

Integrating Explainable Machine Learning and Environmental Epidemiology to Predict Acute Pulmonary Function Decline Associated with Ambient Particle Gamma Radiation Exposure

Banjamin J. Lepozai

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,
OR, USA.

benjamin1975@oregonstate.edu

Botianyi Xiong

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

botianyi.work@binghamton.edu

Abstract

Ambient particulate matter serves as a vector for a range of hazardous constituents, including radon decay progeny that emit ionizing gamma radiation. Short-term elevations in particle-bound gamma activity have been associated with acute decrements in lung function, particularly among individuals with pre-existing respiratory disease. While environmental epidemiology has identified these associations, the translation of such evidence into operational early warning and clinical decision support systems remains nascent. This paper presents a systems-level framework that integrates explainable machine learning with environmental epidemiology to predict acute pulmonary function decline attributable to ambient particle gamma radiation exposure. We examine the architecture of a distributed monitoring and modeling infrastructure, emphasizing trade-offs between sensor fidelity, data granularity, and computational scalability. The design of interpretable predictive models is discussed, focusing on the tension between complex deep learning architectures and inherently transparent models, and on post hoc explanation techniques such as SHAP and LIME that can bridge the gap for clinical and regulatory acceptability. Governance challenges are analyzed with regard to fairness across heterogeneous populations, privacy-preserving health data linkage, and accountability in automated public health alerts. The paper further addresses deployment robustness in the face of concept drift, sensor attrition, and evolving emission profiles, along with sustainability considerations in model training and inference. Policy implications for air quality standards for particulate radioactivity, integration into clinical guidelines for chronic obstructive pulmonary disease management, and the institutional arrangements required to sustain such a system are explored. By synthesizing perspectives from exposure science, machine learning, and socio-technical governance, we delineate a pathway toward a transparent, equitable, and resilient predictive public health infrastructure.

Keywords

Explainable machine learning; environmental epidemiology; particle gamma radiation; pulmonary function; public health surveillance; fairness; system architecture; robustness.

1. Introduction

Ambient particulate matter is a complex and dynamic mixture of chemical and radioactive species whose health effects extend beyond those captured by mass concentration alone. Among the lesser-studied constituents, particle-bound radioactivity arising from the attachment of short-lived radon progeny to fine aerosols represents an exposure pathway that delivers ionizing gamma radiation to the respiratory epithelium. Epidemiological investigations have demonstrated that short-term fluctuations in particle gamma radiation activity are associated with acute pulmonary function decline, especially in vulnerable populations such as patients with chronic obstructive pulmonary disease. While these findings reinforce the need to incorporate radiological parameters into respiratory risk assessment, the transition from retrospective epidemiological analyses to prospective, high-resolution predictive systems requires a fundamental rethinking of data infrastructure, modeling strategy, and institutional governance. Contemporary machine learning offers the capacity to capture complex, non-linear exposure-response relationships across multiple spatio-temporal scales. However, the opacity of many high-performance models conflicts with the transparency demanded by clinical decision-making, regulatory accountability, and public trust. The present work articulates an integrative framework in which explainable machine learning is systematically coupled with environmental epidemiology to forecast acute declines in lung function attributable to ambient particle gamma radiation. Rather than focusing on a single model or dataset, we address the entire socio-technical system needed to operationalize such predictions. This includes sensor network design, data harmonization, model selection with explainability constraints, fairness auditing, and the policy pathways that determine whether predictive insights translate into measurable health protection. By treating prediction as an infrastructural rather than purely algorithmic endeavor, we illuminate the structural trade-offs that shape the feasibility, equity, and longevity of data-driven environmental health interventions.

