Verifying the Measurement Chain for Particle-Associated Gamma Exposure in COPD Studies

Simon Crawford, Vincent Barker, Dominic Warren

Abstract

Environmental gamma-activity studies require a sequence of transformations between instrument counts and the exposure variable used in a health model. Each transformation changes the interpretation and uncertainty of the final association. This article develops a measurement-chain framework using a COPD study of short-term particulate matter gamma activity as its motivating case. We distinguish monitor response, estimated background, particle-associated interpretation, home-exposure assignment, temporal averaging, and biomarker comparison. Particular attention is given to negative background-adjusted values, indoor estimates derived from a particulate tracer, and the limits of assuming that all exposure error attenuates an association. The article is a methodological synthesis, not a new monitor calibration, patient analysis, or radiation-dose calculation. It proposes a verification record in which every derived value retains its source, units, transformation, and validation status. Efficient adaptation of a document-processing model may help maintain that record, but the model cannot replace physical calibration or validate an exposure construct through language alone.

Keywords: gamma monitoring; exposure measurement; COPD; background adjustment; measurement error; provenance.

1. The exposure variable is the end of a pipeline

A health model often receives an exposure column that looks simpler than the process that produced it. A number indexed by day and participant can conceal instrumental response, environmental transport, baseline estimation, and assumptions about where a person spent time. Short-term exposures to particulate matter gamma radiation activities and biomarkers of systemic inflammation and endothelial activation in COPD patients is a useful case for examining this chain (S. Huang et al., 2020a). The scientific question is not merely whether the final column is associated with a biomarker. It is what physical and statistical quantity that column represents after each transformation.

The motivating study uses a background-adjusted ambient signal and estimates an indoor exposure component with a particulate tracer relationship. It then relates temporally summarized exposures to biomarkers measured during repeated visits. These features create distinct validation problems. A successful fit of the health model does not independently validate the background estimate or the indoor assignment. Likewise, an association between an adjusted gamma signal and particulate mass does not identify the contribution of every radionuclide or establish a personal absorbed dose. The present article examines the logic of those links and proposes evidence that would make them reviewable.

This is a methodological synthesis rather than a reanalysis of clinical or monitor observations. We do not download a continuous RadNet time series, fit a background-removal algorithm, estimate infiltration factors, or calculate a new health association. The framework complements a separate descriptive audit of published quartiles, but no numerical result from that audit is treated as a validation of the underlying exposure construct. A variable can have a perfectly reproducible distribution while still being an imperfect representation of the exposure relevant to a biological mechanism. Reproducibility and construct validity are related but distinct requirements.

2. Instrument counts and physical quantities

The U.S. Environmental Protection Agency describes RadNet as an environmental radiation monitoring system with near-real-time gamma monitoring among its public resources (U.S. EPA, 2026). An instrument's recorded counts are a response to incoming radiation under a particular detector configuration and measurement process. They are not interchangeable with activity in a specified volume, energy absorbed by tissue, or an individual's dose. Those quantities require additional calibration and physical information. A health analysis should therefore retain the instrument's original units and explain every conversion, rather than allowing familiar radiation terminology to imply that all of these quantities have already been measured.

Energy bands add information but also require careful interpretation. A band label identifies a part of the instrument's response space under its reporting convention. Its counts can be influenced by detector efficiency, background sources, and environmental conditions. A high count in one band cannot be compared with a count in another as though the instrument had identical response and the biological implications were equal. Identifying specific radionuclides would require evidence beyond a broad-band association. A proposed particle-related measure should state what source interpretation is supported and what source separation remains unresolved.

The verification record should begin with station identity, instrument configuration, sampling schedule, units, and quality flags. Changes in equipment, maintenance, missingness, or reporting procedures can create apparent environmental changes. A day with no valid measurement is different from a day with a low measurement. Aggregation should preserve the number and timing of valid observations, because a daily mean based on an unusual subset of hours can represent a different exposure period. These requirements are ordinary measurement bookkeeping, but they become scientifically consequential when daily differences are linked to short-term changes in physiology.

3. Subtracting background changes the reference point

Background adjustment estimates a component of the measured signal and removes it to emphasize variation attributed to another source. The motivating study uses a moving-background procedure, including a 21-day window selected in relation to particulate matter correlation. The resulting variable is conditional on that procedure. A negative adjusted value means the observed signal falls below the estimated background at that time; it does not represent negative physical radiation. This distinction should remain explicit in tables, extraction systems, and downstream models. Otherwise a mathematically ordinary residual can acquire an impossible physical interpretation.

Window choice determines which variation is considered background and which remains in the signal. A short window can remove local episodes that a longer window preserves; a long window can leave slowly changing background structure in the residual. If a centered window uses future measurements, the resulting historical exposure estimate also has a different information structure from a real-time estimate. Either may be legitimate for a particular retrospective question, but the direction and availability of observations should be documented. The procedure's output should not be described as directly observed particle activity merely because a background component has been subtracted.

