Wearable Sensor-Based Health Risk Modeling for Evaluating the Effects of Transportation-Related Air Pollution on Cardiopulmonary Function

Annunziata Onisto
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
onisto179@ku.edu

Abstract

Transportation-related air pollution produces dynamic, spatially heterogeneous exposures that are poorly represented by fixed-site monitors and residential addresses alone. At the same time, wearable sensors now permit repeated measurement of location, activity, ventilation, heart rhythm, and respiratory signals in free-living conditions. This manuscript develops a methodological framework for wearable sensor-based health risk modeling that connects personal exposure to transportation-related particulate and gaseous pollutants with cardiopulmonary functional endpoints. The central contribution is an integrated analytical design that combines microenvironment-aware exposure reconstruction, ventilation-adjusted deposited dose, physiologically informed response modeling, regularized machine learning for individual risk estimation, calibration analysis, and causal interpretation through target trial logic. The framework specifies inputs, feature construction, objective functions, uncertainty propagation, and evaluation logic without relying on unavailable trial data. It emphasizes the distinction between exposure concentration, inhaled dose, and subclinical physiological response, and it treats sensor calibration, time synchronization, activity misclassification, temporal drift, and confounding as first-order design problems rather than post hoc corrections. The proposed approach is intended to support studies that estimate how transport microenvironments and daily mobility patterns affect heart rate variability, respiratory function, and related cardiopulmonary outcomes. Its scientific significance lies in providing a transparent, testable architecture for combining wearable environmental sensing with wearable physiological sensing, and in clarifying the conditions under which model outputs can support mechanistic inference rather than only correlational prediction.

Keywords

wearable sensors; transportation-related air pollution; cardiopulmonary function; personal exposure modeling; machine learning; causal inference

1. Introduction

Transportation-related air pollution is a persistent cardiopulmonary hazard because it combines combustion-derived particulate matter, nitrogen dioxide, black carbon, ultrafine particles, and volatile organic compounds in spatially concentrated and temporally variable mixtures. Epidemiological evidence has established robust associations between fine particulate matter and cardiovascular morbidity, respiratory symptoms, lung function decline, and mortality [1], [2]. Mechanistic reviews further support pathways involving autonomic imbalance, oxidative stress, endothelial dysfunction, and pulmonary inflammation [3], while long-term cohort evidence has shaped regulatory and public health interpretations of

particulate risk [4]. These findings motivate a more granular question: how can personal, mobility-resolved exposure to transport-related pollution be linked to short-term cardiopulmonary functional change in a way that is both statistically defensible and mechanistically informative?

Traditional exposure assessment often assigns ambient concentrations from fixed monitors to residential locations or uses land-use regression surfaces as proxies for personal exposure. Such approaches are valuable for population epidemiology, but they smooth over the microenvironments where transport-related exposures occur, including road segments, bus stops, vehicle cabins, subway platforms, and cycling corridors. Reviews of personal exposure monitoring show that time-activity patterns and indoor-outdoor differences can substantially alter the exposure contrast experienced by an individual [5], [6]. The growth of low-cost and portable sensing has expanded the possibility of resolving these contrasts, but it has also introduced calibration, drift, and data-quality challenges that are not incidental to modeling [7], [8], [9].

Wearable physiological sensing adds a second source of high-frequency information. Heart rate variability, resting heart rate, respiratory rate, and, in some settings, portable spirometry or photoplethysmography-derived measures can capture subclinical responses to pollution exposure. However, physiological signals are sensitive to posture, sleep, physical activity, temperature, medication use, and psychological stress. A credible model must therefore separate exposure effects from autonomic and behavioral confounders rather than treating wearable time series as direct outcome labels.