2. Background and Related Work

Indoor radon has long been established as the second leading cause of lung cancer, prompting comprehensive public health guidelines and remediation programs worldwide [1]. Outdoor particulate radioactivity, although typically present at lower concentrations, has emerged as a distinct exposure of concern because radon progeny attach to ambient fine particles and can be transported over regional scales. Fine particulate matter itself is causally associated with respiratory and cardiovascular mortality [2]. The radiological dimension of particulate matter adds an additional biological mechanism, as particle-bound radionuclides deliver localized alpha and gamma radiation doses to lung tissues following inhalation. Recent epidemiological studies using continuous measurements of particle gamma activity have reported adverse acute effects on pulmonary function that are independent of particle mass [11, 14]. These findings underscore the need to move beyond mass-based air quality indices toward multipollutant and multi-property exposure metrics. Parallel advances in machine learning have substantially improved the spatial and temporal resolution of air pollution estimates. Deep recurrent neural networks, for instance, can assimilate meteorological, traffic, and land-use covariates to produce hourly pollutant concentration maps at the urban scale [3]. In clinical informatics, interpretable models such as generalized additive models with pairwise interactions have been shown to achieve near state-of-the-art performance while providing clinically transparent risk scores [4]. More recent work has developed model-agnostic explanation frameworks that attribute predictions to input features, most notably SHAP and LIME [5, 6]. These methods open the possibility of applying high-capacity models to health-sensitive tasks without entirely sacrificing interpretability, provided that their outputs are

appropriately contextualized for end users. Meanwhile, the broader machine learning community has drawn attention to fairness concerns that arise when predictive models inadvertently encode or amplify disparities across demographic groups [7]. All of these threads converge when designing a system that forecasts lung function decline from particle radioactivity. The system must reconcile the complexity of atmospheric and dosimetric processes, the sparsity of specialized radiation monitoring, the strict interpretability requirements of clinical practice, and the ethical imperative of equitable protection.

3. System Architecture and Data Infrastructure

Constructing an operational prediction system for particle gamma radiation and pulmonary function decline demands a layered architecture that spans environmental sensing, data integration, epidemiological linkage, and model serving. The first layer is the radiation monitoring network itself. Unlike regulated criteria pollutants for which dense government-operated networks exist, ambient gamma spectrometry is conducted by a much smaller number of research-grade stations and nuclear safety monitors. Recent advances in low-cost sensor technology for air quality [8] have not yet been widely extended to spectrally resolved gamma detection, creating a fundamental resolution gap. A viable system must therefore fuse sparse high-fidelity gamma measurements with more abundant proxies such as radon flux maps derived from geological surveys and meteorological models. This fusion can be accomplished through geostatistical methods or learned embeddings within a multitask neural network. Data integration must also incorporate continuous pulmonary function data, which in clinical cohorts may be collected through daily home spirometry, electronic health records, or wearable devices. The linkage of environmental and health data streams introduces substantial privacy and security challenges, requiring differential privacy mechanisms that limit re-identification risk without destroying the epidemiological signal [17]. From a systems perspective, the data pipeline must handle streaming ingestion, quality flagging, spatio-temporal alignment, missingness imputation, and concept drift monitoring. The architecture should be designed for modularity such that new sensor modalities, cohort registries, and even novel pollutants can be incorporated without requiring a complete retooling of the prediction engine. Scalability is a further concern; the computational cost of producing real-time individual-level risk estimates across an urban population can quickly outstrip available resources unless inference is structured hierarchically, with population-level baseline models refined by personal covariates at the edge. Trade-offs between centralized and federated learning paradigms become pertinent, especially when health data cannot leave institutional firewalls. The governance of such a distributed system implicates multiple agencies: environmental protection authorities that manage ambient monitoring, public health departments that oversee cohort linkages, and clinical providers who ultimately act on the predictions. Clear data stewardship agreements and interoperable application programming interfaces are prerequisites for operational continuity.

