Towards Autonomous Environmental Health: AERIS, a Research Agent for Air-Pollution Exposure Assessment

Romain Lacombe[1] **Ziv Lautman**[1] **Nadav Efrat**[2] **Yougeng Lu**[1] **Michael P. Snyder**[1]
[1]Stanford University [2]University of Illinois Urbana-Champaign
{rlacombe, lautman, yougengl, mpsnyder}@stanford.edu, nefrat2@illinois.edu

Abstract

Air pollution harms human health and contributes to millions of deaths each year, yet many public-health studies omit it even when they record the locations and dates needed to estimate exposure. We introduce AERIS (Airborne Exposure Research & Investigation System), a research harness that gives coding agents the domain tools to link cohort locations and time windows to public air-quality records. To evaluate the system on a complete scientific workflow, we introduce *source-substitution reanalysis*: reconstructing a published study's pollution exposure from independent public measurements, linking it to health outcomes, and re-estimating the health effects. Screening 18,919 papers yielded nine published studies with open health data, reproducible statistical models, and sufficient air-quality coverage. Across these studies, the pipeline processed 1.85 million monitor observations into 20,568 area-time exposure estimates and 14 pollution–health comparisons across four pollutants. With researcher review, reconstructed estimates matched the direction of all seven published positive associations, with a median study-level log-effect difference of 0.040. Across three autonomous campaigns, self-driving AERIS completed 23 of 27 reanalyses compared with 7 of 27 for the same model without the harness. For any study with location and time records, AERIS turns air-pollution exposure assessment into an auditable agent workflow, charting a path toward autonomous environmental health.

1 Introduction

Air pollution is among the most consequential environmental threats to health. The 2024 *State of Global Air* assessment attributed 8.1 million deaths in 2021 to ambient and household air pollution [Health Effects Institute, 2024]. Decades of epidemiological research link short- and long-term exposure to cardiovascular and respiratory disease, hospitalization, and mortality [Dockery et al., 1993, Di et al., 2017, Brook et al., 2010, Brunekreef and Holgate, 2002]. Its relevance extends well beyond studies centered on pollution. Studies of weather, infectious disease, cardiovascular care, mental health, or clinical recovery can yield biased estimates when air pollution confounds the relationship between the primary exposure and the health outcome. Measuring air-pollution exposure can reduce this confounding, sharpen the relationship under study, and reveal health effects and biological pathways that warrant further investigation.

Many public-health studies already record residential addresses, ZIP codes, administrative areas, observation dates, or GPS traces from smartphones and wearables. Converting these records into pollution exposure requires technical skills and domain expertise. Investigators must select a data source, geographic assignment method, temporal aggregation, and completeness standard, then document those choices for review. The research question at hand determines the appropriate air-quality sources and models: a residential study may use monitor-based exposure or a hyperlocal dispersion model at each participant's home, whereas a study of daily activity may require estimates

GenAI for Health Workshop @ NeurIPS 2026, Sydney.

along each person's route [Nyhan et al., 2019, Steinle et al., 2013]. Government monitors provide trusted observations but may not align with where and when people are exposed [Liang and Kumar, 2013]; low-cost sensors add local detail but require calibration [Morawska et al., 2018, Barkjohn et al., 2021]; and models fill geographic gaps through assumptions that affect resolution and accuracy [Hoek et al., 2008, van Donkelaar et al., 2010, Di et al., 2019, Jauvion et al., 2020].

Researchers rebuild these exposure-assessment pipelines for every new study, often across disconnected scripts and data services. This engineering overhead delays statistical analysis, deters investigators from including air-pollution covariates even when location and date records are available, and slows research into how environmental exposures shape disease and recovery. Automating air-pollution exposure assessment would free scientists from routine data engineering, make environmental covariates practical to include across clinical and public-health cohorts, and let researchers test hypotheses that would otherwise go unexplored. To date, however, decisions about data sources, geographic coverage, and fallback specifications have required close researcher oversight, raising a direct question: how much of this workflow can an autonomous agent execute reliably on its own, and where does researcher guidance still matter?

AERIS (Airborne Exposure Research & Investigation System) is a research harness that helps investigators add air-pollution exposure to public-health studies through their preferred command-line agent, either with interactive researcher guidance or in a fully autonomous self-driving mode. Given a study's pollutant, locations, and dates, AERIS retrieves public air-quality data, computes exposure, assesses coverage, and prepares an auditable result for analysis.

Long-running research agents are notoriously difficult to evaluate through isolated question-and-answer benchmarks or bounded, limited-complexity tasks. We introduce the *source-substitution reanalysis* task as an end-to-end, outcome-based benchmark: a chosen scientific study provides context, and the agent is asked to independently recover the study's outcome and location data, reconstruct air-quality exposure from public sources, and reproduce the original analysis. We screened 18,919 PubMed and OpenAlex records, identified 21 candidates, and attempted to reproduce the 14 with sufficient public inputs and methods. Nine had retrievable outcomes and adequate historical air-quality coverage; three others lacked adequate coverage or spatial variation, and two lacked a retrievable outcome export. Comparing the reconstructed and published health-effect estimates evaluates the entire scientific workflow end to end, from data retrieval and exposure computation to statistical modeling and provenance reporting [Chen et al., 2025, Siegel et al., 2024]. Our primary evaluation reflects intended use: a scientist supervises AERIS as it conducts each reanalysis. We then run two autonomous experimental controls, each repeated across three independent nine-study campaigns: self-driving AERIS measures robustness without researcher oversight, while a matched raw-agent ablation isolates the contribution of the harness beyond the underlying model.

Our contributions are: (i) a research-agent harness for computing exposure at fixed locations, along mobility trajectories, and across administrative areas; (ii) source-substitution reanalysis as an outcome-based evaluation, applied to nine published studies with reviewable exposure, analysis, and provenance artifacts; and (iii) autonomous robustness and raw-agent ablation controls that measure self-driving reliability and isolate the contribution of the harness.

2 Related work

Agents for science and reproducible research. Scientific agents connect language models to specialized software, data resources, and laboratory interfaces. Coscientist and ChemCrow established this pattern in chemistry [Boiko et al., 2023, Bran et al., 2024]; later systems extend it to biomedical research and data analysis [Gottweis et al., 2026, Bu et al., 2026], while research-orchestration harnesses such as Distillate coordinate long-horizon experiment loops across command-line coding agents [Lacombe, 2026]. Executable benchmarks such as MLE-bench evaluate complete agent workflows and show that model scaffolding materially affects performance [Chan et al., 2025]. AERIS extends this approach to scientific reanalysis, where success also requires valid data reconstruction, provenance, and agreement with external findings across independent trajectories. Yet ScienceAgentBench and CORE-Bench show that long-horizon, data-driven work remains difficult [Chen et al., 2025, Siegel et al., 2024], while reproducibility studies find that public releases rarely capture every decision needed to rebuild an analysis [Stodden et al., 2018, Seibold et al., 2021]. Source-substitution

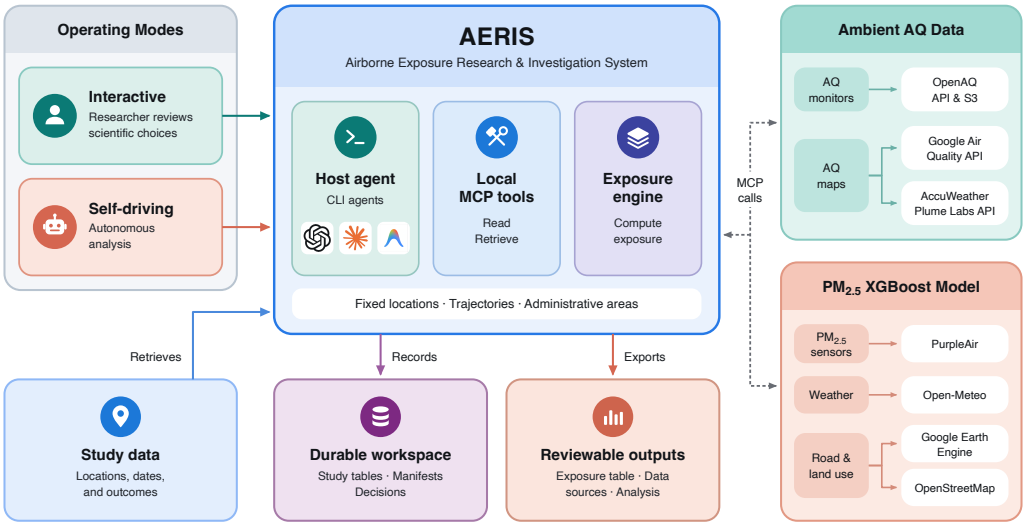


Figure 1: AERIS architecture. In interactive or self-driving mode, a command-line agent loads study data, queries ambient sources, and computes exposure through the AERIS skill and local tool server. Both modes share a workspace that preserves the inputs, decisions, coverage, provenance, and outputs required for scientific review.

reanalysis complements task-based benchmarks by evaluating a complete workflow against external scientific findings while requiring the agent to reconstruct its own inputs.

Exposure assessment for epidemiology. Because regulatory monitors are sparse, epidemiologists estimate exposure at unmonitored locations by interpolating station readings, fitting land-use and dispersion models, or fusing ground monitors with satellite observations [Jerrett et al., 2005, Ryan and LeMasters, 2007, Kloog et al., 2014]. Large cohorts assign modeled concentrations to geocoded residences [Eeftens et al., 2012, Szpiro et al., 2010]; mobility-aware studies link GPS trajectories to time-varying pollution fields [Nyhan et al., 2016, Yu et al., 2020]; and hyperlocal models combine sensors, weather, and land-use features to resolve street-level gradients between stations [Jauvion et al., 2020, Alléon et al., 2020, Cassard et al., 2020]. Existing tools support individual steps such as residential histories, personal tracking, and health-impact assessment [Durney et al., 2026, Donaire-Gonzalez et al., 2019, Breen et al., 2019, Sacks et al., 2018], but investigators must still select sources, align scales, enforce coverage rules, and document provenance for each study. AERIS streamlines that work, turning cohort locations and timelines into auditable exposure data tailored to the research question.

