

# A provenance-aware semantic framework and evidence-aware knowledge graph for micro- and nanoplastic environmental-health evidence

Sudeepa Bhattacharyya

sxbhattach@gmail.com

University of Missouri–Kansas City <https://orcid.org/0000-0003-3493-6029>

Adrija Ghosh

University of Missouri–Kansas City

SumanKalyan Maity

Missouri University of Science and Technology

Maryam Salehi

University of Missouri

---

## Method Article

**Keywords:** Plastic pollution, semantic framework, knowledge graph, provenance, evidence relevance, environmental health, biomedical semantics

**Posted Date:** October 12th, 2026

**DOI:** <https://doi.org/10.21203/rs.3.rs-11309002/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

**Additional Declarations:** The authors declare no competing interests.

---

# Abstract

## Background

Micro- and nanoplastic (MNP) evidence spans environmental fate, experimental toxicology, human biomonitoring, and observational research. These studies differ substantially in particle characterization, exposure conditions, biological models, analytical methods, and study design, creating a risk that cross-study synthesis will lose the context that determines what each finding can support.

## Methods

We developed a provenance-aware semantic framework organized around five upper-level domains: Environment, Agent, Host, Evidence, and Governance, with Exposure Context modeled as a bridge linking environmental and particle conditions to biological systems. Source-reported Evidence Statements were separated from normalized Structured Observations through the provenance chain Publication  $\diamond$  Study  $\diamond$  Evidence Statement  $\diamond$  Observation. Evidence relevance was represented multidimensionally rather than as a single quality score. The framework was instantiated in a purposively selected 10-publication pilot corpus and evaluated for representational coverage, provenance completeness, competency-question performance, and context-sensitive interpretation.

## Results

The frozen schema contained 19 concept classes, 60 unique relationship predicates, and 74 permitted source-relation-target patterns. The pilot comprised nine primary studies and one narrative review. All nine primary study Observations retained complete Publication  $\diamond$  Study  $\diamond$  Evidence Statement  $\diamond$  Observation provenance (9/9, 100%). Six experimental biological studies supported mechanistic and/or model-specific conclusions. The two human studies supported tissue occurrence/internal burden or an exposure-outcome association. None of the nine primary studies supported causal human-health inference. Competency-question testing initially exposed representation errors involving particle-specific mechanism attribution and endpoint/tissue contextualization. After semantic refinement, all five competency questions returned the expected evidence and context.

## Conclusions

A semantic representation of MNP evidence can preserve provenance, exposure and biological context, and study-specific inferential limits within a common knowledge structure. The current implementation is evidence-aware rather than an automated inference-blocking system. It provides a transparent foundation for larger-scale validation and future computational evidence gating in AI-assisted environmental-health synthesis.

### 1. Background

Microplastics and nanoplastics (MNPs) are pervasive environmental contaminants with increasing relevance to human health. Human exposure occurs principally through ingestion and inhalation, and emerging evidence indicates systemic distribution and detection in multiple human biological matrices [1–3]. Recent studies have reported MNPs in carotid atherosclerotic plaques and in postmortem brain, liver, and kidney tissues [4, 5]. These observations have intensified interest in possible health consequences, but tissue detection or an exposure-outcome association does not by itself establish causation.

A central challenge in understanding MNP health effects is the heterogeneity of the evidence base. Relevant information is distributed across environmental occurrence and fate studies, in vitro experiments, animal toxicology, human biomonitoring, epidemiological studies, and evidence syntheses. Studies differ in polymer composition, particle size and shape, surface properties, weathering state, additives and co-contaminants, exposure route, concentration or dose, duration, biological model, and analytical methodology. These physicochemical and exposure characteristics can alter uptake, distribution, biological response, and the scientific relevance of a reported finding [1, 6]. Experimental toxicology has also relied heavily on convenient

model particles, especially pristine polystyrene, raising questions about extrapolation to heterogeneous, aged, or environmentally transformed particles [6].