4. Explainable Machine Learning Model Design

Selecting an appropriate modeling approach for predicting acute pulmonary function decline requires balancing predictive accuracy, calibration, computational efficiency, and interpretability. The exposure-response relationship between particle gamma radiation and lung function is expected to be non-linear, temporally lagged, and modified by co-exposures such as fine particulate matter, ozone, and temperature. Ensemble tree methods such as gradient boosting have historically offered strong performance on tabular environmental health data and provide built-in feature importance metrics. However, their capacity to model

intricate temporal dependencies is limited compared to recurrent or attention-based neural architectures. Deep learning models can capture lagged effects, hourly-seasonal patterns, and interaction terms implicitly, yet their opacity often precludes acceptance in clinical guideline development. Post hoc explanation techniques offer a pragmatic middle ground. SHAP values, grounded in Shapley value theory, can decompose a prediction for an individual patient into additive contributions from each input feature, including gamma activity levels averaged over different time windows [5]. LIME can generate locally faithful linear surrogates, facilitating communication with clinicians who reason about linear risk factors [6]. However, these explanations must be presented with careful uncertainty quantification to prevent overinterpretation, a challenge that becomes acute when predictions drive medical interventions. An alternative design philosophy treats interpretability as a first-order constraint by employing inherently transparent model classes, such as time-varying coefficient models with constrained interactions, which can be estimated via penalized regression. Such models may sacrifice some predictive power but align more naturally with epidemiological conventions and regulatory review processes. A hybrid system could cascade models: a transparent primary risk score governs routine alerts, while a deeper auxiliary model identifies high-complexity patterns and flags cases for expert review. Fairness must be evaluated across the entire modeling pipeline, not merely at the outcome stage. If gamma monitoring stations are preferentially sited in affluent neighborhoods, exposure measurement error may be differential by socioeconomic status, leading to biased risk estimates for disadvantaged communities. Fairness-aware learning algorithms that enforce equality of opportunity or minimize conditional prediction disparities can mitigate some of these structural biases, though they cannot fully compensate for systematic data gaps [7]. Ultimately, the choice of model architecture and explanation strategy is inseparable from the sociotechnical context in which the system will be deployed, and it must be made through deliberative engagement with clinical stakeholders, affected communities, and regulatory bodies.

5. Governance, Fairness, and Ethical Considerations

The deployment of a predictive system linking environmental gamma radiation to individual pulmonary function inevitably surfaces ethical tensions surrounding data ownership, consent, and distributive justice. Participants in epidemiological cohorts who contribute daily lung function measurements and activity trackers may not anticipate that their data will be combined with real-time environmental feeds to generate personalized health risk scores. Transparent informed consent procedures, dynamic opt-out mechanisms, and community advisory boards are essential governance instruments to maintain trust. Fairness concerns extend beyond measurement bias to the very definition of the outcome. Acute pulmonary function decline captured through forced expiratory volume in one second is a clinical endpoint that has different prevalence, measurement reliability, and baseline distributions across age, sex, and racial groups. Models trained predominantly on data from older white male cohorts may systematically underestimate risk for women or ethnic minorities, potentially channeling protective resources away from those most in need. Algorithmic fairness frameworks that require equalized false negative rates across groups can be operationalized during model training, but they entail a trade-off with overall accuracy that must be explicitly negotiated with public health authorities. Furthermore, environmental justice considerations arise because communities living near legacy industrial sites or in housing with poor ventilation may experience higher baseline particle radioactivity burdens. A prediction system that is blind to these structural determinants risks individualizing risk

without addressing its upstream drivers, thereby reinforcing the narrative that health outcomes are solely a matter of personal susceptibility. Strong governance requires that predictive outputs be accompanied by contextual information on environmental inequities and that the system be integrated into broader initiatives aimed at exposure reduction, not merely risk forecasting. Transparency obligations under emerging artificial intelligence regulations, including the need to provide meaningful explanations for automated decisions that affect health, further shape the governance landscape [13]. Establishing audit trails, version control for models, and independent fairness certification can serve as procedural safeguards. The institutional locus of accountability must be clearly defined: when a model fails to predict an exacerbation that leads to hospitalization, it is essential to delineate the responsibilities of the model developer, the clinician, and the public health agency that authorized the system's use.

