Explainable Machine Learning for Predicting Personal Exposure to Fine Particulate Matter and Volatile Organic Compounds in Urban Mobility Environments

Blazej Jaroszuk
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
blazej.work@ucf.edu

Alister A. Durning
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
alister.durning@unh.edu

Hanne-Lore Jungfer
Department of Computer Science, University of Houston, Houston, TX, USA.
hannelorework@uh.edu

Lia Kostolzin
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
lia.work@oregonstate.edu

Abstract

Personal exposure to fine particulate matter (PM2.5) and volatile organic compounds (VOCs) in urban mobility environments is shaped by rapidly changing microenvironments, route choice, transport mode, and pollutant-specific source dynamics. Fixed-site monitoring and conventional land-use regression provide useful ambient estimates, but they cannot directly resolve the exposure sequence experienced by an individual moving through streets, transit stops, vehicles, and indoor transition spaces. Machine learning can capture nonlinear interactions among these factors, yet predictive accuracy alone is insufficient when outputs are intended to inform route selection, mode choice, or exposure-reduction decisions. This manuscript develops an explainable, uncertainty-aware machine learning framework for predicting personal PM2.5 and VOC exposure in urban mobility settings. The framework formalizes exposure as a microenvironment-weighted concentration-time integral, couples multi-output learning with physically motivated monotonic constraints, and treats attribution and counterfactual analysis as model components rather than post hoc decoration. It further specifies conformal uncertainty intervals and a route-level decision rule that uses upper predictive bounds rather than point predictions. The central contribution is a coherent modeling logic that aligns data representation, pollutant-specific mechanisms, explanation fidelity, uncertainty calibration, and practical decision use. The framework yields testable design principles for evaluating whether explanations are stable under correlated features, whether uncertainty intervals remain valid across persons and routes, and whether exposure-aware recommendations preserve both scientific defensibility and interpretability for users and public-health practitioners.

Keywords

personal exposure; fine particulate matter; volatile organic compounds; explainable machine learning; urban mobility; uncertainty quantification

1. Introduction

Urban air pollution is a persistent environmental health concern, and fine particulate matter remains a leading contributor to the global burden of disease [1]. Personal exposure, however, is not equivalent to ambient concentration measured at a central monitor. Individuals move through streets, sidewalks, buses, trains, bicycles, and building interfaces, and each microenvironment imposes different source proximity, ventilation, deposition, and chemical transformation conditions [2]. Personal monitoring can resolve some of this heterogeneity, but deployments are constrained by instrument cost, participant burden, and limited temporal coverage [3]. The growth of low-cost sensing and mobile data has therefore motivated hybrid exposure models that combine stationary monitoring, portable measurements, traffic information, meteorology, and mobility traces [4].

Machine learning is attractive for this problem because exposure depends on nonlinear and context-dependent interactions among traffic intensity, street geometry, ventilation, weather, and individual activity. Tree ensembles such as random forests and gradient-boosted trees are strong general-purpose learners for tabular environmental data [5], [6]. Yet the same flexibility that improves prediction can obscure why a given route or time window is predicted to be high-exposure. Post hoc attribution methods, including Shapley-value explanations and local surrogate models, provide useful diagnostics [7], [8], while interpretable modeling and careful explanation design guard against explanations that appear plausible but do not faithfully represent model behavior [9], [10]. Data stewardship and reproducibility practices are also essential because exposure models combine geolocation, personal sensing, and environmental layers whose provenance and quality directly affect inference [11]. Earth-system applications of machine learning further show that data-driven prediction is most credible when process understanding constrains model structure and evaluation [12].

A major gap remains between pollutant-specific exposure science and explainable prediction practice. Transport-cabin and commuter studies demonstrate substantial variation in pollutant levels across modes and ventilation conditions, including conditions that can produce health-relevant exposures during daily travel [13]. Mobility-aware exposure estimates can differ materially from residence-based or fixed-site estimates when spatiotemporal variability is ignored [14]. The health significance of such differences is reinforced by evidence on pollution-related mortality and global exposure trends [15], [16], [17]. Nevertheless, many prediction studies optimize a single pollutant, use generic feature sets, and report explanations without specifying what those explanations are intended to support. The resulting models may predict well in aggregate while failing to identify actionable route or mode interventions, and they may produce unstable attributions when correlated traffic, land-use, and meteorological variables are present.