Several MNP-specific approaches have therefore been developed to assess study quality, relevance, and utility for risk assessment. These approaches emphasize particle characterization, analytical adequacy, contamination control, experimental design, exposure relevance, and applicability to human-health or ecological risk assessment [7–10]. Recent evaluation of studies measuring microplastics (MPs) in human organs, tissues, and bodily fluids likewise identified substantial variability in sampling, extraction process, quality controls, particle characterization, analytical methodology, and interpretability [10]. Scientific usefulness, however, is not adequately represented by a single global quality label. A study may provide useful mechanistic evidence while being poorly suited to estimating human disease risk. Human tissue study may establish internal occurrence while providing little information about exposure route, intake rate, or causation.

This challenge is also one of knowledge representation. Conventional evidence tables can catalog study characteristics, but they are poorly suited to preserving the relational structure linking particle abundance and state, exposure conditions, biological model, measured effect, and source evidence. Knowledge graphs can represent heterogeneous entities and relationships while supporting traceability and complex querying, and their potential value for environmental-health evidence mapping has been recognized [11]. Existing resources such as the Environmental Conditions, Treatments, and Exposures Ontology (ECTO) provide interoperable representations of exposure events [12], while aggregate exposure pathway approaches have organized environmental and biological exposure information for MPs [13]. More broadly, provenance-, assertion-, and evidence-oriented semantic models emphasize that machine-readable scientific assertions should remain linked to their documentary support so that they remain auditable [14].

Entity connectivity alone, however, does not determine what a finding can legitimately support. A graph assertion such as “polystyrene associated with oxidative stress” may collapse evidence generated under markedly different particle characteristics, doses, exposure routes, biological systems, and study designs. A system can therefore retrieve fully traceable evidence and still overgeneralize across contexts. For MNP evidence, the representation must preserve not only the scientific relationship but also the conditions under which it was observed and the inferential limits imposed by those conditions.

We therefore developed a provenance-aware semantic framework and evidence-aware knowledge graph for MNP environmental-health evidence. The framework integrates five upper-level domains: Environment, Agent, Host, Evidence, and Governance, with Exposure Context represented as a contextual bridge. Scientific assertions are linked through Publication  $\diamond$  Study  $\diamond$  Evidence Statement  $\diamond$  Structured Observation, while evidence relevance is represented across distinct dimensions including particle/analytical adequacy, exposure context, biological-model relevance, evidence credibility, and provenance. We implemented the framework in a purposively selected 10-publication corpus to test whether heterogeneous MNP evidence could be represented while preserving provenance, explicit missingness, scientific context, and distinctions among mechanistic/model-specific, human-occurrence, human-association, and causal human-health interpretations. The objective was not to estimate MNP health effects from this small corpus, but to evaluate a computable representation of what studies report, the context in which findings were generated, and the types of conclusions those findings can appropriately support.

## 2. Methods

### 2.1 Study design and framework-development strategy

The framework was developed as a semantic representation of heterogeneous environmental-health evidence rather than as a formally validated ontology. Development followed an iterative knowledge-engineering process comprising domain scoping, competency-question definition, concept identification, relationship specification, pilot instantiation, scientific review, and refinement.

The organizing structure extends the traditional Agent-Host-Environment model with Evidence and Governance. Agent represents MNP physical and chemical characteristics. Host represents the exposed biological system and downstream biological or health responses. Environment represents the sources, settings, and media in which MNPs occur, undergo transformation, and are transported, such as air, water, soil, and food. Evidence represents the documentary and study-level basis of structured assertions.

Governance is a cross-cutting domain representing the rules and metadata needed to manage evidence eligibility, relevance assessment, review/adjudication, versioning, and auditability. It is not synonymous with governmental policy or regulation.

Exposure Context was modeled as a contextual bridge rather than an additional upper-level domain. It records the conditions connecting Environment and Agent to Host, including route, matrix or administration context, dose or concentration, duration, frequency, and other study-specific descriptors. In other words, Exposure Context links the environmental media and particle characteristics to the biological system through the reported exposure route, dose or concentration, and duration. Each pilot publication was assigned its own exposure record so that one study's dose, route, or duration could not be inadvertently attached to another study. Unknown or not-applicable exposure attributes were retained explicitly rather than inferred.

Candidate concepts and relations were repeatedly compared with the information required to represent particle characteristics, exposure conditions, biological models, tissues, mechanisms, biomarkers, outcomes, and provenance.