The gap addressed here is the absence of a coherent modeling architecture that treats wearable environmental sensing, wearable physiological sensing, and causal interpretation as one problem. Existing prediction models may fit exposure-response associations, but they often lack explicit dose construction, uncertainty propagation, and calibration checks. Conversely, epidemiological designs may use wearable data without exploiting the temporal structure available for lagged and cumulative exposure analysis. This manuscript proposes a wearable sensor-based health risk modeling framework for evaluating the effects of transportation-related air pollution on cardiopulmonary function. The contribution is not an empirical validation study. It is a technically specified analytical framework that defines exposure, dose, response, model estimation, calibration, evaluation, and causal assumptions in a form that can be implemented in a future cohort or panel study.

2. Wearable Sensing and Cardiopulmonary Endpoint Construction

2.1 Microenvironment-Aware Personal Exposure Reconstruction

The first modeling layer converts raw wearable and contextual data into a time-resolved personal exposure series. Inputs include global positioning system trajectories, accelerometer-derived activity intensity, wearable air-quality measurements, temperature and humidity, and map-derived microenvironment labels such as road type, vehicle cabin, transit mode, indoor location, and outdoor walking or cycling. Time synchronization is essential because exposure and physiological responses are analyzed over short windows. A common time base, quality flags for signal loss, and explicit handling of missing intervals should precede feature construction.

Microenvironment reconstruction can be represented as a switching process. At time t, individual i occupies one of several microenvironments m, and the personal breathing-zone

concentration is a weighted combination of microenvironment-specific concentrations after accounting for infiltration, proximity, and sensor response. A compact formulation is

$$C_i(t) = \sum_m S_{im}(t)\lambda_m C_m(t)$$

where $C_i(t)$ is the estimated personal exposure concentration for individual i at time t, $S_{im}(t)$ is an indicator or soft membership probability that individual i occupies microenvironment m at time t, λ_m is a dimensionless attenuation or infiltration factor linking the microenvironment concentration to the breathing zone, and $C_m(t)$ is the measured or modeled pollutant concentration in that microenvironment. The membership term may be probabilistic when GPS is uncertain or when a person moves between indoor and outdoor spaces. The factor λ_m is pollutant-specific and mode-specific; it should be estimated from co-location data rather than fixed by assumption. This relation is deliberately simple because its role is to make the exposure construction auditable. More complex dispersion or cabin-exchange models can replace $\lambda_m C_m(t)$, but the modeling logic remains transparent only if the substitution rules and calibration sources are documented.

The output of this layer is not a single daily exposure value. It is a set of pollutant-specific time series, together with uncertainty flags for sensor drift, humidity effects, and microenvironment misclassification. These time series support lagged, cumulative, and peak exposure features that are more relevant to short-term cardiopulmonary responses than a daily average alone.

2.2 Physiological Endpoints and Measurement Constraints

Cardiopulmonary endpoints should be chosen for biological meaningfulness and wearable feasibility. Heart rate variability is a leading candidate because it reflects autonomic modulation and has established measurement standards [10], [11]. Ambient pollution has been associated with altered heart rate variability in panel studies, supporting its use as a subclinical endpoint [12]. Resting heart rate, respiratory rate, and step- or activity-adjusted heart rate recovery can provide complementary information. Portable spirometry can measure forced expiratory volume in one second and peak expiratory flow, but adherence and forced-effort artifacts are more demanding than heart rhythm monitoring. Blood pressure, when available through validated wearable cuffs, can add a vascular endpoint, but cuff inflation and posture requirements constrain continuous use.

Signal quality is a substantive modeling issue. Electrocardiographic artifacts, motion, poor contact, and ectopic beats can distort heart rate variability. Preprocessing should include beat detection, artifact correction with explicit rules for rejecting rather than imputing low-quality segments, and adjustment for activity and posture. Respiratory signal quality depends on sensor placement and movement. For spirometry, test acceptability criteria and effort are essential. The endpoint model should therefore use quality-weighted observations or hierarchical structures that allow within-person variability to be separated from between-person differences. Treating every wearable-derived value as equally reliable would propagate measurement error into exposure-response estimates.