6. Deployment, Robustness, and Sustainability

Transitioning from a prototype model to a resilient, long-term operational service entails confronting a suite of robustness challenges that are often underappreciated in machine learning research. Ambient particle gamma radiation levels are influenced by meteorological regimes, soil moisture, atmospheric mixing, and changes in the built environment. A model trained during a period of stable weather patterns may experience significant performance degradation when a heatwave or wildfire smoke event alters both exposure and baseline population vulnerability. Concept drift detection methods that monitor the divergence between incoming data and the training distribution are critical to trigger model retraining or human review [9]. Sensor attrition and calibration drift additionally threaten data quality. Environmental gamma detectors require periodic calibration with reference sources, and field-deployed units may suffer from gain shifts due to temperature fluctuations. A systematic quality assurance framework that integrates redundant sensor paths, anomaly detection, and automated flagging of suspect readings is necessary to maintain the integrity of the prediction pipeline. Sustainability also encompasses the computational and energy footprint of the system. Training large transformer-based models on high-resolution spatio-temporal environmental grids can consume substantial electricity and contribute to carbon emissions unless mitigated by efficient architectures and hardware [18]. In a domain where the ultimate goal is health protection, it is incongruous to build systems whose environmental costs are disproportionate to public health gains. Lightweight model distillation, quantized inference, and federated averaging over edge devices offer pathways to reduce the energy profile without unduly compromising performance. Institutional sustainability is equally important. Many promising environmental health informatics systems have foundered after the cessation of research funding because they lacked a clear transition to routine government operation. Embedding the prediction platform within existing public health surveillance programs, such as those operated by national environmental public health tracking networks, can provide a durable operational home. This integration must be accompanied by workforce training, standard operating procedures, and budget lines that recognize predictive services as a continuous public health function rather than a one-time technology project.

7. Policy Implications and Regulatory Pathways

The availability of reliable, explainable predictions of acute pulmonary function decline associated with particle gamma radiation would have direct repercussions for environmental regulation and clinical practice. Current air quality standards for particulate matter do not account for radiological composition, and there is no established ambient limit for particle gamma activity. The development of a prediction system that demonstrates consistent

associations between gamma radiation metrics and clinically meaningful decrements in lung function could motivate the incorporation of radioactive properties into the multipollutant air quality indicators used for public advisories [10]. Epidemiological evidence, including short-term studies in susceptible populations [11, 14], provides a scientific foundation upon which regulatory bodies could begin to consider guideline values for particulate radioactivity. Such a process would require dose-response meta-analyses and health impact assessments that are explicitly informed by the predictive models operating at population scale. Policy mechanisms might include requiring air quality monitoring networks to report particle gamma activity alongside standard criteria pollutants and integrating gamma radiation forecasts into the air quality index color-coded alert system, enabling at-risk individuals to reduce outdoor exposure on high-radioactivity days. In the clinical domain, guidelines for the management of chronic obstructive pulmonary disease could incorporate environmental risk scores derived from the prediction system, prompting preemptive adjustment of medication or recommendation of personal protective measures [15]. However, caution is warranted to avoid shifting the burden of protection entirely onto patients. Regulatory instruments such as radon-resistant building codes for new construction, which already address indoor exposures, could be extended to consider outdoor particulate radioactivity in urban planning and ventilation standards. International harmonization of monitoring protocols is needed, given the transboundary transport of fine particles and attached radionuclides. The European Atlas of Natural Radiation exemplifies the kind of collaborative mapping effort that can underpin both scientific research and policy formation [19]. The governance of the prediction system itself will require new regulatory clarity. As automated environmental health alerts become more prevalent, legal frameworks must specify the evidentiary standards for model-based warnings, the liability regimes when incorrect predictions cause harm or undue alarm, and the rights of individuals to access and contest algorithmic determinations.