Environmental health evidence and study design. Landmark cohort and multi-city studies link particulate air pollution to all-cause and natural-cause mortality [Dockery et al., 1993, Samet et al., 2000, Beelen et al., 2014], cardiopulmonary mortality [Pope et al., 2002, Di et al., 2017], and cardiovascular and respiratory hospital admissions [Dominici et al., 2006]. Systematic reviews and meta-analyses confirm these associations across major pollutants ($PM_{2.5}$, PM_{10}, NO_2, and O_3) and across adult mortality and pediatric respiratory outcomes [Orellano et al., 2020, Khreis et al., 2017]. Because every effect estimate depends on exposure windows, spatial boundaries, coverage thresholds, and covariates, AERIS records these exposure-assessment decisions alongside the resulting estimates so investigators can audit and refine the analysis.

3 AERIS system design and exposure assessment

3.1 What a usable exposure record requires

A reviewable exposure record pairs each pollutant estimate with the data and methodological choices that produced it. AERIS starts from a cohort table, mobility trace, or administrative boundary and

supports three spatial modes: fixed locations (such as participant residences or clinic sites), time-stamped trajectories (from GPS or wearable logs), and administrative areas (such as counties or municipalities).

Before querying external data, the researcher or agent specifies the target pollutants, study timeframe, spatial resolution, spatial matching method (such as nearest station, inverse-distance weighting, or area averaging), and minimum temporal data coverage. AERIS retrieves observations from public air-quality repositories, links them to the study's locations and dates, aggregates them to match the health outcome data, and exports an analysis-ready exposure table alongside a manifest recording the contributing stations or models, temporal coverage, and retrieval decisions.

3.2 How AERIS works

AERIS separates scientific reasoning from reproducible computation across four components. This separation of concerns lets research teams use their preferred command-line agent while preserving a stable exposure workflow and reviewable study record.

1. **Host agent and AERIS skill.** The researcher starts AERIS with `aeris.sh` and selects Codex, Claude Code, or Antigravity. The launcher opens that host in the selected study directory, supplies the AERIS workflow instructions, and connects the local tool server. The host manages conversation, planning, and tool selection.
2. **Local tool server and exposure tools.** AERIS uses the Model Context Protocol (MCP) to expose typed tools over local stdio. Any compatible host can work with public air-quality repositories and study data through one interface. Its tools inspect inputs, approve exposure requests, retrieve observations, assign exposure, diagnose coverage, export outputs, and retain decisions. Shared definitions keep these operations consistent across hosts and independent of chat history.
3. **Durable study workspace.** Loading a cohort creates a local `.aeris` workspace. Input tables, exposure outputs, and run manifests receive stable identifiers. The host can inspect schemas, sample rows, missing values, value counts, and numeric summaries. Each run manifest records inputs, parameters, source settings, coverage, and decisions. Following the external-state principle of recursive language models [Zhang et al., 2025], the workspace also gives subagents shared state for bounded inspection and monitoring tasks.
4. **Exposure computation layer.** AERIS performs geocoding, boundary resolution, monitor discovery, retrieval, spatial and temporal aggregation, quality checks, and export. It queries OpenAQ for monitoring data and falls back to OpenAQ's public S3 archive for historical dates or bulk retrieval. Request pacing, workload planning, and a local SQLite cache prevent duplicate queries and allow interrupted downloads to resume without data loss.

A study begins with a research question and a cohort or area table. In interactive mode, the researcher reviews the proposed study design and approves the exposure request before retrieval begins. In self-driving mode, the agent operates within fixed safety and resource limits: it locks the study's target question, statistical model, and exposure contrast, records any material substitutions, and critiques its own analysis before finalizing. Failure modes observed during interactive use informed the autonomous recovery rules—retrying after failed queries, probing alternative public archives, and explicitly labeling methodological departures. Because both modes share the same exposure engine and workspace, shifting from interactive to autonomous execution changes only who approves scientific decisions, not how exposure is computed or documented (Figure 1).

3.3 Data sources, retrieval, and quality control

AERIS links study locations and dates to modular exposure sources detailed in Appendix C. Its OpenAQ pathway retrieves station metadata and hourly or daily measurements through the API and Python SDK (version 0.4.0), with the public S3 archive as a fallback for historical data and bulk downloads [OpenAQ, 2026a, Biggs and OpenAQ, 2025, OpenAQ, 2026b]. Google Maps provides recent, coordinate-level hourly histories, while the AccuWeather Plume Labs API provides current and forecast air-quality indices [Google Maps Platform, 2026, AccuWeather, 2026, Jauvion et al., 2020]. For reproducible hyperlocal $PM_{2.5}$ estimates, a packaged XGBoost model extends the random-forest method of Lu et al. [2021] by combining nearby low-cost sensor readings with weather, road

proximity, land cover, and day of the week. Researchers can also substitute their own exposure model while retaining AERIS's assignment, quality-control, and provenance workflow.

Before retrieval, AERIS validates identifiers, dates, and locations, resolving ZIP codes with `pgeocode` [Yurchak, 2024], addresses with Nominatim [Nominatim developer community, 2026], and administrative units with geoBoundaries [Runfola et al., 2020]. Exported tables retain the contributing stations or boundaries, spatial matching settings, temporal coverage, data gaps, and underlying daily or hourly measurements.

3.4 Exposure computation pathways

Airborne pollutant exposure is estimated through three pathways depending on the study's spatial and temporal resolution. All three share the same source selection, retrieval, coverage validation, provenance, and export workflows.

Fixed-location participant exposure. The fixed-location pathway accepts a participant identifier, ZIP code, address, or coordinates with an exposure window. It supports studies where a residence, school, clinic, or other stable setting represents the relevant location. AERIS geocodes each record, finds candidate monitors, and assigns either the nearest monitor or an inverse-distance-weighted estimate:

$$C_p = \frac{\sum_{i \in M_p} w_i C_i}{\sum_{i \in M_p} w_i}, \qquad w_i = d_i^{-\beta},$$

where M_p is the set of candidate monitors for participant p. The output identifies the contributing monitors, distances, exclusions, coverage, and aggregation rule. When the required features are available, the packaged XGBoost model can estimate $PM_{2.5}$ concentrations at fixed coordinates over time.

Trajectory-based participant exposure. The trajectory pathway accepts time-stamped coordinates from GPS traces, travel diaries, or similar records. AERIS validates coordinates and timestamps, converts time to UTC, removes duplicate participant-hours, orders each trajectory, and estimates conditions at every retained location and time. The XGBoost pathway loads packaged weights on demand and combines nearby sensor, weather, road-network, and land-cover features. Its output preserves both the location-time variation and source metadata.

Administrative-area exposure. The administrative-area pathway accepts a named unit or researcher-provided polygon with an exposure window. It supports ecological studies whose outcomes are indexed by municipality, county, department, region, or country. AERIS resolves the boundary, selects monitors within it, aggregates daily measurements, and records the polygon source, contributing stations, coverage, and aggregation method.

4 Outcome-Based Evaluation Using Published Studies

4.1 Experimental design

Long-running research agents are notoriously difficult to evaluate through isolated question-and-answer benchmarks or bounded, limited-complexity tasks [Chen et al., 2025, Siegel et al., 2024]. To evaluate AERIS on a complete scientific workflow, we test it on published public-health studies. Each article defines the pollutant, health outcome, study period, geographic unit, and statistical model. Without access to the original authors' exposure file, AERIS retrieves independent public measurements of the target pollutant, constructs a new exposure record, joins it to public outcome data, and fits the specified model. This *source-substitution reanalysis* tests the entire workflow against an external scientific result while exposing failure modes that bounded tasks miss, including incomplete historical coverage, mismatched geographic units, and unsupported substitutions.

We count a study as completed when the agent retrieves usable health outcomes, constructs an exposure record, and produces an identifiable fitted effect. Completion measures end-to-end execution, whereas direction, null recovery, and effect magnitude measure scientific agreement. For every run, we preserve the article, data sources, tool trace, manifest, analysis panel, and fitting script.

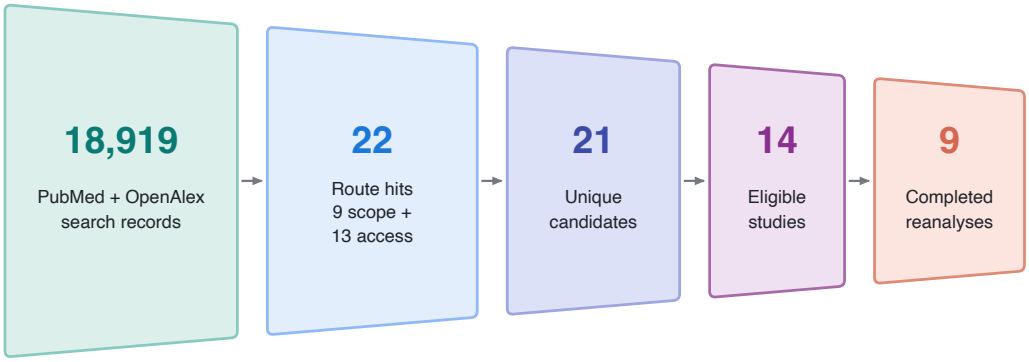


Figure 2: Study-selection funnel from 18,919 deduplicated PubMed and OpenAlex records to 21 audited candidates, 14 attempted reanalyses, and nine completed fits (see Appendix A for the stage-by-stage breakdown).

4.2 Study-selection and reanalysis pipeline

We searched five years of PubMed and OpenAlex for outdoor air-pollution and human-health studies across two complementary screening routes (Figure 2): one prioritized scientific scope and the other publicly accessible health outcomes. After removing one overlap, we audited 21 candidates and attempted the 14 with sufficient public data and methodological detail; nine had retrievable outcomes and adequate historical air-quality coverage for a fitted reanalysis (all using administrative-area outcomes). COVID-19 studies dominate because governments released cases and deaths with the dates and geographic keys needed for exposure reconstruction, whereas comparable public outcome exports remain rare for chronic disease and hospital admissions. GPT-5 Codex subagents audited each article and named source on 13–14 August 2026; Appendix A and the supplementary material provide the full screening ledger and exclusion records.