3. Methodological Framework for Integrated Risk Modeling

3.1 Exposure and Deposited Dose

Personal exposure concentration is necessary but not sufficient for health risk modeling because the dose delivered to the respiratory tract depends on ventilation, deposition fraction, and duration. A physically interpretable dose construct is

$$D_i = \int_0^T V_i(t)\, C_i(t)\, \delta_i(t)\, dt$$

where D_i is the deposited dose over an observation window from time 0 to T, $\dot{V}_i(t)$ is minute ventilation, $C_i(t)$ is the personal exposure concentration, and $\delta_i(t)$ is the pollutant- and particle-size-dependent deposition fraction. Minute ventilation can be estimated from accelerometer intensity and heart rate through activity-specific calibration, but this estimate carries uncertainty. The deposition fraction depends on particle size distribution, breathing route, and respiratory rate; for gaseous pollutants it follows different uptake kinetics. The framework therefore treats D_i as an estimated dose with propagated uncertainty, not as a directly observed quantity. This distinction matters because an individual who cycles through traffic may receive a higher deposited dose than a sedentary individual in a vehicle cabin at the same ambient concentration.

Cumulative, lagged, and peak dose features can be derived from D_i. For short-term cardiopulmonary responses, windows of minutes to hours are plausible, whereas longer windows may capture delayed inflammatory or autonomic effects. The choice of window length should be prespecified and subjected to sensitivity analysis rather than optimized only for predictive performance.

3.2 Nonlinear Physiologically Informed Response Model

The response layer links dose and physiological state to cardiopulmonary endpoints. A general representation is

$$Y_i(t) = f\big(P_i(t), D_i(t-\tau), X_i(t)\big) + \varepsilon_i(t)$$

where $Y_i(t)$ is a cardiopulmonary endpoint such as log-transformed heart rate variability, respiratory rate, or spirometric function; $P_i(t)$ is physiological state including activity, posture, heart rate, and sleep status; $D_i(t-\tau)$ is deposited dose lagged by τ; $X_i(t)$ is a vector of time-varying covariates such as temperature, humidity, time of day, and medication use; f is a flexible but constrained function; and $\varepsilon_i(t)$ is residual variation. The form of f should reflect physiological constraints. For example, heart rate variability is strongly influenced by activity and posture, so the model should either restrict analysis to stable resting windows or include activity interaction terms. Lagged effects should be represented through distributed lag structures rather than a single arbitrary lag when the biological response is expected to persist.

This formulation preserves the distinction between exposure, dose, and response. A model that predicts $Y_i(t)$ from ambient concentration at a fixed monitor can conflate these layers. A model that uses $D_i(t-\tau)$ and $P_i(t)$ explicitly makes the dose construction visible and allows sensitivity analyses to show how ventilation estimation or deposition assumptions affect conclusions.

3.3 Feature Fusion and Risk Estimation

Wearable data are heterogeneous: environmental signals are continuous but sensor-specific; mobility data are categorical or geospatial; physiological data are high-frequency and autocorrelated. The framework therefore constructs a feature vector that summarizes exposure, dose, physiological response, and covariates over biologically relevant windows. A generic fusion map is

$$z_i = \varphi(E_i, D_i, Y_i, X_i)$$

where z_i is the fused feature vector for individual i, E_i contains exposure summaries such as peak, cumulative, and time-in-transport metrics, D_i contains deposited-dose summaries, Y_i contains physiological endpoint summaries, X_i contains covariates, and φ denotes a prespecified feature construction and normalization function. The function φ must be fitted or defined within training folds to avoid information leakage. Because autocorrelated time series can inflate apparent performance, cross-validation should group observations by participant and, where possible, by time block.