## 2.2 Competency questions

Competency questions were used to operationalize what the representation should be able to retrieve while preserving scientific context. Five questions were used in the final pilot evaluation: (CQ1) What mechanisms are associated with polystyrene particles? (CQ2) What biomarkers are linked to oxidative stress? (CQ3) What liver effects were reported after oral particle exposure? (CQ4) What human evidence supports a health outcome? and (CQ5) What is the full provenance chain supporting the Dong *et al.* [18] cellular conclusion?

These questions were selected to test particle specificity, biomarker retrieval, tissue/endpoint contextualization, human-evidence interpretation, and provenance tracing. When a competency question exposed a scientifically mis-scoped relationship, the relevant semantic structure was refined and the full set of questions was rerun.

## 2.3 Pilot corpus selection

A purposively selected corpus of 10 MNP publications was assembled to maximize variation in evidence source, particle characteristics, exposure conditions, biological model, and type of scientific interpretation. The corpus was designed to test representational coverage rather than to provide a representative sample or systematic review of the broader MNP literature.

The corpus comprised nine primary studies of which four were animal studies, two in vitro studies, one environmental fate/process study, one human tissue/biomonitoring study, and one prospective human observational study and the remaining one was a narrative review [4, 5, 15–22]. The narrative review was represented as a secondary evidence source rather than as equivalent to primary experimental or human evidence.

## 2.4 Extraction and scientific review

Pilot records were generated through a human-directed, prompt-guided extraction workflow. A structured extraction template was developed from the semantic-framework requirements and competency questions to capture publication and study characteristics, particle and analytical characteristics, exposure conditions, biological model and tissue, mechanisms, biomarkers, outcomes, source Evidence Statements, Structured Observations, provenance, and evidence-relevance fields. Full-text papers were processed with task-specific, schema-constrained prompts in GitHub Copilot/Visual Studio Code using GPT-5.6 Sol and/or Claude Sonnet 5 to generate candidate structured representations. A separate critique prompt was used during iterative refinement to identify missing, inconsistent, or overinterpreted content. Candidate outputs were reviewed against the source publications and corrected before inclusion in the final pilot dataset. Missing information was recorded as not reported, unknown, or not applicable rather than inferred.

An MNP domain expert reviewed designated scientific-relevance fields across the primary pilot corpus, including particle-characterization completeness, exposure relevance, biological-model relevance, limitations/provenance considerations, and supported-versus-not-supported interpretations. A separate 26-field extraction audit was later performed for the two human studies (Nihart *et al* [5] and Marfella *et al* [4]) to examine extraction fidelity in greater detail; this audit was not used as a primary evaluation endpoint in the present study. Because the framework, schema, and prompts were refined during development, the exact model and prompt version responsible for every individual extracted field was not consistently retained. Accordingly, this

study evaluates the reviewed and corrected semantic representation rather than the accuracy or reproducibility of the underlying AI extraction process.

## 2.5 Provenance-aware evidence model

A central design feature was the separation of source evidence from normalized representation. An Evidence Statement preserves the source-reported supporting text or result, whereas a Structured Observation represents a normalized, machine-readable finding derived from that evidence. Each primary-study Observation was linked to at least one Evidence Statement, which was linked to its Study and Publication, yielding the provenance chain Publication  $\diamond$  Study  $\diamond$  Evidence Statement  $\diamond$  Observation.

Where available, provenance also records the source location of the supporting evidence, such as a Results section, table, figure, or page. This design allows a structured assertion to be reconstructed back to the source passage from which it was derived. Provenance completeness was treated separately from scientific relevance. A traceable assertion may still be limited by particle characterization, exposure comparability, biological model, analytical method, or study design.

## 2.6 Multidimensional evidence-relevance and inference-boundary representation

Evidence relevance was represented as a multidimensional profile rather than a single quality score. Four principal scientific dimensions were retained: particle/analytical adequacy, exposure context, biological-model relevance, and evidence credibility, together with provenance and source-versus-reviewer status.

Particle/analytical adequacy includes polymer identity, size, shape, surface properties, weathering state, additives or co-contaminants, particle source, and the analytical method for particle identification. Exposure context includes route, concentration or dose, matrix or administration method, duration, frequency, and the degree to which the tested exposure can be compared with environmental or human exposure. Biological-model relevance records the study system, species or population, tissue/cell type, and physiological or disease context. Evidence credibility captures study design, controls, reporting completeness, replication where applicable, major limitations, and other features affecting interpretation.