4.3 Evaluation conditions and autonomous controls

Our primary evaluation reflects intended use: AERIS retrieves data, builds exposure panels, and fits models while a researcher reviews its scientific choices (Appendix F). Two autonomous controls then execute the same nine study targets in fresh workspaces without human intervention, each repeated across three independent campaigns using Codex CLI with GPT-5.6 Terra at medium reasoning effort: (1) *self-driving AERIS* operates with the full harness, standing permissions, retry policies, and internal critique; and (2) a *raw-agent ablation* runs Codex with the article bundle, task contract, network access, and output schema, but without AERIS tools, prior code, or results.

Because effect metrics apply only to runs that produce comparable estimates, we report completion alongside retention of the published direction or null conclusion, absolute log-effect difference, geometric mean reconstructed-to-published ratio, and campaign-level Pearson correlation (Appendix E).

5 Results

5.1 AERIS with researcher review

Across the nine completed reanalyses (Table 1), AERIS processed 1,847,519 air-quality monitor records across 705 monitor-study associations and produced 20,568 area-time exposure estimates. These analyses covered 14 effect contrasts, four pollutants, and administrative units across Europe and the Americas.

AERIS recovered the direction of all seven published positive contrasts, with three reconstructed confidence intervals excluding 1 (Figure 3). Among the seven null or no-evidence contrasts, four AERIS intervals included 1 (all four OECD endpoints), whereas the Brazil cancer-mortality and two Northern Italy estimates excluded 1; Appendix E places each contrast beside its source and reconstruction settings.

Table 1: Primary evaluation and autonomous controls. Completion is reported as mean ± SD across campaigns, with pooled attempts in parentheses. Agreement metrics include only completed, comparable estimates and should be read alongside completion rates. Pearson correlations are listed by campaign; an em dash marks unavailable logging or fewer than two comparable estimates. Token counts include all retries.

Measure	AERIS with researcher review	Self-driving AERIS	Raw agent
Role	Primary evaluation	Robustness control	Harness ablation
AERIS tools and context	Yes	Yes	No
Human researcher intervention	Occasional	None	None
Independent campaigns	1	3	3
Study attempts	9	27	27
Fitted studies per campaign	9 (9 / 9)	7.7 ± 1.5 (23 / 27)	2.3 ± 1.2 (7 / 27)
Complete nine-study campaigns	1 / 1	1 / 3	0 / 3
AERIS provenance record	9 / 9	26 / 27	0 / 27
Comparable study-level estimates	8	17	7
Published positive direction retained	7 / 7	11 / 16	3 / 4
Published null conclusion retained	4 / 7	14 / 21	3 / 3
Median $\lvert\Delta \log R\rvert$	0.040	0.169	0.073
Geometric mean ratio	0.998	0.901	1.011
Pearson r by campaign	0.834	$0.507 \mid 0.091 \mid -0.170$	$0.016 \mid -0.675 \mid$ —
Campaign runtime, hours	—	2.10 ± 0.52	1.02 ± 0.01
Model output tokens per campaign	—	$275\text{k} \pm 52\text{k}$	$103\text{k} \pm 5\text{k}$

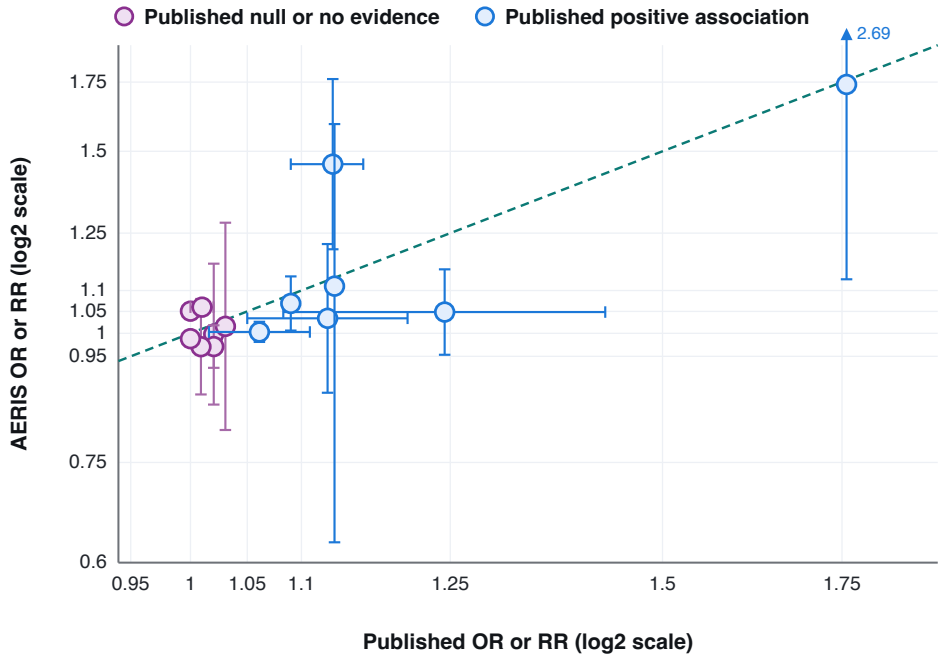


Figure 3: Published and AERIS effect ratios for all 14 contrasts from reanalyses conducted with researcher review, shown on $\log_2$ axes. Error bars show available 95% CIs; colors distinguish positive from null or no-evidence contrasts. The identity line marks exact agreement.

For numerical agreement, Table 5 gives each study equal weight. For a study with k contrasts, we compute its mean absolute log-effect difference and signed mean log ratio:

$$e_s = \frac{1}{k}\sum_{c=1}^{k}\left|\log\left(\frac{R_{\mathrm{AERIS},c}}{R_{\mathrm{published},c}}\right)\right|, \qquad d_s = \frac{1}{k}\sum_{c=1}^{k}\log\left(\frac{R_{\mathrm{AERIS},c}}{R_{\mathrm{published},c}}\right),$$

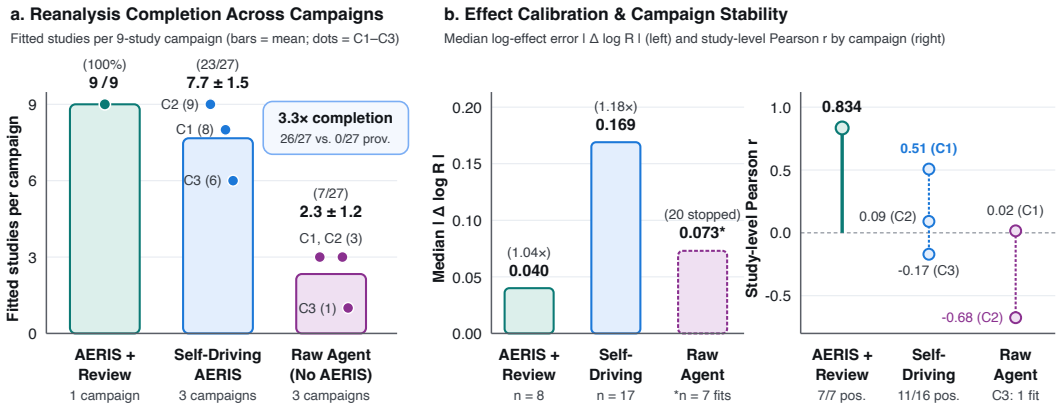


Figure 4: Comparison of AERIS with researcher review, self-driving AERIS, and the raw-agent ablation across nine benchmark studies. **(a)** Fitted studies completed per nine-study campaign, with points marking individual campaigns C1–C3. **(b)** Pooled median study-level absolute log-effect difference $|\Delta \log R|$ among comparable fits (left) and study-level Pearson correlation r by campaign (right).

where $\exp(d_s)$ yields the geometric mean reconstructed-to-published ratio, and averaging within studies prevents the four OECD contrasts from dominating.

Excluding the two Northern Italy contrasts reported as 1.00, the median study-level absolute log-effect difference across the remaining eight studies was 0.040, the geometric mean AERIS-to-published ratio was 0.998, and six estimates were within 10% on the ratio scale. Published and reconstructed estimates had a study-level Pearson correlation of $r = 0.834$ and a rank correlation of Spearman $\rho = 0.548$ ($\rho = 0.637$ across all 14 contrasts). Leave-one-study-out Pearson r remained positive for every dropped study ($r \in [0.401, 0.954]$); the Italian cities study exerts higher linear leverage only because its published contrast is scaled per doubling of PM_{10} exceedance days (1.76 vs. 1.74) rather than per 1 or 10 µg/m^3. Because ratio-scale proximity can still reflect a change in excess risk (for São José ozone, the RR fell from 1.061 to 1.003), we evaluate direction, null recovery, log-scale calibration, and correlation jointly.

5.2 Autonomous controls

Figure 4 compares completion and calibration across all three conditions; Appendix G and Appendix H report the individual self-driving and raw-agent runs.

Self-driving AERIS fitted 8, 9, and 6 studies across the three campaigns: 23 of 27 attempts, or 7.7 ± 1.5 per campaign. One campaign completed all nine studies, and 26 of 27 attempts retained an AERIS exposure and provenance record. Among fitted, interpretable contrasts, self-driving AERIS retained 11 of 16 published positive directions and 14 of 21 null conclusions, with a median study-level absolute log-effect difference of 0.169, a geometric mean ratio of 0.901, and campaign Pearson correlations of 0.507, 0.091, and -0.170. Campaigns took 2.10 ± 0.52 hours and generated $275{,}000 \pm 52{,}000$ output tokens, including retries.

The raw-agent harness-ablation campaigns fitted 7 of 27 attempts (3, 3, and 1), compared with 23 for self-driving AERIS, a 3.3-fold difference. Six of the seven fits retained the published conclusion (median absolute log-effect difference 0.073; geometric mean ratio 1.011; campaign correlations 0.016 and -0.675 in the two campaigns with multiple fits), while the raw agent produced no estimate in the other 20 attempts. The ablation shows that AERIS helps the agent recover from incomplete public inputs and finish analyses that the raw model usually abandons.

6 Discussion

With researcher review, AERIS completed all nine feasible reanalyses, recovered every published positive direction, and reached a median log-effect difference of 0.040. Agreement across four

pollutants, several countries, and varied health outcomes shows that the exposure pipeline can preserve broad spatial and temporal patterns under source substitution. It does not imply 96% accuracy in recovering a health effect: a small ratio-scale difference can represent a large change in excess risk. As an evaluation method, source substitution yields a weight of evidence across direction, uncertainty, calibration, and study-level failures rather than a single benchmark score.

