, domestic combustion, and vehicular traffic play a stronger role in shaping elevated concentrations in densely populated and industrial neighborhoods. At the same time, areas with lower measured concentrations tend to be those without major industrial activity or dense traffic corridors, and where local emissions are largely absent. The framework thus distinguishes between macro-scale contributors that elevate baseline levels across the city and local-scale contributors that dominate specific regions.

Policy Implications This causal interpretation has direct relevance for policy. High-pollution areas in New Delhi are often located near clusters of industries or dense residential activity, suggesting that local interventions in these areas—such as emissions standards for industrial stacks or clean-fuel transitions in households—could have measurable benefits. Conversely, policies aimed at power plants are likely to reduce city-wide baseline concentrations but may not strongly affect spatial heterogeneity. Importantly, even under sparse sensing and incomplete inventories, the parameterized GPDM highlights how different sources interact with meteorology to shape observed concentrations. This provides a pathway for designing policies that target the most influential sources at the appropriate spatial scale, while also underscoring the need for better inventories and higher-resolution meteorological data to refine causal attribution.

4 Discussion and Conclusion

We have presented a causal framework for urban air pollution source attribution that combines mechanistic dispersion physics with parameterized formulations adapted to sparse sensing. By embedding the Gaussian Plume Dispersion Model within a causal inference perspective, our approach moves beyond correlational analysis and emphasizes how emissions, meteorology, and unobserved factors interact to generate observed concentrations. Even with incomplete inventories and limited meteorological resolution, the framework yields interpretable insights into the relative roles of different sources and clarifies the distinction between city-wide and local contributors.

While our illustration highlights the potential of mechanistic causal models, it also exposes their limitations. Dispersion models are sensitive to uncertainties in wind fields and source inventories, and cannot fully explain observed variability without richer data. Nevertheless, their causal interpretability makes them valuable for guiding targeted policy interventions, particularly in data-constrained settings where purely statistical methods may misattribute sources.

Looking ahead, we envision this work as a step toward establishing air pollution source attribution as a benchmark causal challenge at the intersection of machine learning and environmental science. Future efforts could integrate more accurate inventories, higher-resolution meteorological inputs, or hybrid statistical–mechanistic models that combine causal interpretability with predictive power. By framing source attribution as a causal inference problem, we hope to encourage collaboration between the causal learning community and domain scientists, and to seed further research that connects methods to pressing societal challenges.

References

Cpcb data portal. https://app.cpcbccr.com/ccr/#/caaqm-dashboard-all/caaqm-landing/caaqm-comparison-data.

Pinhas Alpert, Olga Shvainshtein, and Pavel Kishcha. AOD trends over megacities based on space monitoring using MODIS and MISR. *American Journal of Climate Change*, 01(03):117–131, 2012. doi: 10.4236/ajcc.2012.13010.

Joshua S. Apte, Kyle P. Messier, Shahzad Gani, Michael Brauer, Thomas W. Kirchstetter, Melissa M. Lunden, Julian D. Marshall, Christopher J. Portier, Roel C.H. Vermeulen, and Steven P. Hamburg. High-Resolution Air Pollution Mapping with Google Street View Cars: Exploiting Big Data. *Environmental Science & Technology*, 51(12):6999–7008, June 2017. ISSN 0013-936X. doi: 10.1021/acs.est.7b00891. URL https://doi.org/10.1021/acs.est.7b00891. Publisher: American Chemical Society.

Ankit Bhardwaj, Shiva R. Iyer, Sriram Ramesh, Jerome White, and Lakshminarayanan Subramanian. Understanding sudden traffic jams: From emergence to impact. *Development Engineering*, 8: 100105, 2023. ISSN 2352-7285. doi: https://doi.org/10.1016/j.deveng.2022.100105. URL https://www.sciencedirect.com/science/article/pii/S2352728522000148.

Ankit Bhardwaj, Ananth Balashankar, Shiva Iyer, Nita Soans, Anant Sudarshan, Rohini Pande, and Lakshminarayanan Subramanian. Comprehensive monitoring of air pollution hotspots using sparse sensor networks. *ACM Journal on Computing and Sustainable Societies*, 2024.

Srinivas Bikkina, August Andersson, Elena N. Kirillova, Henry Holmstrand, Suresh Tiwari, A. K. Srivastava, D. S. Bisht, and Örjan Gustafsson. Air quality in megacity Delhi affected by countryside biomass burning. *Nature Sustainability*, 2(3):200–205, March 2019. ISSN 2398-9629. doi: 10.1038/s41893-019-0219-0. URL https://www.nature.com/articles/s41893-019-0219-0. Number: 3.

Center for International Earth Science Information Network . CIESIN Columbia University. Gridded population of the world, version 4 (gpwv4): Population density, revision 11, 20230129 2018. URL https://doi.org/10.7927/H49C6VHW.