Source-reported interpretation and framework assessment were stored separately. An author's description of an exposure as environmentally relevant was preserved as source information but did not automatically determine the framework's relevance assessment. Similarly, weathered or environmentally collected particles could increase environmental realism without converting a non-human experiment into human causal evidence.

Each Observation could therefore be assigned a supported inference class, such as environmental/process evidence, mechanistic or model-specific biological evidence, human exposure/occurrence evidence, human association evidence, or causal human-health inference. In the present proof of concept, these annotations classify and flag evidence boundaries. They do not computationally block graph traversal or generation.

## 2.7 Knowledge-graph implementation and interface

The semantic framework was instantiated as a directed knowledge graph with linked provenance and scientific-context components. The provenance component connects Publications, Studies, Evidence Statements, and Observations. The scientific component represents particle/material characteristics, Exposure Context, host/model, tissue, mechanisms, biomarkers, biological effects, and health outcomes. Evidence-relevance annotations remain linked to the Observation or its study context so that retrieval does not detach a scientific relationship from the conditions under which it was generated.

The schema-conformant explorer graph was treated as the canonical implementation for manuscript reporting. The interactive dashboard was regenerated against the finalized corpus and reconciled with the same terminology and provenance chain. It is presented as a user-facing exploratory interface rather than as an independently evaluated system. The public explorer is available at [https://ag230602.github.io/microplastic\\_graph/microplastics\\_explorer.html?author=bhattacharyya#home](https://ag230602.github.io/microplastic_graph/microplastics_explorer.html?author=bhattacharyya#home).

## 2.8 Evaluation strategy

The framework was evaluated descriptively across four dimensions: representational coverage, provenance completeness, competency-question performance, and evidence-relevance interpretation.

Representational coverage assessed whether the pilot records could explicitly represent the major information categories required by the framework, including particle/Agent characteristics, Exposure Context, biological model/Host, tissue, mechanism, biomarker, and biological effect or health outcome. Presence reflected explicit representation from source-derived information; unknown and not-applicable values were retained as such rather than imputed.

Provenance completeness was defined as the proportion of primary-study Structured Observations linked through a complete Publication → Study → Evidence Statement → Observation chain. Competency-question testing assessed whether the graph returned the expected scientific concepts while preserving particle specificity, exposure and tissue context, human-versus-experimental evidence type, null findings, and provenance. When a competency question exposed a representation error, the semantic relationships were refined and all five questions were rerun.

Evidence-relevance evaluation examined whether the pilot records distinguished environmental/process evidence, mechanistic/model-specific biological evidence, human occurrence evidence, human association evidence, and causal human-health inference. Because inference boundaries were not computationally enforced, evaluation focused on whether supported and unsupported interpretations could be represented explicitly and consistently.

### 3. Results

#### 3.1 Final semantic framework

The finalized framework organizes MNP environmental-health evidence around five upper-level domains: Environment, Agent, Host, Evidence, and Governance. Exposure Context functions as a bridge connecting environmental and particle characteristics with the biological system. Fate and transformation processes, including weathering, fragmentation, aggregation, and contaminant adsorption, are represented at the Environment-Agent interface. The Evidence domain preserves the provenance chain linking source documents to normalized Observations, while Governance captures the rules and metadata used to manage evidence inclusion, relevance assessment, validation status, versioning, and auditability.

The frozen computable schema contained 19 concept classes, 60 unique relationship predicates, and 74 unique permitted source-relation-target patterns (Supplementary Table S2). The five upper-level domains are organizing constructs and are not required to map one-to-one onto graph node classes. For example, Governance is implemented through rules, review fields, logs, and version/audit metadata rather than as a required node in every study.

The framework separates evidence source type from evidence relevance and supported interpretation. This allows the same study to be useful for one scientific purpose while remaining limited for another and prevents missing exposure or particle information from being silently inferred.