Self-driving AERIS completed 23 of 27 attempts, compared with 7 for raw Codex, and preserved provenance in 26. Yet its median log-effect difference rose to 0.169, while study-level correlations ranged from -0.170 to 0.507 across campaigns. Researchers still improve source selection, estimand fidelity, and calibration. This gap separates two requirements that completion scores often conflate: a harness can help an agent finish a long analysis, while scientific judgment still determines whether it answers the intended question precisely and accurately. A practical path toward autonomy preserves review at consequential branch points while automating retrieval, computation, recovery, and documentation.

Our autonomy results remain specific to this evaluation. We designed the retry policies after observing failures in the nine-study campaign with researcher review, and every control used GPT-5.6 Terra. We have not yet tested whether the same policies transfer to new studies or other agents. The evaluation also favors COVID-19 and administrative-area analyses because their health outcomes were publicly accessible. Future work should test other agents, unseen disease settings, and private cohort data under appropriate research governance.

Scientific uncertainty remains in monitor placement, exposure-model validity, indoor exposure, and the design of each health study. The reanalyses evaluate the administrative-area pathway; fixed-location and mobility exposure received implementation tests only. Autonomous analyses should therefore preserve their sources, traces, and substitutions and receive expert review before informing research or policy.

7 Conclusion

Air pollution remains absent from many public-health studies because exposure assessment is technically demanding and difficult to audit. AERIS lowers this barrier by giving coding agents the domain tools to retrieve public air-quality data, compute exposure across participant and population settings, assess coverage, and preserve provenance. Evaluated through source-substitution reanalysis across nine published studies, the system processed 1.85 million monitor records, recovered every published positive direction, and reached a median absolute log-effect difference of 0.040, while autonomous controls completed 23 of 27 attempts compared with 7 for the raw agent.

For any study with location and time records, AERIS turns air-pollution exposure assessment into an auditable agent workflow, charting a path toward autonomous environmental health so that the air we breathe is no longer left unmeasured in medical and public-health research.

AI Use Statement

We used generative AI for research design, literature retrieval, software development, data preparation and analysis, interpretation, reanalysis execution, and manuscript drafting and editing. Codex ran the reanalyses with researcher review, autonomous AERIS, and the raw agent with GPT-5.6 Terra at medium reasoning; Gemini, Codex, and Claude Code also supported implementation and writing. We take full responsibility for all content and AI-assisted artifacts.

Code and Data Availability

The project page (`https://rlacombe.github.io/aeris/`) and appendices document the literature search, eligibility decisions, study targets, source substitutions, statistical comparisons, and agent configuration. The AERIS agent skill, Model Context Protocol (MCP) server, exposure tools, test suite, packaged XGBoost weights, and per-study reanalysis scripts are publicly available at `https://github.com/rlacombe/aeris`. The 18,919-record literature-screening ledgers, benchmark comparison tables, data-source attributions, and complete execution trajectories for the self-driving and raw-agent control runs are publicly available at `https://huggingface.co/`

datasets/rlacombe/aeris and mirrored in the repository's evals/ directory. To respect upstream provider terms, both archives provide automated retrieval scripts and SHA-256 run manifests rather than redistributing raw station observations.

Ethics Statement

This study reanalyzes public, aggregate, or previously released research data; we did not recruit participants or access private clinical records. Exposure estimates remain subject to monitor coverage, source quality, geographic aggregation, and the assumptions of each health study. AERIS therefore retains provenance and coverage diagnostics so that researchers can review its work before drawing scientific or clinical conclusions. The system does not provide medical advice, and its outputs require expert review before use in research or policy.

Acknowledgments

We thank OpenAQ for maintaining open programmatic access to global air-quality observations and open-source software tools. This work relies on measurements collected and published by national, regional, municipal, and scientific environmental monitoring agencies and sensor networks worldwide, including the U.S. Environmental Protection Agency AirNow program, the European Environment Agency, Arpae Emilia-Romagna, ARPA Lazio, the Companhia Ambiental do Estado de São Paulo (CETESB), the Instituto de Energia e Meio Ambiente (IEMA), and PurpleAir contributors. We also thank Google (Google Air Quality API and Google Earth Engine), AccuWeather and Plume Labs (AccuWeather Air Quality API), Open-Meteo, OpenStreetMap, and geoBoundaries for air-quality, meteorological, and geospatial data services, as well as the public-health agencies and study authors whose open outcome datasets made source-substitution reanalysis possible. Their commitment to reproducible science makes it possible to study, research, and ultimately reduce the health burden of air pollution worldwide.

References

AccuWeather. Guide | AccuWeather Enterprise API, 2026. URL https://apidev.accuweather.com/developers/air-quality/guide. Accessed 5 September 2026.

Antoine Alléon, Grégoire Jauvion, Boris Quennehen, and David Lissmyr. PlumeNet: Large-scale air quality forecasting using a convolutional LSTM network. *arXiv preprint arXiv:2006.09204*, 2020. URL https://arxiv.org/abs/2006.09204.

Karoline K. Barkjohn, Brett Gantt, and Andrea L. Clements. Development and application of a United States-wide correction for $PM_{2.5}$ data collected with the PurpleAir sensor. *Atmospheric Measurement Techniques*, 14(6):4617–4637, 2021. doi: 10.5194/amt-14-4617-2021.

Rob Beelen, Ole Raaschou-Nielsen, Massimo Stafoggia, et al. Effects of long-term exposure to air pollution on natural-cause mortality: An analysis of 22 European cohorts within the multicentre ESCAPE project. *The Lancet*, 383(9919):785–795, 2014. doi: 10.1016/S0140-6736(13)62158-3.

Russell Biggs and OpenAQ. OpenAQ Python SDK, 2025. URL https://github.com/openaq/openaq-python/releases/tag/v0.4.0. Version 0.4.0; computer software; OpenAQ, ROR: https://ror.org/03a79pm79; MIT License (v0.4.0); released 31 March 2025.

Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023. doi: 10.1038/s41586-023-06792-0.

Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D. White, and Philippe Schwaller. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5):525–535, 2024. doi: 10.1038/s42256-024-00832-8.

Michael S. Breen, Catherine Seppanen, Vlad Isakov, Saravanan Arunachalam, Miyuki Breen, James M. Samet, and Haiyan Tong. Development of TracMyAir smartphone application for modeling exposures to ambient $PM_{2.5}$ and ozone. *International Journal of Environmental Research and Public Health*, 16(18):3468, 2019. doi: 10.3390/ijerph16183468.

Robert D. Brook, Sanjay Rajagopalan, C. Arden Pope, III, et al. Particulate matter air pollution and cardiovascular disease: An update to the scientific statement from the American Heart Association. *Circulation*, 121(21):2331–2378, 2010. doi: 10.1161/CIR.0b013e3181dbece1.

Bert Brunekreef and Stephen T. Holgate. Air pollution and health. *The Lancet*, 360(9341):1233–1242, 2002. doi: 10.1016/S0140-6736(02)11274-8.

Dechao Bu et al. Empowering AI data scientists using a multi-agent LLM framework with self-evolving capabilities for autonomous, tool-aware biomedical data analyses. *Nature Biomedical Engineering*, 2026. doi: 10.1038/s41551-026-01634-6.

Thibaut Cassard, Grégoire Jauvion, and David Lissmyr. High-resolution air quality prediction using low-cost sensors. *arXiv preprint arXiv:2006.12092*, 2020. URL `https://arxiv.org/abs/2006.12092`.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Aleksander Madry, and Lilian Weng. MLE-bench: Evaluating machine learning agents on machine learning engineering. In *International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=6s5uXNWGIh`.

Ziru Chen, Shijie Chen, Yuting Ning, et al. ScienceAgentBench: Toward rigorous assessment of language agents for data-driven scientific discovery. In *International Conference on Learning Representations*, 2025. URL `https://proceedings.iclr.cc/paper_files/paper/2025/hash/f12b4df26344f3be803c06b555252efe-Abstract-Conference.html`.

Qian Di, Yan Wang, Antonella Zanobetti, Yun Wang, Petros Koutrakis, Christine Choirat, Francesca Dominici, and Joel D. Schwartz. Air pollution and mortality in the Medicare population. *The New England Journal of Medicine*, 376(26):2513–2522, 2017. doi: 10.1056/NEJMoa1702747.

Qian Di et al. An ensemble-based model of $PM_{2.5}$ concentration across the contiguous United States with high spatiotemporal resolution. *Environment International*, 130:104909, 2019. doi: 10.1016/j.envint.2019.104909.

Douglas W. Dockery, C. Arden Pope, III, Xiping Xu, John D. Spengler, James H. Ware, Martha E. Fay, Benjamin G. Ferris, Jr., and Frank E. Speizer. An association between air pollution and mortality in six U.S. cities. *The New England Journal of Medicine*, 329(24):1753–1759, 1993. doi: 10.1056/NEJM199312093292401.

Francesca Dominici, Roger D. Peng, Michelle L. Bell, Luu Pham, Aidan McDermott, Scott L. Zeger, and Jonathan M. Samet. Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *JAMA*, 295(10):1127–1134, 2006. doi: 10.1001/jama.295.10.1127.

David Donaire-Gonzalez et al. ExpoApp: An integrated system to assess multiple personal environmental exposures. *Environment International*, 126:494–503, 2019. doi: 10.1016/j.envint.2019.02.054.

Clinton H. Durney, Aadetri Tawara, Michael Brauer, Sukhinder Atkar-Khattra, Renelle Myers, Rafael Meza, and Stephen Lam. APEX: A web-based tool for assessing long-term outdoor $PM_{2.5}$ exposure—brief report. *JTO Clinical and Research Reports*, 7(5):100975, 2026. doi: 10.1016/j.jtocrr.2026.100975.

Marloes Eeftens, Rob Beelen, Kees de Hoogh, et al. Development of land use regression models for $PM_{2.5}$, $PM_{2.5}$ absorbance, PM_{10} and PM_{coarse} in 20 European study areas; results of the ESCAPE project. *Environmental Science & Technology*, 46(20):11195–11205, 2012. doi: 10.1021/es301948k.