For a binary adverse response or event within a defined horizon, the model output can be written as

$$r_i = \sigma\big(g_\theta(z_i)\big)$$

where r_i is the estimated probability of the response, σ is the logistic function, g_θ is a risk function parameterized by θ, and z_i is the fused feature vector. Tree-based ensembles are credible candidates for tabular feature sets because they handle nonlinearity and interactions without requiring a fully specified functional form [13], [14]. Gradient boosting can improve predictive accuracy when sufficient observations are available [14]. Interpretability methods such as Shapley-value summaries can identify which exposure or physiological features contribute to predictions [15], but they should be used to interrogate model behavior, not to establish causal effects. If the analysis uses raw time-series sequences, a recurrent or temporal convolutional architecture may be appropriate, but it should be compared against simpler regularized and tree-based baselines. Complexity is justified only when it improves calibration and out-of-sample performance under participant-level validation.

3.4 Regularized Learning and Model Interpretation

When the risk model is trained on an approved cohort, an appropriate objective combines discrimination with regularization:

$$L(\theta) = -\sum_i [y_i \log r_i + (1 - y_i) \log(1 - r_i)] + \lambda \|\theta\|_1$$

where L(θ) is the penalized negative log-likelihood, y_i is the observed binary response, r_i is the model-estimated probability, λ is a nonnegative regularization parameter, and $\|\theta\|_1$ is the L1 norm of the parameter vector. The L1 penalty encourages sparse models when many derived features are correlated. Alternative penalties, such as group lasso for blocks of exposure or physiological features, may be more appropriate when the goal is to retain or exclude entire sensor domains. The objective does not resolve confounding by itself. It produces a prediction rule that must be evaluated for calibration, transportability, and fairness before any health-risk interpretation is attached to it.

3.5 Causal Interpretation and Target Trial Logic

Prediction and causal inference answer different questions. A model may predict cardiopulmonary response accurately from correlated features while misestimating the effect of changing transportation exposure. To move from prediction toward effect estimation, the analysis should specify a target trial: the eligibility criteria, exposure contrast, follow-up period, and outcome definition that would have been used in a randomized study. Wearable cohort data can then be analyzed as an observational emulation of that target trial [16]. Causal inference in public health requires explicit assumptions about consistency, positivity, exchangeability, and correct adjustment for time-varying confounders [17].

For a discrete exposure contrast d, a g-computation-style target parameter can be expressed as

$$\mathrm{E}\left[Y_i^{\,d}\right] = \sum_x \mathrm{E}[Y_i | D_i = d, X_i = x]\mathrm{P}(X_i = x)$$

where Y_i^d is the potential outcome under deposited dose or exposure contrast d, D_i is the exposure variable, X_i is a vector of measured confounders, and the summation is over confounder strata. In continuous-time wearable data, X_i may include time-varying activity, location, temperature, and prior physiological state. This relation is an identification formula under stated assumptions; it is not evidence that those assumptions hold. Negative-control outcomes, lagged exposure models, and sensitivity analyses for unmeasured confounding should accompany any causal claim. The framework should report both predictive performance and effect estimates, with the distinction clearly maintained.

4. Technical Discussion

4.1 Transport Microenvironments and Exposure Error

Transport microenvironments are not merely high-exposure settings; they are structured environments in which pollutant composition, breathing rate, and time-activity patterns change together. A cyclist and a car passenger may experience different cabin or roadway concentrations, different ventilation rates, and different deposition patterns. A global analysis of land, sea, and air transport cabins has reported substantial health risks from air pollution in these settings [18], reinforcing the need to treat transport mode as a causal and exposure-relevant variable rather than a nuisance covariate. Within the proposed framework, transport mode enters through the microenvironment membership $S_{im}(t)$, the attenuation factor λ_m, and the ventilation estimate $\dot{V}_i(t)$. This separation allows the model to distinguish a high ambient concentration from a high personal dose caused by activity or cabin ventilation.

Exposure error is not random only. Low-cost optical particulate sensors are sensitive to humidity and particle composition; electrochemical gas sensors can drift and cross-react; GPS can lose signal in tunnels, underground transit, and dense urban canyons. If such errors are correlated with transport mode, they can bias exposure contrasts. The framework should therefore include co-location calibration, field blank or reference comparisons, and mode-specific error models. Soft membership probabilities for microenvironment assignment are preferable to hard labels when location uncertainty is high. Sensitivity analysis should compare results under alternative calibration models and microenvironment definitions.