Daniel H Cusworth, Loretta J Mickley, Melissa P Sulprizio, Tianjia Liu, Miriam E Marlier, Ruth S DeFries, Sarath K Guttikunda, and Pawan Gupta. Quantifying the influence of agricultural fires in northwest india on urban air pollution in delhi, india. *Environmental Research Letters*, 13 (4):044018, mar 2018. doi: 10.1088/1748-9326/aab303. URL https://doi.org/10.1088%2F1748-9326%2Faab303.

Prachi Goyal, Sunil Gulia, and S.K. Goyal. Identification of air pollution hotspots in urban areas - an innovative approach using monitored concentrations data. *Science of The Total Environment*, 798:149143, 2021. ISSN 0048-9697. doi: https://doi.org/10.1016/j.scitotenv.2021.149143. URL https://www.sciencedirect.com/science/article/pii/S0048969721042169.

Sharath Guttikonda. A primer on source apportionment. Urbanemissions.info, 2011.

Sarath K. Guttikunda and Giuseppe Calori. A gis based emissions inventory at 1 km × 1 km spatial resolution for air pollution analysis in delhi, india. *Atmospheric Environment*, 67:101–111, 2013. ISSN 1352-2310. doi: https://doi.org/10.1016/j.atmosenv.2012.10.040. URL https://www.sciencedirect.com/science/article/pii/S1352231012010229.

Shiva R. Iyer, Ananth Balashankar, William H. Aeberhard, Sujoy Bhattacharyya, Giuditta Rusconi, Lejo Jose, Nita Soans, Anant Sudarshan, Rohini Pande, and Lakshminarayanan Subramanian. Modeling fine-grained spatio-temporal pollution maps with low-cost sensors. *npj Climate and Atmospheric Science*, 5(1):76, Oct 2022. ISSN 2397-3722. doi: 10.1038/s41612-022-00293-z. URL https://doi.org/10.1038/s41612-022-00293-z.

Anil Namdeo, Ibrahim Sohel, Justin Cairns, Margaret Bell, Mukesh Khare, and Shiva Nagendra. Performance evaluation of air quality dispersion models in delhi, india. In Sébastien Rauch and Gregory M. Morrison, editors, *Urban Environment*, pages 121–130, Dordrecht, 2012. Springer Netherlands. ISBN 978-94-007-2540-9.

NCEP, NWS, NOAA, U.S. DoC. Ncep fnl operational model global tropospheric analyses, continuing from july 1999, 2000. URL https://doi.org/10.5065/D6M043C6.

Pentti Paatero and Unto Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, 1994.

Adam Reff, Shelly I Eberly, and Prakash V Bhave. Receptor modeling of ambient particulate matter data using positive matrix factorization: review of existing methods. *Journal of the Air & Waste Management Association*, 57(2):146–154, 2007.

Mukesh Sharma and Onkar Dikshit. Comprehensive study on air pollution and green house gases (GHGs) in Delhi. Technical report, IIT Kanpur, Jan 2016. URL https://cerca.iitd.ac.in/uploads/Reports/1576211826iitk.pdf.

Sangeeta Sharma, Gulsha Rai Sethi, Ashish Rohtagi, Anil Chaudhary, Ravi Shankar, Jawahar Singh Bapna, Veena Joshi, and Debarati Ghua Sapir. Indoor air quality and acute lower respiratory infection in indian urban slums. *Environmental health perspectives*, 106(5):291–297, 1998.

John G Watson, Judith C Chow, and Thompson G Pace. Chemical mass balance. In *Data Handling in Science and Technology*, volume 7, pages 83–116. Elsevier, 1991.

John G Watson, Judith C Chow, and Eric M Fujita. Review of volatile organic compound source apportionment by chemical mass balance. *Atmospheric Environment*, 35(9):1567–1584, 2001.

WHO, UNAIDS, et al. *Air quality guidelines: global update 2021*. World Health Organization, 2021.

A Deriving Approximate Emissions Inventory

The general bottom-up approach to source apportionment starts with building a detailed emissions inventory, which is a detailed spatiotemporal map of all pollution-causing activities in the city, along with estimates of their fuel consumption and emission [Guttikonda, 2011]. This emissions inventory is then used as input for pollution dispersion models to compute the pollution map of the city. The process of collating an accurate and detailed emissions inventory is a significant undertaking, requiring on-ground surveys of emission activities. In [Guttikunda and Calori, 2013], Guttikunda et. al. build a detailed GIS-based emissions inventory where they divided their study area into grids of 0.01^ox0.01^o (approximately 1km x 1km). According to them, the eight major sources of pollution emissions in New Delhi are brick kilns, industries, power plants, domestic emissions, vehicular emissions, road dust, construction activities, and garbage burning. Every cell in their inventory consists of the details of emissions from the eight sources. Unfortunately, the inventory is not publicly available, so we had to reconstruct an approximate version. In the process, we were unable to find the updated distribution maps for road dust, garbage burning, and construction activities, which according to Guttikunda et. al., account for about 20% of New Delhi's emissions.