This paper addresses that gap by developing a technically explicit framework for explainable prediction of personal PM2.5 and VOC exposure in urban mobility environments. The contribution is not a claim of empirical performance. It is a modeling and evaluation architecture that connects exposure physics to machine learning design. First, it formalizes exposure as a microenvironment-weighted concentration-time integral and separates PM2.5

from VOC mechanisms. Second, it specifies a multi-output additive learner with pollutant-specific constraints and an objective that penalizes physically implausible sensitivities. Third, it integrates Shapley attribution and counterfactual route analysis as first-class outputs. Fourth, it defines calibrated uncertainty intervals and a decision rule that uses upper exposure bounds. Finally, it sets out an evaluation logic that tests prediction, calibration, explanation fidelity, stability, and decision utility without conflating association with causation.

2. Exposure Process and Data Representation

2.1 Personal exposure as a spatiotemporal integral

Personal exposure should be defined at the level of an individual trajectory rather than a static address. For individual i, let T_i,m denote the time spent in microenvironment m during an observation window, C_m(t) the pollutant concentration in that microenvironment at time t, and R_i,m(t) a breathing-rate or weighting function. A general exposure target is

$$E_i \;=\; \sum_m \int_\{T_i, m\}\, C_m(t)\; R_i, m(t)\; dt$$

where E_i is the individual exposure burden over the window, m indexes microenvironments such as sidewalk, bus interior, subway platform, or building foyer, and t is time. If R_i,m(t) is set to one, the expression yields concentration-time exposure; if it represents inhalation rate, the expression approximates inhaled dose, subject to additional deposition and absorption assumptions. The separation of C_m(t) and R_i,m(t) matters because route recommendations can change both the concentration sequence and the duration of exposure. A model that predicts only daily average concentration discards the timescale on which mobility decisions operate.

The integral also clarifies why mobility traces are not merely auxiliary covariates. Global positioning and transit records determine the time weights T_i,m, while microenvironment classification determines which concentration regime applies. Misclassification of a waiting period at a bus stop as an indoor period, or of a vehicle cabin as ambient street air, can bias the target even when concentration predictions are accurate. The uncertain geographic context problem is therefore intrinsic: the relevant context is not only where a person is located but also when, for how long, and under what ventilation and source conditions [22]. For this reason, the modeling target should be constructed from synchronized time-activity, route, and microenvironment labels rather than from static buffers alone.

2.2 Pollutant-specific mechanisms

PM2.5 and VOCs require different representations because their governing processes differ. PM2.5 exposure in urban mobility is influenced by near-road traffic emissions, brake and tire wear, construction dust, resuspension, and secondary aerosol formation. In vehicle cabins, filtration and air-exchange rates modulate the relationship between outdoor and indoor concentrations, while open windows or poor filtration can bring cabin levels close to street levels. VOCs are more source-specific and chemically diverse. Fuel evaporation, vehicle exhaust, solvent use, consumer products, and building emissions can produce sharp spatial and temporal gradients, and sorption or desorption from surfaces can create lagged responses. Commuter studies show that mode, fuel type, and route affect particulate exposure [19], while VOC exposures vary across transport modes and are sensitive to ventilation and source proximity [18]. A common feature set may support joint learning, but a shared model must not force PM2.5 and VOC to share the same monotonic or interaction structure.

This distinction has a direct modeling consequence. The output vector should be treated as multi-task rather than as a single scalar with a pollutant indicator. Shared layers or shared tree structure can borrow strength when features such as traffic and meteorology affect both pollutants, but pollutant-specific residual structure is necessary for VOC source patterns that are weakly related to PM2.5 mass. The framework therefore permits partial parameter sharing while retaining separate calibration and explanation outputs for each pollutant.

3. Explainable Multi-Output Learning Architecture

3.1 Feature construction and data quality

The input representation for person i at time t should be organized into four blocks. The first block contains meteorological variables such as temperature, relative humidity, wind speed, and boundary-layer proxies. The second contains traffic and street-scale variables, including traffic volume, heavy-duty vehicle share, road class, intersection density, and distance to roadway. The third contains mobility and microenvironment variables, including mode, cabin ventilation state, route segment, stop duration, and indoor-outdoor transition indicators. The fourth contains sensing and temporal variables, including low-cost sensor readings, time of day, day of week, and lagged concentrations. Provenance metadata should record instrument type, calibration status, sampling interval, geolocation accuracy, and missingness mechanism, consistent with reproducible data stewardship [11].