Google Maps Platform. History | Air Quality API | Google for Developers, 2026. URL `https://developers.google.com/maps/documentation/air-quality/history`. Accessed 25 August 2026.

Juraj Gottweis et al. Accelerating scientific discovery with Co-Scientist. *Nature*, 655(8122):487–496, 2026. doi: 10.1038/s41586-026-10644-y.

Health Effects Institute. State of Global Air 2024: Special report. Technical report, Health Effects Institute, Boston, MA, 2024. URL https://www.stateofglobalair.org/resources/archived/state-global-air-report-2024. Produced in partnership with UNICEF. Accessed 25 August 2026.

Gerard Hoek, Rob Beelen, Kees de Hoogh, Danielle Vienneau, John Gulliver, Paul Fischer, and David Briggs. A review of land-use regression models to assess spatial variation of outdoor air pollution. *Atmospheric Environment*, 42(33):7561–7578, 2008. doi: 10.1016/j.atmosenv.2008.05.057.

Grégoire Jauvion, Thibaut Cassard, Boris Quennehen, and David Lissmyr. DeepPlume: Very high resolution real-time air quality mapping. *arXiv preprint arXiv:2002.10394*, 2020. URL https://arxiv.org/abs/2002.10394.

Michael Jerrett, Altaf Arain, Pavlos Kanaroglou, Bernardo Beckerman, Dimitri Potoglou, Talar Sahsuvaroglu, Jason Morrison, and Chris Giovis. A review and evaluation of intraurban air pollution exposure models. *Journal of Exposure Analysis and Environmental Epidemiology*, 15(2): 185–204, 2005. doi: 10.1038/sj.jea.7500388.

Haneen Khreis, Charlotte Kelly, James Tate, Roger Parslow, Karen Lucas, and Mark Nieuwenhuijsen. Exposure to traffic-related air pollution and risk of development of childhood asthma: A systematic review and meta-analysis. *Environment International*, 100:1–31, 2017. doi: 10.1016/j.envint.2016.11.012.

Itai Kloog, Alexandra A. Chudnovsky, Allan C. Just, Francesco Nordio, Petros Koutrakis, Brent A. Coull, Alexei Lyapustin, Yujie Wang, and Joel Schwartz. A new hybrid spatio-temporal model for estimating daily multi-year $PM_{2.5}$ concentrations across northeastern USA using high resolution aerosol optical depth data. *Atmospheric Environment*, 95:581–590, 2014. doi: 10.1016/j.atmosenv.2014.07.014.

Romain Lacombe. Distillate: A meta-agent harness for autonomous ML research. Working paper, 2026. URL https://distillate.dev/paper.pdf.

Dong Liang and Naresh Kumar. Time-space kriging to address the spatiotemporal misalignment in the large datasets. *Atmospheric Environment*, 72:60–69, 2013. doi: 10.1016/j.atmosenv.2013.02.034.

Yougeng Lu, Genevieve Giuliano, and Rima Habre. Estimating hourly $PM_{2.5}$ concentrations at the neighborhood scale using a low-cost air sensor network: A Los Angeles case study. *Environmental Research*, 195:110653, 2021. doi: 10.1016/j.envres.2020.110653.

Lidia Morawska et al. Applications of low-cost sensing technologies for air quality monitoring and exposure assessment: How far have they gone? *Environment International*, 116:286–299, 2018. doi: 10.1016/j.envint.2018.04.018.

Nominatim developer community. Nominatim 5.3.2 Manual, 2026. URL https://nominatim.org/release-docs/5.3.2/. Accessed 5 September 2026.

Marguerite Nyhan, Sebastian Grauwin, Rex Britter, Bruce Misstear, Aonghus McNabola, Francine Laden, Steven R. H. Barrett, and Carlo Ratti. Exposure Track—the impact of mobile-device-based mobility patterns on quantifying population exposure to air pollution. *Environmental Science & Technology*, 50(17):9671–9681, 2016. doi: 10.1021/acs.est.6b02385.

Marguerite M. Nyhan, Itai Kloog, Rex Britter, Carlo Ratti, and Petros Koutrakis. Quantifying population exposure to air pollution using individual mobility patterns inferred from mobile phone data. *Journal of Exposure Science & Environmental Epidemiology*, 29(2):238–247, 2019. doi: 10.1038/s41370-018-0038-9.

OpenAQ. OpenAQ API, 2026a. URL https://api.openaq.org. Organization ROR: https://ror.org/03a79pm79. Air-quality data aggregated and distributed by OpenAQ; original providers retain their own licenses. Accessed 26 September 2026.

OpenAQ. About Open Data on AWS, 2026b. URL https://docs.openaq.org/aws/about. Accessed 5 September 2026.

Pablo Orellano, Julieta Reynoso, Nancy Quaranta, Ariel Bardach, and Agustin Ciapponi. Short-term exposure to particulate matter (PM_{10} and $PM_{2.5}$), nitrogen dioxide (NO_2), and ozone (O_3) and all-cause and cause-specific mortality: Systematic review and meta-analysis. *Environment International*, 142:105876, 2020. doi: 10.1016/j.envint.2020.105876.

C. Arden Pope, III, Richard T. Burnett, Michael J. Thun, Eugenia E. Calle, Daniel Krewski, Kazuhiko Ito, and George D. Thurston. Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution. *JAMA*, 287(9):1132–1141, 2002. doi: 10.1001/jama.287.9.1132.

Daniel Runfola, Austin Anderson, Heather Baier, Matt Crittenden, Elizabeth Dowker, Sydney Fuhrig, Seth Goodman, Grace Grimsley, Rachel Layko, Graham Melville, Maddy Mulder, Rachel Oberman, Joshua Panganiban, Andrew Peck, Leigh Seitz, Sylvia Shea, Hannah Slevin, Rebecca Youngerman, and Lauren Hobbs. geoBoundaries: A global database of political administrative boundaries. *PLOS ONE*, 15(4):e0231866, 2020. doi: 10.1371/journal.pone.0231866.

Patrick H. Ryan and Grace K. LeMasters. A review of land-use regression models for characterizing intraurban air pollution exposure. *Inhalation Toxicology*, 19(Suppl 1):127–133, 2007. doi: 10.1080/08958370701495998.

Jason D. Sacks, Jennifer M. Lloyd, Yun Zhu, Jim Anderton, Carey J. Jang, Bryan Hubbell, and Neal L. Fann. The Environmental Benefits Mapping and Analysis Program—Community Edition (BenMAP-CE): A tool to estimate the health and economic benefits of reducing air pollution. *Environmental Modelling & Software*, 104:118–129, 2018. doi: 10.1016/j.envsoft.2018.02.009.

Jonathan M. Samet, Francesca Dominici, Frank C. Curriero, Ivan Coursac, and Scott L. Zeger. Fine particulate air pollution and mortality in 20 U.S. cities, 1987–1994. *The New England Journal of Medicine*, 343(24):1742–1749, 2000. doi: 10.1056/NEJM200012143432401.

Heidi Seibold, Severin Czerny, Siona Decke, Roman Dieterle, Thomas Eder, Steffen Fohr, Nico Hahn, Rabea Hartmann, Christoph Heindl, Philipp Kopper, Dario Lepke, Verena Loidl, Maximilian Mandl, Sarah Musiol, Jessica Peter, Alexander Piehler, Elio Rojas, Stefanie Schmid, Hannah Schmidt, Melissa Schmoll, Lennart Schneider, Xiao-Yin To, Viet Tran, Antje Völker, Moritz Wagner, Joshua Wagner, Maria Waize, Hannah Wecker, Rui Yang, Simone Zellner, and Malte Nalenz. A computational reproducibility study of PLOS ONE articles featuring longitudinal data analyses. *PLOS ONE*, 16(6):e0251194, 2021. doi: 10.1371/journal.pone.0251194.

Zachary S. Siegel, Sayash Kapoor, Nitya Nadgir, Benedikt Stroebl, and Arvind Narayanan. CORE-Bench: Fostering the credibility of published research through a computational reproducibility agent benchmark. *Transactions on Machine Learning Research*, 2024. URL `https://openreview.net/forum?id=BsMMc4MEGS`.

Susanne Steinle, Stefan Reis, and Clive E. Sabel. Quantifying human exposure to air pollution—moving from static monitoring to spatio-temporally resolved personal exposure assessment. *Science of the Total Environment*, 443:184–193, 2013. doi: 10.1016/j.scitotenv.2012.10.098.

Victoria Stodden, Jennifer Seiler, and Zhaokun Ma. An empirical analysis of journal policy effectiveness for computational reproducibility. *Proceedings of the National Academy of Sciences*, 115(11): 2584–2589, 2018. doi: 10.1073/pnas.1708290115.

Adam A. Szpiro, Paul D. Sampson, Lianne Sheppard, Thomas Lumley, Sara D. Adar, and Joel D. Kaufman. Predicting intra-urban variation in air pollution concentrations with complex spatio-temporal dependencies. *Environmetrics*, 21(6):606–631, 2010. doi: 10.1002/env.1014.

Aaron van Donkelaar, Randall V. Martin, Michael Brauer, Ralph Kahn, Robert Levy, Carolyn Verduzco, and Paul J. Villeneuve. Global estimates of ambient fine particulate matter concentrations from satellite-based aerosol optical depth: Development and application. *Environmental Health Perspectives*, 118(6):847–855, 2010. doi: 10.1289/ehp.0901623.

Xiaonan Yu et al. Quantifying the impact of daily mobility on errors in air pollution exposure estimation using mobile phone location data. *Environment International*, 141:105772, 2020. doi: 10.1016/j.envint.2020.105772.

Roman Yurchak. pgeocode — pgeocode 0.5.0 documentation, 2024. URL https://pgeocode.readthedocs.io/. Version 0.5.0. Accessed 5 September 2026.

Alex L. Zhang, Tim Kraska, and Omar Khattab. Recursive language models. *arXiv preprint arXiv:2512.24601*, 2025. URL https://arxiv.org/abs/2512.24601.

A Study-identification pipeline in full

Figure 2 condenses the study-identification pipeline to its principal stages. This appendix reports the intermediate counts, explains each screening rule, and shows how the two discovery routes combine.