4.2 Calibration and Uncertainty

A risk model that discriminates well but is poorly calibrated can misstate individual probabilities. Calibration should be assessed on held-out participants, not only on held-out observations from the same participants. A common calibration model is

$$logit\, P(y = 1|r) = a + b\, logit(r)$$

where r is the model-estimated probability, y is the observed outcome, $\mathrm{logit}(p)=\log[p/(1-p)]$, a is the calibration intercept, and b is the calibration slope. An intercept different from zero indicates systematic over- or under-estimation; a slope different from one indicates that predicted probabilities are too extreme or too compressed. Calibration is central to predictive analytics in medicine and should be reported alongside discrimination metrics [19]. In the wearable setting, calibration can be assessed overall and by transport mode,

season, activity level, and participant subgroup. Poor subgroup calibration may reveal sensor drift or context-dependent exposure error rather than a true difference in risk.

Uncertainty should be propagated from sensor calibration, microenvironment assignment, ventilation estimation, and model estimation. A practical approach combines Monte Carlo simulation for exposure and dose features with participant-level bootstrap resampling for model parameters. The resulting intervals should describe both measurement uncertainty and sampling uncertainty. If the uncertainty is wide enough to change the direction of an exposure-response contrast, the model should not be used for individual risk communication.

4.3 Evaluation, Baselines, and Clinical Translation

Evaluation should include baselines that isolate the value of wearable data. A fixed-site monitor model, a residential exposure model, a personal exposure model without physiological features, a physiological-only model, and a full wearable fusion model should be compared under the same participant-level validation scheme. This design clarifies whether improvements come from finer exposure resolution, physiological context, or overfitting to idiosyncratic wearable patterns. Metrics should include discrimination, calibration, net benefit at clinically relevant thresholds, and subgroup performance. Temporal validation, in which training uses earlier periods and testing uses later periods, is important because sensor drift and seasonal pollutant mixtures can change over time.

Clinical translation requires more than predictive accuracy. A wearable risk model should not be presented as a diagnostic tool unless it has been prospectively validated in an appropriate population and shown to alter management or improve outcomes. In the research setting, the model is best used to estimate short-term exposure-response associations, identify susceptible subgroups, and generate hypotheses about transport-mode interventions. Machine learning in medicine has advanced rapidly, but deployment requires attention to dataset shift, human factors, and workflow integration [20]. For transportation and environmental health, the analogue is a model that supports exposure reduction and study design rather than replacing clinical judgment.

4.4 Privacy, Equity, and Deployment Conditions

Wearable environmental and physiological monitoring generates sensitive location, health, and behavioral data. Privacy protection requires data minimization, on-device preprocessing where feasible, encryption, controlled access, and clear consent for secondary use. Equity is also a modeling issue. Low-cost sensors may perform differently across environmental conditions, and study volunteers may not represent populations with the highest exposure burdens. Ethical machine learning in medicine emphasizes transparency, accountability, and attention to biased or unrepresentative data [21]. In this framework, fairness monitoring should examine whether model performance and calibration differ by age, sex, socioeconomic status, occupation, and neighborhood. Deployment conditions should include battery constraints, participant burden, and the need for rapid feedback that does not overinterpret uncertain risk estimates.

5. Conclusion

This manuscript has developed a wearable sensor-based health risk modeling framework for evaluating the effects of transportation-related air pollution on cardiopulmonary function. The framework treats personal exposure as a microenvironment-weighted process, converts exposure into ventilation-adjusted deposited dose, models cardiopulmonary response with

explicit physiological covariates and lag structures, fuses heterogeneous wearable features into a regularized risk model, and distinguishes predictive performance from causal effect estimation. It also specifies calibration, uncertainty propagation, baseline comparisons, and deployment safeguards. The central technical argument is that wearable environmental sensing and wearable physiological sensing must be connected through a transparent dose and response model; otherwise, high-frequency data can create an illusion of precision without resolving the exposure contrasts that matter. The framework does not claim empirical validation, but it provides a defensible analytical architecture for future studies that combine transportation exposure reconstruction, cardiopulmonary endpoint measurement, machine learning, and target trial emulation. Its practical value lies in making assumptions visible, error modes testable, and model outputs interpretable within the limits of wearable measurement.