Data leakage is a central risk. If same-day central-monitor values are used as features when the target is personal exposure measured during the same period, the model may learn the monitor's temporal pattern rather than the personal microenvironment. If route identifiers are used as features and the same route appears in training and test sets, evaluation can overstate generalization to new routes. The framework therefore distinguishes features available before a decision from features available only after the exposure period. For route-level prediction, only pre-travel or real-time information should enter the decision model. For retrospective exposure reconstruction, lagged and co-located measurements may be used but must be evaluated with group-wise splits by person, route, and time.

3.2 Additive prediction function

The prediction model is specified as an additive ensemble

$$f_\theta(x_i, t) = \sum_{k=1}^{K} g_k(x_i, t)$$

where x_i,t is the feature vector for individual i at time t, θ collects model parameters, K is the number of base learners, and g_k denotes a tree or smooth component. The output f_θ is a vector whose first component targets a transformed PM2.5 exposure and whose second targets a transformed VOC exposure. Logarithmic or generalized-linear link transformations can stabilize right-skewed concentrations, but the transformation must be inverted before decision use. The additive form is chosen because it supports exact or near-exact attribution for tree ensembles and because it permits monotonic constraints on selected features. It does not imply that the underlying exposure process is additive; rather, it is a model class that balances flexibility, interpretability, and computational feasibility.

The model should be trained with grouped cross-validation that separates persons and routes, not only random observations. For each fold, preprocessing steps such as imputation, scaling, and sensor-drift correction must be fitted only on the training partition. Baselines

should include fixed-site monitor interpolation, land-use regression, random forest, and gradient boosting under the same feature-availability rules [2], [5], [6]. This comparison is necessary because a complex model is justified only if it improves out-of-sample exposure prediction or decision-relevant calibration relative to transparent alternatives.

3.3 Objective with physical regularization

The training objective should combine predictive loss with complexity control and physically motivated sensitivity penalties. A suitable form is

$$L(\theta) = \sum_i w_i \ell(y_i, f_\theta(x_i)) + \lambda_1 \Omega(\theta) + \lambda_2 \sum_{j \in M} max(0, -\partial f_\theta / \partial x_j)^2$$

where y_i is the observed or reconstructed exposure vector, ℓ is a multi-output loss such as weighted squared error or Huber loss, w_i are sampling weights that may account for uneven time coverage, Ω(θ) is a complexity penalty, λ_1 and λ_2 are nonnegative regularization strengths, M is a set of features with expected nonnegative directional effects, and ∂f_θ/∂x_j is the model sensitivity to feature j. The derivative penalty discourages predictions that decrease when a feature known to increase exposure rises, conditional on the stated feature set. For example, for fixed ventilation and meteorology, increasing near-road heavy-duty traffic should not reduce primary PM2.5 exposure in a local sensitivity sense. The set M must be pollutant-specific and context-specific; forcing a universal monotonic relationship for ventilation would be indefensible because increased ventilation can either dilute indoor-generated VOCs or admit outdoor PM2.5.

The penalty is not a substitute for process modeling. It is a soft constraint that encodes only directional expectations that are defensible across the training domain. Strong monotonic constraints can reduce variance and improve extrapolation, but they can also introduce bias when the assumed direction is wrong. The framework therefore treats λ_2 and M as design choices to be stress-tested, not as fixed truths. A model with no monotonic penalty should be retained as a diagnostic baseline so that the effect of constraints on accuracy and explanation stability can be assessed.

4. Attribution and Counterfactual Explanation

4.1 Shapley attribution for local exposure predictions

For a prediction f(x), the contribution of feature j can be expressed using a Shapley decomposition

where F is the full feature set, S is a subset of features excluding j, f_S is the model restricted to the feature subset S, x_S is the corresponding feature vector, x_{S ∪ {j}} adds feature j, and w_S = |S|!(|F|−|S|−1)!/|F|! is the combinatorial weight. The sum of φ_j over all features equals f(x) minus a baseline expectation under the chosen value function. This property makes the attribution useful for communicating why a specific route segment is predicted to have higher PM2.5 or VOC exposure than a reference segment. For tree ensembles, exact or efficient algorithms can compute these values under specified assumptions [7]. The interpretation remains conditional on the model and feature distribution; it is not a causal effect of changing traffic or ventilation in the real world.