8. Conclusion

The convergence of explainable machine learning and environmental epidemiology holds transformative potential for protecting respiratory health in the face of ambient particle gamma radiation exposure. Realizing this potential, however, demands more than the development of an accurate model. It requires constructing an integrated socio-technical system that spans heterogeneous monitoring networks, privacy-preserving data architectures, interpretable prediction engines, and robust governance frameworks that actively address fairness, sustainability, and long-term institutional viability. The structural trade-offs between predictive performance and transparency, between centralized efficiency and federated data sovereignty, and between individual-level risk communication and population-level environmental justice cannot be resolved through technical optimization alone. They require deliberate, participatory design processes that engage clinicians, environmental regulators, community advocates, and machine learning engineers in an ongoing dialogue. By foregrounding these system-level considerations, this paper provides a roadmap for building predictive public health infrastructures that are not only scientifically rigorous but also ethically defensible and operationally resilient. The integration of particle gamma radiation as a distinct exposure dimension within the broader air pollution surveillance apparatus could ultimately reduce the burden of acute respiratory morbidity and advance the vision of precision environmental health.

References

1. World Health Organization. (2009). WHO handbook on indoor radon: A public health perspective. World Health Organization.
2. Pope, C. A., Burnett, R. T., Thun, M. J., Calle, E. E., Krewski, D., Ito, K., & Thurston, G. D. (2002). Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution. *JAMA*, 287(9), 1132-1141.
3. Li, X., Peng, L., Yao, X., Cui, S., Hu, Y., You, C., & Chi, T. (2017). Long short-term memory neural network for air pollutant concentration predictions: Method development and evaluation. *Environmental Pollution*, 231, 997-1004.
4. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1721-1730). ACM.
5. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765-4774). Curran Associates.
6. 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.
7. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
8. Snyder, E. G., Watkins, T. H., Solomon, P. A., Thoma, E. D., Williams, R. W., Hagler, G. S., ... & Preuss, P. W. (2013). The changing paradigm of air pollution monitoring. *Environmental Science & Technology*, 47(20), 11369-11377.
9. Gama, J., Žliobaite, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1-37.
10. European Commission. (2013). Council Directive 2013/59/Euratom of 5 December 2013 laying down basic safety standards for protection against the dangers arising from exposure to ionising radiation. *Official Journal of the European Union*, L31, 1-73.
11. Blomberg, A. J., Coull, B. A., Jhun, I., Vieira, C. L. Z., Zanobetti, A., Garshick, E., ... & Koutrakis, P. (2019). Effect modification of ambient particle radioactivity by air pollution on lung function in an elderly cohort. *Environmental Research*, 170, 345-353.
12. Zheng, Y., Liu, F., & Hsieh, H. P. (2015). Forecasting fine-grained air quality based on big data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 2267-2276). ACM.
13. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
14. Vieira, C. L., Koutrakis, P., Huang, S., Grady, S., Hart, J. E., Coull, B. A., ... & Garshick, E. (2019). Short-term effects of particle gamma radiation activities on pulmonary function in COPD patients. *Environmental research*, 175, 221-227.
15. Gan, W. Q., FitzGerald, J. M., Carlsten, C., Sadatsafavi, M., & Brauer, M. (2013). Associations of ambient air pollution with chronic obstructive pulmonary disease

hospitalization and mortality. *American Journal of Respiratory and Critical Care Medicine*, 187(7), 721-727.

16. Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., & Stewart, W. F. (2016). RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems* 29 (pp. 3504-3512). Curran Associates.
17. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211-407.
18. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54-63.
19. Cinelli, G., Tollefsen, T., Bossew, P., Gruber, V., Bogucarskis, K., De Felice, L., & De Cort, M. (2019). The European Atlas of Natural Radiation: state of the art and future perspectives. *Journal of Environmental Radioactivity*, 196, 240-252.
20. Hasenfratz, D., Saukh, O., Walser, C., Hueglin, C., Fierz, M., & Thiele, L. (2015). Deriving high-resolution urban air pollution maps using mobile sensor nodes. *Pervasive and Mobile Computing*, 16, 268-285.