Selecting a window to maximize correlation with particulate mass introduces another interpretive dependency. The adjustment is partly tuned to resemble the pollutant with which the final signal will be discussed. That can be useful for constructing a hypothesis-driven proxy, but it reduces the independence of a later correlation offered as validation. A stronger assessment would freeze the algorithm and evaluate it against additional reference measurements or a separate time interval. The proposed comparison should

include a clear target quantity and should avoid treating agreement with the tuning criterion as a new discovery about the constructed variable.

Uncertainty in the estimated baseline also carries into the adjusted signal. Treating the baseline as a known constant can understate uncertainty when the scientific question depends on small differences around it. Sensitivity analyses using documented alternative windows can describe algorithm dependence, although they do not alone establish which window is physically correct. The entire pattern should be reported rather than selecting the version with the strongest health association. If a health result appears only under one narrow background choice, that is relevant evidence about its robustness, even when the chosen version has a plausible rationale.

4. From a monitor to a home and a person

A central monitor measures conditions at a location, while a participant experiences a changing sequence of environments. Distance, atmospheric transport, building characteristics, ventilation, and time spent indoors can all affect the relationship. The motivating study estimates an indoor component using an indoor-to-ambient sulfur relationship. That step should be described as an exposure assignment based on a particulate tracer, rather than as a direct measurement of indoor gamma activity for every participant-day. The tracer may capture aspects of particle infiltration while leaving uncertainties about composition, source, and gamma-emitting constituents unresolved.

The validity of an assignment can vary across homes and periods. A ratio derived under one ventilation regime may behave differently when windows are open, filtration changes, or indoor sources contribute. A participant's location during the relevant window adds another layer. A home-based estimate may represent residential conditions well while missing work, transport, or other environments. These are testable concerns rather than reasons to discard all assigned exposures. A validation subsample with appropriately chosen colocated measurements could estimate agreement and identify conditions under which the assignment performs poorly. The relevant comparison must still be made in compatible units and time windows.

Exposure measurement error has several forms. Classical error describes a measured value varying around a person's true exposure under particular assumptions. Berkson-type error describes individual exposures varying around an assigned value (Berkson, 1950). Real environmental assignments can contain mixtures of both, along with systematic error related to covariates or sources. General measurement-error frameworks show why the direction of bias depends on the model and error structure (Carroll et al., 2006). It is therefore unsafe to assume that every uncertainty in the exposure column simply attenuates the health coefficient toward zero. That claim requires assumptions that should be stated and assessed.

Environmental epidemiology has long distinguished these error structures in exposure assessment (Armstrong, 1998). Time-series air-pollution work also shows how shared temporal patterns and assignment errors interact with inference (Zeger et al., 2000). The relevance here is methodological: a central-monitor proxy may contain useful common temporal information while imperfectly representing an individual's physical exposure. A validation design should estimate the features of error that matter for the intended health model. A high correlation between two ambient signals alone cannot characterize person-specific error or justify a universal correction factor.

5. Time windows belong to the biological claim

A biomarker measured at a visit can reflect processes operating over several time scales. A same-day exposure, a preceding-week average, and a distributed lag are different explanatory variables. Choosing among them should draw on a plausible temporal mechanism and the sampling schedule. A daily monitor value cannot identify the precise time at which an inflammatory response began. Nor does a stronger association for a longer average prove a cumulative biological mechanism, because averaging also changes measurement noise and temporal correlation. The exposure window is part of the estimand and should remain attached to every reported coefficient.

Distributed lag models offer a framework for describing associations across time without reducing every question to one arbitrary average (Gasparrini et al., 2010). They still require decisions about lag range, functional form, and the amount of information available to estimate that structure. A small repeated-visit panel can support only limited complexity. Searching many windows and retaining the most favorable one produces uncertainty beyond the standard error of the selected coefficient. A prospective plan should identify a primary window and preserve alternatives as sensitivity analyses, with their full result pattern rather than only their significant entries.

Temporal confounding is also a measurement and design problem. Season, weather, infections, behavior, and co-pollutants can vary on overlapping schedules. A causal diagram can make the assumed pathways and adjustment goals explicit (Greenland et al., 1999). Adjusting for particulate mass may address one question while introducing sensitivity to correlated measurement errors or changing the interpretation of the gamma coefficient. A numerically stable two-pollutant model does not prove that the radiation-related component acts independently of every other particle property. The model conditions on measured covariates; it does not experimentally separate the environmental mixture.

6. Verification should follow the chain in both directions

A practical verification record would let a reader trace a health-model row backward to the daily assigned exposure, the aggregation window, the background-adjusted station values, and the valid original measurements. Each transition should carry units, timestamps, parameter settings, missing-data rules, and a version identifier. The same record can be read forward to understand how a maintenance gap or a baseline revision changes the final analytic sample. This bidirectional traceability is more useful than a single final CSV whose column names no longer reveal how its numbers were constructed.