Because urban features are correlated, attribution can split credit among variables in unstable ways. Traffic volume, road class, and distance to roadway may all carry similar information. A responsible explanation layer should therefore report grouped attributions for semantically related feature blocks as well as individual feature values. If the grouped

attribution is stable but the individual attribution is not, the explanation should be presented at the group level. This design reduces the risk of false precision in route-level explanations.

4.2 Counterfactual route and microenvironment analysis

Attribution explains a prediction as it stands. Counterfactual analysis asks how the prediction would change under an alternative route, mode, or waiting location. For a feature j and two feasible values a and b, define

$$\Delta_j(x;\ a, b)\ =\ f_\theta(x\ with\ x_j = b)\ -\ f_\theta(x\ with\ x_j = a)$$

where x with x_j = b denotes the feature vector with feature j set to b while other features are held at their observed or scenario-specified values. For route decisions, j may represent route segment, transport mode, or cabin ventilation state; a and b must belong to the feasible action set for that individual. The counterfactual is model-based and should be labeled as such. It does not estimate the effect of a randomized intervention, and it is sensitive to the assumptions used for features that cannot remain fixed when a route changes. For example, changing from bus to bicycle changes travel time, traffic proximity, and breathing rate simultaneously. A defensible counterfactual should therefore alter a coherent bundle of variables rather than a single feature in isolation.

4.3 Explanation fidelity and stability

Explanation quality should be evaluated rather than assumed. Fidelity asks whether the explanation actually represents the model's decision in a local neighborhood; stability asks whether similar inputs yield similar explanations; sanity checks ask whether explanations respond to known structural changes, such as removing a route segment or increasing traffic exposure along a corridor. Local surrogate explanations can be useful for communication [8], but they should not replace model-native attribution when the model class permits it. In high-stakes exposure decisions, interpretable models or constrained additive models may be preferable to post hoc explanation of an unconstrained black box [9]. General explanation design principles, including the distinction between local and global explanation and the need to match explanation type to audience and decision, are well established [10]. The framework therefore treats explanation fidelity and stability as model outputs with their own evaluation protocol.

5. Uncertainty, Calibration, and Decision Use

5.1 Predictive intervals and conformal calibration

Point predictions are insufficient for route selection because a small predicted difference between two routes may be overwhelmed by model uncertainty. A distribution-free approach uses a calibration set of residuals to construct intervals. For a nominal miscoverage rate α, let q_{1−α} be an appropriate empirical quantile of absolute residuals from held-out calibration data. A simple split-conformal interval is

$$C_\alpha(x) = [f_\theta(x) - q_{1-\alpha}, f_\theta(x) + q_{1-\alpha}]$$

where C_α(x) is the predictive interval, f_θ(x) is the point prediction, and q_{1−α} is the calibration quantile. The interval has a coverage guarantee under exchangeability between calibration and test data. That assumption is strong for personal exposure because persons, routes, seasons, and sensor platforms may not be exchangeable. The framework should therefore use Mondrian or group-conditional calibration by mode, season, and sensor type where sample sizes permit, and it should report coverage and interval width by group.

Calibrated regression methods provide a related route for models that output predictive distributions [20], while conformal methods offer a practical distribution-free alternative when distributional assumptions are uncertain [21].

Uncertainty should be decomposed into at least three sources: measurement uncertainty from low-cost sensors, representativeness uncertainty from assigning a microenvironment label, and model uncertainty from limited coverage of routes or conditions. Conformal intervals capture the aggregate effect of these sources on prediction errors if the calibration set reflects the deployment conditions, but they do not identify which source dominates. Diagnostic stratification can reveal whether coverage fails for a particular mode or season, which is more useful for model improvement than a single aggregate coverage number.