A.1 Corpus

We searched PubMed and OpenAlex for public-health studies published during the preceding five years that associated an outdoor ambient pollutant with a human-health outcome. We deduplicated results by DOI, then by normalized title, year, and first author when no DOI was available, leaving 18,919 records. Both discovery routes searched this archived corpus in parallel. A study found by both routes therefore represents overlap between search strategies rather than a duplicated record.

A.2 Route 1: scientific-scope screening

In Route 1, we screened titles, abstracts, and metadata for empirical studies of outdoor ambient exposure and human-health outcomes. We then assessed data access, scientific fit, and whether the available inputs could support source-substitution reanalysis.

Rapid title, abstract, and metadata screen. The screen retained records that might describe an empirical human-health study of outdoor ambient exposure. It excluded clear reviews, protocols, editorials, laboratory or non-human studies, and studies outside outdoor ambient exposure.

Decision	Records
Retained	10,020
Uncertain	5,985
Excluded	2,914
Total	18,919

The metadata did not support a decision for 5,985 records. We retained them as unresolved rather than scientific exclusions, leaving some risk that the search missed eligible studies.

Open-access confirmation. We checked the 10,020 retained records for a PubMed Central identifier or OpenAlex metadata identifying an open full-text location.

Decision	Records
Confirmed open access	4,658
Open-access status unresolved	5,359
Confirmed non-open-access	3
Total	10,020

We could not confirm access for 5,359 records. The prespecified Route 1 workflow continued with the 4,658 records whose open-access status we could confirm.

Strict linked-data path. The strict path searched the 4,658 confirmed-open-access papers for a concrete public-data URL. Full-text review then tested scientific fit, followed by field-level validation of outcomes, dates, locations, covariates, and the analysis specification.

Stage	Records
Confirmed open access	4,658
Concrete public-data URL	138
Passed scientific-fit review	22
Passed field-level validation	3

Public-source rescue path. The rescue path screened the 4,520 confirmed-open-access papers with PubMed Central identifiers for named public systems from which the required inputs could be reconstructed even when the article did not link a complete analysis dataset. Full text was cached for 3,707 papers. A high-precision screen selected 133 candidates; three already belonged to the strict candidate pool and were removed before rescue review, leaving 130 new candidates.

Stage	Records
Confirmed-open-access PMC papers	4,520
Cached full texts screened	3,707
High-precision candidates	133
New candidates after removing strict-path overlap	130
Passed full-text source review	13
Passed field-level validation	6

The strict and rescue inputs overlap: the 4,520 PMC papers are a subset of the 4,658 confirmed-open-access papers, and 138 PMC papers passed the strict URL gate. We removed three overlapping records when we combined the candidate lists. The two paths yielded $3 + 6 = 9$ validated scientific-scope candidates. The rescue path's 13 full-text passes are distinct from the 13 Route 2 candidates below.

A.3 Route 2: screening for publicly accessible health outcomes

Title and abstract screening. A script examined each title and available abstract for a named source that might provide publicly accessible health outcomes. Signals included vital statistics and civil-registration systems; DATASUS and SIH/SUS; NHANES, DHS, MICS, and BRFSS; cancer and disease registries; notifiable-disease and hospital-discharge systems; and named government or international COVID-19 reporting systems. This pass set review priority rather than study eligibility.

Routing label	Records
Named public outcome source	203
Likely public aggregate outcome	3,430
Likely individual-level or restricted outcome	6,251
Insufficient metadata to identify a route	9,035
Total	18,919

These labels did not exclude records. We left 9,035 records unresolved because their metadata did not identify an outcome source. We prioritized the 203 records that named one for article review.

Article and source review. Agents read the Methods and data-availability material for all 203 papers and inspected each cited outcome source.

Access disposition	Papers
Direct public download	45
Public source with a clear retrieval route	47
Registration or account required	52
Agreement, application, or request required	36
No usable retrieval link found	21
Unresolved	2
Total	203

The first two groups formed 92 papers with a potentially accessible public outcome. The $45+47 = 92$ split describes two forms of access and is not an additional filtering stage.

Health-outcome access and ambient-pollution review. Agents reviewed all 92 papers again to identify studies of outdoor ambient air pollution, wildfire smoke, or desert dust whose cited outcomes could be retrieved without authentication.

Decision	Papers
Ambient pollution or wildfire smoke	17
Unauthenticated public outcome retrieval	13
Agreement or request required	2
Public route not working or not found	2
Non-ambient exposure or unrelated subject	75
Total	92

Route 2 therefore produced 13 candidates. Figure 2 combines the ambient-scope and access checks into the $92 \rightarrow 13$ transition.

A.4 Combining the routes and conducting the full audit

Routes 1 and 2 produced 9 and 13 candidates, respectively, for 22 route hits. One study appeared in both, leaving 21 unique candidates. The limited overlap shows that the routes captured substantially different parts of the literature.

We then audited the paper and data sources for every candidate. We required an outcome source available without an account, request, agreement, approval, or collaborator; recoverable dates and geographic identifiers; a described primary model and required non-exposure covariates; public monitoring data for the reported pollutant; and enough information for AERIS to reconstruct exposure independently.

Audit decision	Studies
Eligible for attempted reanalysis	14
Excluded during audit	7
Missing a required input or method	6
Historical outcome export not recovered	1
Total	21

A.5 Execution

AERIS attempted all 14 eligible studies; nine produced fitted reanalyses.

Execution result	Studies
Completed with fitted results	9
No fitted result	5
Inadequate historical air-quality coverage or variation	3
Required outcome export not retrieved	2
Total	14

The five execution exclusions were:

- **China COVID-19:** we found two usable prefecture panels for 2014–2016, compared with 287 in the published analysis.
- **Israel COVID-19:** the available $PM_{2.5}$ panel had no spatial variation across municipalities for the study dates.
- **Argentina COVID-19:** the completed boundary run yielded too few $PM_{2.5}$ observations for a fitted panel.
- **Taubaté COVID-19 admissions:** the required daily SIH/SUS outcome export was not recovered.
- **Alaska wildfire asthma:** we could not download the required regional outcome table from a public source.

The audit established that an analysis appeared runnable from the article and identified sources. Execution tested whether those sources supplied adequate data in practice. Three failures arose from inadequate air-quality coverage or variation; two arose because we could not retrieve the required outcome export. The records do not establish that those outcome exports had previously been downloadable and later disappeared.

A.6 Count reconciliation

Stage	Count	Yield
Unique search records	18,919	—
Route 1 candidates	9	0.048% of corpus
Route 2 candidates	13	0.069% of corpus
Route hits	22	—
Unique candidates after one overlap	21	—
Eligible after audit	14	66.7% of candidates
Completed reanalyses	9	64.3% of attempts
Completed among unique candidates	9	42.9%
Corpus-to-completion yield	9	0.048%

GPT-5 Codex subagents conducted the article-level reviews on 13–14 August 2026. The public repository includes the search protocol and scripts, review prompt, source records, exclusion reasons, and execution artifacts. Figure 2 omits the uncertain and unresolved records, Route 1's overlapping strict and rescue subpaths, the $45 + 47$ access subdivision, and Route 2's intermediate count of 17 ambient-exposure studies; the appendix reports all of them above.

B Trajectory-pathway implementation details

The open-source trajectory pathway estimates hourly $PM_{2.5}$ at each time-stamped coordinate with a packaged XGBoost regression model. The model contains 600 gradient-boosted trees and minimizes squared error. Its 14 ordered inputs are: the nearest PurpleAir $PM_{2.5}$ measurement within 5 km; temperature, relative humidity, wind speed, wind direction, and precipitation from Open-Meteo; distance to the nearest highway and arterial road from OpenStreetMap; open space, water, impervious surface, tree cover, and elevation from Google Earth Engine; and a weekday indicator. When PurpleAir has no nearby measurement, AERIS can substitute $PM_{2.5}$ from an OpenAQ regulatory monitor within 25 km. The run record preserves the input sources and missing-feature warnings for each estimate.

The feature design extends the neighborhood-scale random-forest method of Lu et al. [2021], which combined low-cost sensor measurements with meteorological, temporal, and spatial predictors. AERIS loads the model weights locally, so a trajectory need not be sent to a third-party prediction service. Researchers may also supply locally trained XGBoost weights with the same ordered inputs. The packaged weights support implementation testing and reproducible inference; their accuracy outside the training domain depends on local sensor conditions and requires external validation. We verified that the complete pathway runs on a 24-hour New York City trajectory. This execution test does not measure predictive accuracy or personal dose.

C Ambient-data sources and local $PM_{2.5}$ inference

Table 2 summarizes each external source, the information it contributes, and its role in exposure reconstruction. The local inference pathway groups the packaged model with the sensor, weather, and geographic inputs it expects.

OpenAQ aggregates observations from governments, researchers, and sensor networks rather than originating every measurement. We therefore cite OpenAQ as the distribution platform. Provider-specific licenses govern downstream reuse. The supplement's `OPENAQ_ATTRIBUTION_MANIFEST` audits the provider and license evidence for all nine fitted reanalyses and flags records that cannot be verified from the preserved traces. The artifact does not redistribute raw observations.

Table 2: Ambient-data sources and model inputs. AERIS records the selected pathway, source settings, and assignment method in the run manifest.

Source	What it provides	Why AERIS uses it
OpenAQ API and public S3 archive	Global repository of hourly and daily station observations from public and regulatory monitoring networks	Supplies historical, station-level concentrations, monitor metadata, and coverage diagnostics for fixed-location and area-based exposure assignment
Google Maps Air Quality API	Location-specific hourly air-quality history and pollutant information	Supplies a credentialed coordinate-level pathway when a trajectory requires location-time estimates beyond a monitor assignment
AccuWeather Air Quality API	Credentialed coordinate-level current and forecast air-quality index and category information	Adds a localized current and forecast source-inspection pathway. Its index is retained as an index; concentration-based epidemiological models require a source that returns physical pollutant concentrations [AccuWeather, 2026]
Local XGBoost $PM_{2.5}$ inference pathway		
Packaged XGBoost weights	A local hourly $PM_{2.5}$ model evaluated at a requested coordinate and time	Extends the low-cost-sensor, weather, and spatial-feature approach of Lu et al. [2021] with gradient-boosted trees; produces a reproducible estimate without sending participant locations to a third-party prediction service
PurpleAir	Dense low-cost sensor observations	Supplies nearby particulate measurements that capture local variation between regulatory stations
Open-Meteo weather	Meteorological conditions at the requested place and time	Supplies weather features that influence pollutant formation, dispersion, and transport
OpenStreetMap	Roads and other built-environment features	Supplies local source-proxy features associated with traffic and urban form
Land-cover inputs	Surface-cover features around the requested coordinate	Supplies geographic context associated with dispersion and background concentration

D Exposure defaults and MCP tool contracts

Table 3 states the defaults used when a researcher does not supply a study-specific choice. They are starting points rather than universal scientific rules: AERIS records any override in the run manifest. In particular, AERIS imposes no default completeness cutoff.