Validation experiments should target the weakest physical or statistical link. Instrument calibration addresses detector response. Colocated sampling addresses spatial agreement. Indoor reference measurements address the home-assignment model. Personal monitoring, where technically appropriate, addresses another level of exposure representation. Repeating a health regression under several covariate sets is not a substitute for any of these measurements. Conversely, a well-calibrated instrument does not eliminate confounding in a health study. The framework encourages these contributions to be reported together while preserving the distinct question that each one can answer.

Negative controls or related diagnostic outcomes can sometimes challenge an explanation, but they require their own assumptions. A control chosen after observing the primary association may not provide the intended independent test. The same caution applies to comparing gamma bands that share sources or construction steps. A useful falsification strategy states in advance why the proposed control should respond differently under competing explanations. The goal is to reduce the number of plausible interpretations through observable consequences, rather than to accumulate reassuring analyses whose relationship to the exposure mechanism remains unspecified.

7. Language models as record keepers

Maintaining provenance across papers, instrument documentation, tables, and code can become labor-intensive. A document-processing model could assist by extracting units, windows, transformation parameters, and statements about whether exposure was measured or assigned. Low-rank adaptation is relevant to limiting the trainable updates required for a specialized extractor (Hu et al., 2022). Hadamard Tensor Ring provides a candidate efficiency approach for such fine-tuning (H. Tong et al., 2026). This is a proposed application only. No model was adapted or benchmarked in this article, and computational efficiency would not validate the physical interpretation of the extracted fields.

The extractor should preserve uncertainty rather than resolve it through fluent wording. For example, it should retain a distinction between background-adjusted counts and dose, between a residential assignment and a personal observation, and between a tuning correlation and independent validation. Every field should point to a supporting passage or table. Human review can then focus on disagreements and missing links. A benchmark should include these difficult distinctions and score them separately from ordinary number copying. An efficient system that erases qualifiers can make a research record easier to read while making its scientific claims less accurate.

8. Interpretation at the end of the chain

The motivating COPD research identifies an important observational question at the intersection of environmental monitoring and inflammatory biology. Its interpretation depends on a constructed exposure measure whose components deserve direct inspection. The framework developed here does not supply a new effect estimate, a radionuclide attribution, a clinical threshold, or advice about patient exposure. It identifies the records and comparisons needed to understand what a measured association could mean. That is a substantive contribution when the alternative is to let a convenient column name conceal several untested physical and statistical assumptions.

Background subtraction, home assignment, and temporal averaging are models embedded before the health regression begins. Their outputs can be useful and reproducible while remaining uncertain representations of the exposure of interest. A stronger evidence chain preserves that uncertainty, evaluates each transformation against an appropriate reference, and limits the final claim to what the chain supports. The resulting study is easier to reproduce and harder to overinterpret. Most importantly, it makes clear which additional observation would meaningfully strengthen the connection between an environmental signal and a biological response.

References

Huang, Shaodan, et al. “Short-term exposures to particulate matter gamma radiation activities and biomarkers of systemic inflammation and endothelial activation in COPD patients.” *Environmental Research*, vol. 180, 2020a, pp. 108841. https://doi.org/10.1016/j.envres.2019.108841.

Tong, H., Xu, G., & Chen, Y. (2026, May). Hadamard Tensor Ring for Efficient Low-Rank Fine-Tuning. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

U.S. Environmental Protection Agency. “RadNet.” *EPA*, updated 5 Oct. 2026. https://www.epa.gov/radnet. Accessed 22 Oct. 2026.

Berkson, Joseph. “Are there Two Regressions?.” *Journal of the American Statistical Association*, vol. 45, no. 250, 1950, pp. 164-180. https://doi.org/10.1080/01621459.1950.10483349.

Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. “Measurement Error in Nonlinear Models.” 2006. https://doi.org/10.1201/9781420010138.

Armstrong, B G. “Effect of measurement error on epidemiological studies of environmental and occupational exposures..” *Occupational and Environmental Medicine*, vol. 55, no. 10, 1998, pp. 651-656. https://doi.org/10.1136/oem.55.10.651.

Zeger, S L, et al. “Exposure measurement error in time-series studies of air pollution: concepts and consequences..” *Environmental Health Perspectives*, vol. 108, no. 5, 2000, pp. 419-426. https://doi.org/10.1289/ehp.00108419.

Gasparrini, A., B. Armstrong, and M. G. Kenward. “Distributed lag non-linear models.” *Statistics in Medicine*, vol. 29, no. 21, 2010, pp. 2224-2234. https://doi.org/10.1002/sim.3940.

Greenland, Sander, Judea Pearl, and James M. Robins. “Causal Diagrams for Epidemiologic Research.” *Epidemiology*, vol. 10, no. 1, 1999, pp. 37-48. https://doi.org/10.1097/00001648-199901000-00008.

Hu, Edward J., et al. “LoRA: Low-Rank Adaptation of Large Language Models.” *International Conference on Learning Representations*, 2022. https://openreview.net/forum?id=nZeVKeeFYf9.