5.2 Exposure-aware route and mode decisions

For decision support, the relevant quantity is not the point prediction alone but an upper bound on exposure under the chosen action. Let A(x) be the feasible action set for a traveler, such as a set of routes or modes, and let U_α(x,a) be the upper endpoint of the predictive interval for action a. A conservative route rule is

$$a^{\wedge} * (x) \; = \; argmin_\{a \in A(x)\} \, f_\theta(x, a) \; subject \; to \; U_\alpha(x, a) \; \leq \; E_max$$

where a^(x) is the selected action, f_θ(x,a) is the predicted exposure for action a, U_α(x,a) is the upper confidence or conformal bound, and E_max is a user- or policy-defined exposure threshold. If no action satisfies the constraint, the rule should return the action with the smallest upper bound and flag that the threshold cannot be met under the available options. This rule is deliberately conservative: it prefers a route whose worst plausible exposure is acceptable over a route with a slightly lower point prediction but a wide uncertainty interval. The threshold itself is a normative input. It may be tied to air-quality guidelines [1] or to personal health goals, but it should not be hidden inside the model.

Decision rules also require equity analysis. A route recommendation that reduces exposure for one traveler may shift waiting time, travel cost, or physical activity in ways that affect other risks. The model should expose the trade-offs rather than collapsing them into a single uninterpreted score. Group-level explanations can show whether recommendations systematically favor or disadvantage certain neighborhoods, modes, or travel times.

6. Evaluation Protocol and Technical Discussion

A credible evaluation of this framework must separate prediction, calibration, explanation, and decision utility. Prediction should be assessed with mean absolute error, root mean squared error, and bias by pollutant, mode, and season, using group-wise splits by person, route, and time. A random split is insufficient because it allows the same route or person to appear in both training and test sets. Baselines should include fixed-site interpolation, land-use regression, random forest, and gradient boosting. Because personal exposure targets are often noisy and unevenly sampled, evaluation should also report error by concentration range and by microenvironment rather than only aggregate averages.

Calibration should report empirical coverage of predictive intervals at multiple nominal levels, interval width, and coverage by mode, season, and sensor platform. A model that achieves 90 percent marginal coverage but only 60 percent coverage for cyclists or winter commutes is not adequate for decision support. Explanation evaluation should report fidelity to the model, stability under small input perturbations, stability under correlated-feature resampling, and group-level consistency. Counterfactual explanations should be checked

against feasibility constraints and should be presented with their assumptions. These evaluations do not require unrealistically precise sensors; they require transparent protocols and honest reporting of uncertainty.

Several error modes are specific to this domain. Geolocation error can assign a traveler to the wrong street canyon. Microenvironment classifiers can confuse a covered transit stop with an indoor space. Low-cost PM sensors are sensitive to humidity and particle composition, and VOC sensors often suffer from cross-sensitivity and drift. Route and traffic features may be missing or aggregated at a resolution too coarse to represent near-road gradients. The model should include missingness indicators and should avoid imputing a microenvironment label when the probability of misclassification is high. Instead, predictions with low microenvironment confidence should be widened or flagged.

Explainability introduces its own failure modes. Shapley values can be unstable when features are strongly correlated, and they explain the model rather than the physical world. Counterfactual route changes may violate assumptions when travel time, breathing rate, and exposure duration change jointly. A route recommendation can be technically valid under the model yet operationally infeasible if transit schedules or sidewalk availability are ignored. These limitations argue for a layered explanation system: model-native attributions for analysts, grouped feature contributions for decision-makers, and simple route comparisons for travelers. Each layer should carry the same uncertainty and assumption disclosures.

The framework also has implications for data governance and public health practice. Personal exposure prediction combines sensitive mobility data with environmental measurements, and the granularity that makes prediction useful also creates privacy risk. Aggregation, purpose limitation, and clear consent are not peripheral concerns; they determine whether such models can be deployed ethically. The scientific value of explainable prediction lies in making model behavior contestable. A model that can show which route segments, sources, and conditions drive a prediction is easier to audit, improve, and challenge than one that reports only an accuracy score.

7. Conclusion

This manuscript has developed an explainable machine learning framework for predicting personal PM2.5 and VOC exposure in urban mobility environments. The framework begins from a microenvironment-resolved exposure integral rather than a static ambient concentration, which makes route, mode, duration, and ventilation explicit parts of the prediction problem. It then specifies a multi-output additive learner with pollutant-specific physical regularization, a Shapley attribution layer, counterfactual route analysis, conformal uncertainty intervals, and a conservative decision rule based on upper exposure bounds. The central argument is that explainability in exposure prediction is not a presentation layer added after modeling. It is a design constraint that determines which features are admissible, which explanations are reportable, how uncertainty is calibrated, and how recommendations are made.