The local MCP server returns each result in a common envelope containing `status`, warnings, tool name, workspace identifier, artifact identifier, and run or plan identifier. Invalid arguments, authorization failures, missing data, and upstream failures return typed error fields instead of fabricated exposure values. Table 4 groups the public endpoints; the artifact includes their complete JSON schemas.

Table 3: Default exposure computations in the released implementation.

Pathway	Default	Computation and edge cases
Fixed location	Nearest monitor within 25 km	The caller may request more monitors or a radius up to 200 km. With multiple monitors, AERIS aligns observations by UTC hour and applies inverse-distance weighting with power 2. An exactly co-located monitor receives the full weight. Hours without observations are omitted.
Administrative area	All eligible monitors within the resolved boundary	AERIS first averages observations within each monitor, then takes the unweighted mean across monitor means. Its daily series likewise averages within monitor-day before averaging across monitors. If a station budget is supplied, deterministic spatial spread selects the subset.
Temporal completeness	No cutoff	AERIS reports observed versus expected periods. A user-supplied threshold flags or excludes incomplete monitor assignments; no hidden 75% rule applies. A query with no usable observations returns a missing estimate and an explicit quality-control code.
Historical recovery	OpenAQ API, then public S3 archive	An empty historical API response triggers bounded archive retrieval. The default automatic archive window is at most 400 days and 500 objects; larger jobs require an explicit plan.

Table 4: Compact MCP contract. Arguments shown are the principal scientific inputs; all outputs carry the common provenance envelope.

Tool	Principal inputs	Output
`load_cohort`	Local table path	Validated study schema, location mode, date range, workspace and dataset IDs
`plan_exposure`	Pollutant, country, monitor count, temporal resolution, location strategy	Proposed estimand, source plan, workload, and plan ID
`estimate_exposure`	Approved plan; pollutant, backend, monitor count, completeness rule	Fixed-location or trajectory exposure artifact, coverage summary, and run ID
`estimate_boundary_exposure`	Approved boundary, pollutant, dates, resolution, completeness rule	Area estimate, monitor coverage, boundary provenance, artifact and run IDs
`get_exposure_summary`; `get_daily_data`; `get_raw_data`; `get_station_info`	Artifact, subject or pollutant selector	Bounded summaries, normalized observations, and contributing-monitor records
`record_study_design`; `get_decision_log`	Question, outcome, pollutant, location and time units, assignment and analysis plan	Durable design and decision records
`list_artifacts`; `describe_artifact`; `sample_artifact`; `query_artifact`	Artifact ID and bounded inspection operation	Schema, provenance, samples, or aggregate statistics without loading the full table into chat
`export_results`	Artifact selector, local path, format	Reviewable exposure table and export provenance

E Completed source-substitution reanalyses

E.1 Summary metrics

Table 5: AERIS reanalysis summary with researcher review. Direction and null recovery use fitted contrasts; numerical agreement gives each study equal weight. We exclude the two Northern Italy contrasts from continuous comparisons because the article reports them only as 1.00.

Measure	Result
Contrast-level direction and null recovery	
Fitted contrasts	14
Published positive associations, same direction	7 / 7
Published positive associations, AERIS CI excludes 1	3 / 7
Published null or no-evidence contrasts, AERIS CI includes 1	4 / 7 (all OECD)
Study-weighted numerical agreement	
Studies	8
Pearson r (Fisher 95% CI)	0.834 (0.314–0.969)
Leave-one-study-out Pearson r range	0.225–0.958
Median absolute log difference	0.040
Mean absolute log difference	0.078
Within 10% on the ratio scale	6 / 8
Geometric mean AERIS / published ratio	0.998

E.2 Contrast-level results

E.3 Study registry

For workload accounting, a monitor record is counted each time an observation enters a study-specific exposure calculation; the same physical observation may therefore enter more than one study or geographic assignment. An area-time exposure estimate is one pollutant value assigned to one administrative area and time interval.

United Kingdom COVID-19. M. Pear Hossain, Wen Zhou, Marco Y. T. Leung, and Hsiang-Yu Yuan. *Association of air pollution and weather conditions during infection course with COVID-19 case fatality rate in the United Kingdom*. DOI: 10.1038/s41598-023-50474-w. The study relates country-level infection-course pollution and weather to daily COVID-19 case-fatality rates from March through September 2020. We reconstruct the four country series with public monitor-derived $PM_{2.5}$.

Italian provincial capitals COVID-19. Mario Coccia. *Effects of the spread of COVID-19 on public health of polluted cities: results of the first wave for explaining the déjà vu in the second wave of COVID-19 pandemic and epidemics of future vital agents*. DOI: 10.1007/s11356-020-11662-7. This ecological study links pre-pandemic PM_{10} exceedance days in Italian cities to first-wave COVID-19 outcomes. We combine public city outcomes with independently reconstructed 2018 monitor exposure.

France COVID-19. Anastase Tchicaya, Nathalie Lorentz, Hichem Omrani, Gaetan de Lanchy, and Kristell Leduc. *Impact of long-term exposure to $PM_{2.5}$ and temperature on coronavirus disease mortality: observed trends in France*. DOI: 10.1186/s12940-021-00784-1. The study models department-level COVID-19 mortality with long-term $PM_{2.5}$, temperature, and public health-system covariates. We join public department outcomes to a 2019 AERIS monitor-based exposure panel.

OECD COVID-19. Zohar Barnett-Itzhaki and Adi Levi. *Effects of chronic exposure to ambient air pollutants on COVID-19 morbidity and mortality: a lesson from OECD countries*. DOI: 10.1016/j.envres.2021.110723. This cross-country study evaluates $PM_{2.5}$ against COVID-19 cases and deaths at prespecified days after the first case. We fit the public outcome series for countries with adequate 2017 AERIS coverage.

Table 6: AERIS source-substitution reanalyses with researcher review on a common multiplicative scale. Bold values are odds ratios (OR) or rate ratios (RR). Rows follow each article's conclusion; the setting columns expose differences between the published and AERIS analyses.

Study and contrast	Published setting	AERIS setting	Published OR/RR	AERIS OR/RR
Published positive evidence				
France COVID-19				
RR / 1 μg/m^3 $PM_{2.5}$	96 departments, 1999–2016	63 departments, 2019	**1.24** (1.08–1.43)	**1.05** (0.95–1.15)
Italy COVID-19				
OR / doubling PM_{10} days	55 cities, 2018	8 cities, 2018	**1.76**‡	**1.74** (1.13–2.69)
São José admissions	SIH/SUS, 2020–2021	SIVEP-Gripe, 89 days in 2021		
O_3 lag 2 RR / IQR			**1.061** (1.016–1.108)	**1.003** (0.981–1.026)
NO_2 lag 3 RR / IQR			**1.125** (1.050–1.205)	**1.034** (0.876–1.220)
São Paulo AMI				
High-pollution RR	39 municipalities, 2012–2015	39 municipalities, 193 days	**1.090**	**1.069** (1.006–1.135)
UK COVID-19				
OR / 10 μg/m^3 $PM_{2.5}$	4 nations, 696 days	4 nations, 635 nation-days	**1.13** (1.09–1.16)	**1.46** (1.21–1.76)
US stroke mortality				
RR / 10 μg/m^3 $PM_{2.5}$†	US counties, 2016–2020	46 counties, 2019	**1.132**	**1.111** (0.628–1.594)
Published null or no evidence				
Brazil cancer mortality				
OR / 10 μg/m^3 $PM_{2.5}$	Brazil, 2015–2019	5 states, 1,824 days; 929,790 deaths	**1.010** (1.00–1.02)	**1.060** (1.052–1.069)
Italy NO_2 COVID-19	2 regions, Feb–May 2020	2 regions, 14 days		
Incidence RR / 1 μg/m^3			**1.00***	**1.051** (1.042–1.060)
Mortality RR / 1 μg/m^3			**1.00***	**0.988** (0.987–0.990)
OECD COVID-19	36 countries	11 countries, 2017		
Day 60 cases			**1.020**	**0.998** (0.853–1.168)
Day 60 deaths			**1.030**	**1.016** (0.806–1.279)
Day 80 cases			**1.020**	**0.971** (0.926–1.018)
Day 80 deaths			**1.009**	**0.971** (0.873–1.079)

* The article reports a null effect rounded to two decimals. † Stroke ratios use the reported linear-model coefficient and mean outcome rate. ‡ For the published estimate, the article reports the underlying log-log coefficient without a confidence interval or p-value.

Northern Italy COVID-19. Chi-Chang Ho, Shih-Chieh Hung, and Wen-Chao Ho. *Effects of short- and long-term exposure to atmospheric pollution on COVID-19 risk and fatality: analysis of the first epidemic wave in northern Italy*. DOI: `10.1016/j.envres.2021.111293`. The study models daily incidence and mortality in Lombardia and Veneto. We combine public regional outcomes with the recovered AERIS NO_2 series.

São José dos Campos COVID admissions. A.O.R. Santos, B.R. Lucarevschi, C.J.D. Cunha, P.C. Ribeiro, A.C.G. Cesar, and L.F. Nascimento. *Exposure to air pollution and hospitalization due to COVID-19 in São José dos Campos, Brazil*. DOI: `10.1590/1414-431x2021e12273`. The article

estimates lagged pollutant associations with daily admissions. Because the full SIH/SUS period could not be recovered, we use public SIVEP-Gripe admissions.