We denote our emissions inventory (E) as a matrix

$$E[i, j, t] = \{Q_b, Q_i, Q_p, Q_d, Q_v\} \tag{7}$$

where, i and j are latitude and longitude indexes, ant t is the time. Q_b, Q_i, Q_p, Q_d, Q_t refer to the intensities of pollution sources corresponding to brick kilns, industries, power plants, domestic emissions, and vehicular emissions.

Brick Kilns, Industries, Power Plants and Domestic Emissions: Using the distribution map of industries and brick kilns from [Guttikunda and Calori, 2013], we reconstructed the positional map of these pollution sources by using the number of colored pixels to approximately infer the number of sources in different cells. Thus

$$E[i, j][Q_b] = P_r(i, j) \tag{8}$$

where $P_r(i, j)$ is the number of red dots in cell i,j of brick kilns map, and

$$E[i, j][Q_i] = P_b(i, j) \tag{9}$$

where $P_b(i, j)$ is the number of black dots in cell i,j of industries map from [Guttikunda and Calori, 2013]. Note that we have made the simplifying assumption that all instances of brick kilns and

industries emit equal amounts of $PM_{2.5}$ pollutants per unit of time, making the inventory independent of time for Q_b and Q_i. For the case of power plants, we assumed that the emissions of power plants are proportional to their power capacity. Adjusting for the historical data on the decommissioning and changes in capacities of power plants, we were able to create a similar inventory for power plants.

$$E[i,j][Q_p] = \sum \text{Capacity}_{pp'}, \text{ for } pp' \in \text{ cell i,j.} \tag{10}$$

For domestic sources of pollution, we utilized the Gridded Population of the World (GPWv4) population density data [Columbia University, 2018] as representative of domestic pollution emission intensity. This data was adjusted to fit the 0.01^ox0.01^o positional map. The same assumption of spatiotemporal uniformity was applied to convert this map into emission intensities.

$$E[i,j][Q_d] = \text{Population}_{i,j} \tag{11}$$

where $\text{Population}_{i,j}$ is the population density in cell i,j.

Vehicular Emissions: We could not find open-source detailed traffic measurement data for New Delhi. Therefore, we used Google maps to obtain a representation of New Delhi's traffic at 4 different time snapshots, 6 AM, 12 PM, 6 PM, and 10 PM for a section of the study area. The typical traffic patterns on different days were only slightly different, hence, we assume that these maps are an average representation of the traffic in the city on all days. From Google's explanation, we know that green roads signify no delays in traffic, which is equivalent to allowing vehicles to move at the speed limit. Evoking prior work on traffic curves [Bhardwaj et al., 2023], we can reasonably bound the maximum possible link density (the density of vehicles on the road) in such a case to less than 0.1. Thus, we assume that pollution caused due to traffic on green-marked roads is negligible. At 6 AM, we have near-zero traffic-based emissions. Therefore,

$$E[i,j,6am][Q_v] = 0. \tag{12}$$

Using the 6 AM map as the baseline, and using image differencing, we computed the additional traffic at other times. In the differenced images, we found roads marked 'orange', 'red', and 'maroon' corresponding to progressively increasing levels of traffic delay. Based on the structure of the traffic curves in [Bhardwaj et al., 2023], we posit that 'orange' roads corresponding to mild delay be in the free-flow traffic regime, 'red' roads corresponding to significant congestion would be in the spiraling regime, and 'maroon' roads corresponding to traffic jams would be in the jam regime. For traffic curves derived for typical segments in other cities in developing countries like Nairobi and Sao Paulo, the link densities for the three regions would approximately be in the ratio 1:2:3. Thus, we use this ratio for our analysis.

$$E[i,j,t][Q_v] = P_o(I_t - I_{6am}, i, j) + 2 * P_r(I_t - I_{6am}, i, j) + 3 * P_m(I_t - I_{6am}, i, j) \tag{13}$$

where I_t refers to traffic snapshot at time t, and P_o, P_r, P_m are pixel-counting functions.

Note that here we have assumed the traffic emissions to be proportional to link density (number of vehicles). The resulting emissions inventory was sum-normalized by the total emission statistics in [Guttikunda and Calori, 2013]

$$E[i,j,t] = \left\{ \frac{Q_b T_b}{\sum_{i,j,t} Q_b}, \frac{Q_i T_i}{\sum_{i,j,t} Q_i}, \frac{Q_p T_p}{\sum_{i,j,t} Q_p}, \frac{Q_d T_d}{\sum_{i,j,t} Q_d}, \frac{Q_v T_v}{\sum_{i,j,t} Q_v} \right\} \tag{14}$$

where, T_b, T_i, T_p, T_d, T_v are corresponding total emission statistics.