The framework also clarifies where empirical work must be careful. Correlated urban features can make local attributions unstable; exchangeability assumptions can fail across persons, seasons, and modes; and counterfactual route changes can confound exposure with travel time and activity. These are not reasons to abandon explainable prediction. They are reasons to evaluate explanation fidelity, group-conditional calibration, and decision utility as rigorously as prediction error. By aligning exposure science, machine learning, and decision

analysis, the proposed framework provides a defensible basis for models that not only estimate personal exposure but also make the basis of that estimate visible and contestable.

References

1. World Health Organization. (2021). WHO global air quality guidelines: Particulate matter (PM2.5 and PM10), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide. World Health Organization.

2. Jerrett, M., Arain, A., Kanaroglou, P., Beckerman, B., Potoglou, D., Sahsuvaroglu, T., Morrison, J., & Giovis, C. (2005). A review and evaluation of intraurban air pollution exposure models. Journal of Exposure Analysis and Environmental Epidemiology, 15(2), 185–204.

3. Koehler, K. A., & Peters, T. M. (2015). New methods for personal exposure monitoring for airborne particles. Current Environmental Health Reports, 2(4), 399–411.

4. Snyder, E. G., Watkins, T. H., Solomon, P. A., Thoma, E. D., Williams, R. W., Hagler, G. S. W., Shelow, D., Hindin, D. A., Kilaru, V. J., & Preuss, P. W. (2013). The changing paradigm of air pollution monitoring. Environmental Science & Technology, 47(20), 11369–11377.

5. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.

6. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). ACM.

7. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30, pp. 4765–4774).

8. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM.

9. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206–215.

10. Molnar, C. (2022). Interpretable machine learning: A guide for making black box models explainable (2nd ed.). Independently published.

11. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3, 160018.

12. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat. (2019). Deep learning and process understanding for data-driven Earth system science. Nature, 566(7743), 195–204.

13. Wang, H., Wei, W., Wang, K., Kong, H., Vieira, C. L. Z., Koutrakis, P., ... & Giorio, C. (2025). A global multidimensional analysis of air pollution in land, sea, and air transport cabins reveals substantial health risks. One Earth, 8(1).

14. Park, H. S., & Kwan, M.-P. (2017). Individual exposure estimates may be erroneous when spatiotemporal variability of air pollution and human mobility are ignored. Health & Place, 43, 85–94.

15. Lelieveld, J., Evans, J. S., Fnais, M., Giannadaki, D., & Pozzer, A. (2015). The contribution of outdoor air pollution sources to premature mortality on a global scale. Nature, 525(7569), 367–371.

16. Burnett, R., Chen, H., Szyszkowicz, M., Fann, N., Hubbell, B., Pope, C. A., ... & Spadaro, J. V. (2018). Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. Proceedings of the National Academy of Sciences, 115(38), 9592–9597.

17. Shaddick, G., Thomas, M. L., Mudu, P., Ruggeri, G., & Gumy, S. (2020). Half the world's population are exposed to increasing air pollution. npj Climate and Atmospheric Science, 3, 23.

18. Jo, W. K., & Park, K. H. (2010). Commuter exposure to volatile organic compounds in different modes of transportation in Korea. Atmospheric Environment, 44(24), 3017–3025.

19. Zuurbier, M., Hoek, G., Oldenwening, M., Lenters, V., Meliefste, K., van den Hazel, P., & Brunekreef, B. (2010). Commuters' exposure to particulate matter air pollution is affected by mode of transport, fuel type, and route. Environmental Health Perspectives, 118(6), 783–789.

20. Kuleshov, V., Fenner, N., & Ermon, S. (2018). Accurate uncertainties for deep learning using calibrated regression. In Proceedings of the 35th International Conference on Machine Learning (pp. 2796–2804). PMLR.

21. Angelopoulos, A. N., & Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511.

22. Kwan, M.-P. (2012). The uncertain geographic context problem. Annals of the Association of American Geographers, 102(5), 958–968.