São Paulo acute myocardial-infarction admissions. Marcelo H. Ribeiro, André Grossi, Adriano Caixeta, Marcelo Franken, Marcelo Katz, Vinicius Seleme, Expedito Ribeiro, Antonio Eduardo Pesaro, Jose Fabri, Sameer Mehta, and Carlos M. Campos. *Cluster of climatic and pollutant characteristics increases admissions for acute myocardial infarction: Analysis of 30,423 patients in the metropolitan area of Sao Paulo*. DOI: 10.1016/j.hrtlng.2020.04.018. The study contrasts daily admission counts across weather–pollution clusters. We use SIH/SUS admissions, Open-Meteo weather, and a reduced-input NO_2 cluster sensitivity analysis.

Brazil cancer mortality. Pei Yu, Rongbin Xu, Yao Wu, Wenzhong Huang, Micheline S.Z.S. Coelho, Paulo H.N. Saldiva, Tingting Ye, Bo Wen, Yanming Liu, Zhengyu Yang, Shanshan Li, Michael J. Abramson, and Yuming Guo. *Cancer mortality risk from short-term $PM_{2.5}$ exposure and temporal variations in Brazil*. DOI: 10.1016/j.jhazmat.2024.134606. The paper uses a national time-stratified case-crossover design for daily cancer mortality. Our reanalysis covers 1,824 days and 929,790 deaths using public IEMA monitoring records from five states.

United States stroke mortality. Jean-Luc K. Kabangu, Danny Fowler, Amanda Hernandez, Takara Newsome-Cuby, Benson Joseph, John Dugan, Lane Fry, Momodou G. Bah, and Sonia V. Eden. *Inequities in air pollution on stroke mortality among Older Americans: a U.S. nationwide analysis*. DOI: 10.3389/fpubh.2024.1364165. The article estimates county-level $PM_{2.5}$ associations with ischemic-stroke mortality among adults aged 65+ years. We use 2019 public all-stroke mortality for 46 counties and the stated covariate families.

F Researcher interventions during the primary evaluation

AERIS found the papers and public data, wrote retrieval and analysis code, reconstructed exposure, built the analysis panels, and fit the models. The researchers intervened at scientific branch points: they asked the agent to continue after a failed route, challenged unnecessarily narrow coverage, corrected an estimand or reporting scale, or approved an explicitly labeled fallback. Table 7 records these interventions. We reconstructed the ledger from the project conversation, per-study code and result notes, prompt summaries, and the later supervision audit. Because the campaign did not retain one canonical turn-by-turn transcript, we report documented intervention classes rather than infer a precise count.

Across studies, the interventions served five recurring purposes: persistence after a failed access route, broader geographic coverage, fidelity to the article's estimand, conversion to the reported effect scale, and approval of a clearly labeled bounded analysis. They did not provide retrieval code, source URLs, exposure values, or fitted answers. This retrospective record supports the paper's claim that the primary campaign used AERIS with researcher review; it cannot establish a precise intervention count. The autonomous campaigns in Appendix G instead used instructions fixed before evaluation and received no scientific intervention.

G Autonomous AERIS evaluation

Protocol. We fixed the AERIS instructions, retry policy, schemas, safety envelope, and container before evaluation. Studies ran in fresh workspaces with Codex CLI, GPT-5.6 Terra, and medium reasoning. The agent received the article bundle and scientific target, but no prior code, source routes, manifests, exposure values, or results. It could retrieve public data, inspect endpoints, write code, revise its design, and run bounded analyses without approval. A host-side critic retried unsupported blockers and invalid handoffs. No run received human researcher intervention.

Results and stability. The three autonomous campaigns produced pollutant-effect estimates for 8, 9, and 6 studies, or 23 of 27 attempts (Table 8 lists the first campaign). One campaign completed all nine studies. AERIS retained an exposure and provenance record for 26 attempts. Across the fitted, interpretable contrasts, 11 of 16 published positive directions and 14 of 21 null conclusions were retained. The pooled median study-level absolute log-effect difference was 0.169, the geometric

Table 7: Researcher-intervention ledger for the nine-study primary evaluation. Evidence: D, direct project conversation; R, reconstruction from study artifacts and prompt summaries; I, interpretation or reporting after fitting.

Study (evidence)	Researcher intervention	Scientific effect
France COVID-19 (D, R)	Continue the department-level run, broaden coverage, and disclose the use of 2019 monitor data for the article's 1999–2016 exposure period.	Increased coverage and made the period substitution explicit.
Italian cities COVID-19 (D)	Recover the largest defensible city sample; preserve the PM_{10} estimand; report the article's doubling contrast.	Changed coverage and presentation scale while preserving the target estimand.
São José admissions (D, R)	Continue after the SIH/SUS route failed; permit a labeled public-outcome fallback while retaining the lag structure.	Substituted SIVEP-Gripe outcomes and shortened the period; the analysis was labeled bounded.
São Paulo AMI (D, R)	Continue when the full pollutant-cluster specification could not be recovered; permit a reduced-input sensitivity analysis.	Reduced model fidelity while retaining an interpretable pollution–health analysis.
UK COVID-19 (R)	Recover all four national outcome series, fit the viable core model, and add public weather covariates where available.	Preserved the target geography and pollutant with the strongest executable public-data model.
US stroke mortality (D, R)	Continue the county analysis; permit a deterministic sample and public outcome proxy; match the article's per-10-μg/m^3 scale.	Changed outcome specificity and coverage and corrected the comparison scale.
Brazil cancer mortality (D, R)	Pursue national coverage and search regulator and public-monitor archives beyond the default API.	Broadened the analysis beyond the initial municipality-level attempt.
Northern Italy COVID-19 (D, R, I)	Use the article's NO_2 analysis after $PM_{2.5}$ coverage failed; classify its rounded 1.00 estimates as reported null effects.	Changed the pollutant within the article's analyses and set the meta-analysis classification.
OECD COVID-19 (R, I)	Continue country-level retrieval, retain countries meeting the declared coverage rule, and follow the article's null conclusion.	Reduced country coverage and set the qualitative classification.

Table 8: Study-level results from the first autonomous AERIS campaign. Effects are the principal retained contrast or contrasts; bounded analyses record material substitutions.

Study	Result	Principal retained effect
France COVID-19	Fit	RR 1.355 (1.193–1.540)
Italian cities	Bounded fit	Ratio 1.238 (1.020–1.503) per doubling
Northern Italy COVID-19	Bounded fit	Case RR 1.006; death RR 1.005
OECD COVID-19	Fit	Four RRs, 1.026–1.135
São José admissions	Bounded fit	O_3 RR 0.692; NO_2 RR 2.198
United Kingdom COVID-19	Bounded fit	OR 1.070 (1.068–1.073)
Brazil cancer mortality	Bounded fit	OR 0.971 (0.929–1.015)
United States stroke mortality	Bounded fit	Ratio 0.989 (0.973–1.005)
São Paulo AMI	Evidenced stop	Historical pollutant route not recovered

mean reconstructed-to-published ratio was 0.901, and study-level Pearson correlations were 0.507, 0.091, and -0.170 by campaign. Campaigns took 137, 151, and 91 minutes and generated 270,297, 329,971, and 225,718 output tokens, including retries. Autonomous AERIS completed 3.3 times as many attempts as raw Codex, but did not match the numerical consistency of the campaign with researcher review.

São Paulo exposed the remaining run-to-run variance. Three independent runs produced two bounded pollutant fits and one evidenced stop. The two fits also chose different public-data windows and returned different estimates. This variation reflects long-horizon source discovery and scientific judgment, not sampling uncertainty alone. It motivates reporting complete traces, substitutions, and failures rather than a single completion score.

Table 9: Study-level results from the first unaugmented-agent campaign. Ratios for completed fits use the article's reported contrast.

Study	Outcome	Effect or stopping point
Italy COVID-19	Fit	Published 1.76; Codex 1.64 (1.43–1.87)
OECD COVID-19	Fit	Published 1.02; Codex 1.08 (0.94–1.23)
UK COVID-19	Fit	Published 1.13; Codex 2.50 (2.44–2.57)
São Paulo AMI	Stopped	Pollution series and exact outcome definition unresolved
Northern Italy COVID-19	Stopped	Agent declined inference from two regional clusters
France COVID-19	Stopped	Historical $PM_{2.5}$ raster and adjustment panel unresolved
São José admissions	Stopped	Outcome retrieval failed; selected exposure source returned nulls
Brazil cancer mortality	Stopped	Historical mortality microdata retrieval failed
US stroke mortality	Stopped	Target outcome and national pollution coverage unresolved

H Unaugmented-agent ablation

Design. We tested whether the underlying agent could complete the nine fitted reanalyses without AERIS in three independent campaigns. All attempts ran in fresh containers using Codex CLI 0.155.1 with GPT-5.6 Terra at medium reasoning. The agent received the article bundle, common task contract, public network access, fixed retry rules, and structured output schema. It received no AERIS instructions, MCP tools, exposure engine, prior reanalysis code, or prior results. Runs had an eight-hour limit and no human researcher intervention.

Results. The three campaigns produced 3, 3, and 1 fitted effects, for 7 fits across 27 attempts. Six of seven retained the article's qualitative conclusion. Their median absolute log-effect difference was 0.073, their geometric mean reconstructed-to-published ratio was 1.011, and study-level Pearson correlations were 0.016 and -0.675 in the two campaigns with multiple fits. The third campaign produced one fit, so its correlation is undefined. The campaigns generated 103,007, 107,929, and 98,749 output tokens; the last includes no tokens for an attempt stopped at preflight. The first campaign's three fits appear in Table 9; the later campaigns added one Italian, three OECD, and one UK fit, with overlap across studies. The other 20 attempts stopped or failed before a fit even though AERIS had completed each task with researcher review. A preflight disk threshold stopped two attempts, which we classify as infrastructure-limited. In several scientific stops, the agent recognized an inadequate input but did not recover the source or bounded substitution that enabled the AERIS-assisted analysis.

Preserved artifacts. The supplementary archive preserves the 27 run manifests and each available transcript, generated program, retrieved source, analysis panel, model output or blocker record, and decision log. We store large run directories as checksum-verified archives. The archive also includes the container specification, package versions, and artifact checksums.