References

1. Brook, R. D., Rajagopalan, S., Pope, C. A., III, Brook, J. R., Bhatnagar, A., Diez-Roux, A. V., Holguin, F., Hong, Y., Luepker, R. V., Mittleman, M. A., Peters, A., Siscovick, D., Smith, S. C., Jr., Whitsel, L., & Kaufman, J. D. (2010). Particulate matter air pollution and cardiovascular disease: An update to the scientific statement from the American Heart Association. Circulation, 121(21), 2331–2378.

2. Pope, C. A., III, & Dockery, D. W. (2006). Health effects of fine particulate air pollution: Lines that connect. Journal of the Air & Waste Management Association, 56(6), 709–742.

3. Rajagopalan, S., Al-Kindi, S. G., & Brook, R. D. (2018). Air pollution and cardiovascular disease: JACC state-of-the-art review. Journal of the American College of Cardiology, 72(17), 2054–2070.

4. Dockery, D. W., Pope, C. A., Xu, X., Spengler, J. D., Ware, J. H., Fay, M. E., Ferris, B. G., Jr., & Speizer, F. E. (1993). An association between air pollution and mortality in six U.S. cities. New England Journal of Medicine, 329(24), 1753–1759.

5. Steinle, S., Reis, S., & Sabel, C. E. (2013). Quantifying human exposure to air pollution—Moving from static monitoring to spatio-temporally resolved personal exposure. Science of the Total Environment, 443, 184–193.

6. Nieuwenhuijsen, M. J. (2015). Personal exposure monitoring of air pollution and health: A review. International Journal of Environmental Research and Public Health, 12(2), 1815–1833.

7. 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.

8. Castell, N., Dauge, F. R., Schneider, P., Vogt, M., Lerner, U., Fishbain, B., Broday, D., & Bartonova, A. (2017). Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates? Environment International, 99, 293–302.

9. Rai, A. C., Kumar, P., Pilla, F., Skouloudis, A. N., Di Sabatino, S., Ratti, C., Yasar, A., & Rickerby, D. (2017). End-user perspective of low-cost sensors for outdoor air pollution monitoring. Science of the Total Environment, 607–608, 691–705.

10. Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology. (1996). Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. Circulation, 93(5), 1043–1065.

11. Sassi, R., Cerutti, S., Lombardi, F., Malik, M., Huikuri, H. V., Peng, C. K., Schmidt, G., & Yamamoto, Y. (2015). Advances in heart rate variability signal analysis: Joint position statement by the e-Cardiology ESC Working Group and the European Heart Rhythm Association co-endorsed by the Asia Pacific Heart Rhythm Society. Europace, 17(9), 1341–1353.

12. Gold, D. R., Litonjua, A., Schwartz, J., Lovett, E., Larson, A., Nearing, B., Allen, G., Verrier, M., Cherry, R., & Verrier, R. (2000). Ambient pollution and heart rate variability. Circulation, 101(11), 1267–1273.

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

14. 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.

15. 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).

16. Hernán, M. A., & Robins, J. M. (2016). Using big data to emulate a target trial when a randomized trial is not available. American Journal of Epidemiology, 183(8), 758–764.

17. Glass, T. A., Goodman, S. N., Hernán, M. A., & Samet, J. M. (2013). Causal inference in public health. Annual Review of Public Health, 34, 61–75.

18. 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).

19. Van Calster, B., McLernon, D. J., van Smeden, M., Wynants, L., & Steyerberg, E. W. (2019). Calibration: The Achilles heel of predictive analytics. BMC Medicine, 17(1), 230.

20. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.

21. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. PLoS Medicine, 15(11), e1002689.