**Table 1. Core components of the MNP semantic framework.**

| Framework component | Type                             | Representative concepts                                                                                                         | Role in framework                                                                                                           |
|---------------------|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Environment         | Upper-level domain               | Environmental source, matrix, geographic/contextual setting, fate/transformation context                                        | Represents where MNPs originate, occur, transform, or are encountered.                                                      |
| Agent               | Upper-level domain               | Polymer type, particle size, shape, surface properties, weathering, additives/co-contaminants                                   | Represents physical and chemical characteristics of the MNP material.                                                       |
| Host                | Upper-level domain               | Population/model, species, tissue, mechanism, biomarker, biological effect, health outcome                                      | Represents the exposed biological system and downstream responses.                                                          |
| Evidence            | Upper-level domain               | Publication, Study, Evidence Statement, Observation, study design, evidence-source type                                         | Preserves the documentary and study-level basis of each structured finding.                                                 |
| Governance          | Upper-level/cross-cutting domain | Eligibility, metadata requirements, review/adjudication, relevance assessment, versioning, audit status                         | Defines how evidence is included, assessed, documented, updated, and interpreted.                                           |
| Exposure Context    | Contextual bridge                | Route, matrix/administration, dose or concentration, duration, frequency, regimen                                               | Connects Environment and Agent to Host while preserving conditions under which exposure occurred.                           |
| Evidence Relevance  | Cross-cutting assessment         | Particle/analytical adequacy, exposure relevance, biological-model relevance, credibility, provenance, permitted interpretation | Represents the scientific context and limits of use for each Observation without collapsing dimensions into a single score. |

### 3.2 Pilot corpus characteristics and representational coverage

The pilot corpus comprised nine primary studies and one narrative review. The primary evidence included four animal studies, two in vitro studies, one environmental fate/process study, one human tissue/biomonitoring study, and one prospective human observational study [4,5,15–21], the narrative review was retained as secondary evidence [22]. The corpus intentionally contrasted pristine and aged or field-collected particles, experimental and real-world contexts, different biological models, and evidence supporting mechanistic, environmental-process, occurrence, and association interpretations.

Representational coverage varied according to study type, as expected. Particle/Agent characteristics were explicitly represented in 9/10 publications. A study-specific Exposure Context record was instantiated for all 10 publications, where route, dose, or duration was not reported or was not applicable, that status was retained explicitly. Biological model/Host and tissue context were represented in 8/10 publications, mechanisms in 6/10, biomarkers in 7/10, and a biological effect or health outcome in 8/10 (Supplementary Table S2). The environmental fate study appropriately lacked a biological host, and the narrative review did not generate a primary Observation.

**Table 2. Characteristics of the purposively selected 10-publication pilot corpus.**

| Citation                               | Evidence type                           | Particle / exposure context                                                                                | Biological model                                  | Outcome type                                                                     | Framework feature demonstrated                                                                                       |
|----------------------------------------|-----------------------------------------|------------------------------------------------------------------------------------------------------------|---------------------------------------------------|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| Nihart et al. [5], 2025                | Human tissue / biomonitoring, primary   | Internally accumulated polymers, PE-dominant nanoscale particles, external exposure route and dose unknown | Human postmortem brain, liver, and kidney tissues | Tissue occurrence / internal burden, dementia-associated differences, noncausal  | Human tissue occurrence distinguished from causality, unknown exposure route retained explicitly                     |
| Li et al. [15], 2024                   | Animal toxicology, primary              | Virgin and photoaged PS, waterborne exposure, 0.1-100 µg/L to 120 hpf                                      | Zebrafish larvae                                  | Locomotion, motor-neuron signal, neurotransmission, and neurodevelopment markers | Particle-aging context and model-specific developmental / neurobiological effects                                    |
| Chen et al. [16], 2023                 | Animal toxicology, primary              | 100-nm PS, drinking-water exposure, 1 or 10 mg/L during gestation                                          | Pregnant mice and fetuses                         | Fetal growth, placental and fetal-muscle histology and omics                     | Dose-specific maternal exposure and tissue-specific outcomes                                                         |
| Pannetier et al. [17], 2020            | Animal toxicology, primary              | Field-collected PE/PP mixtures, dietary exposure for 14 or 30 days                                         | Japanese medaka larvae / juveniles                | Growth, survival, DNA damage, EROD, and swimming outcomes                        | Field-collected mixed-particle exposure, tissue/endpoint scoping and null-result retention                           |
| Dong et al. [18], 2020                 | In vitro toxicology, primary            | PS, $1.72 \pm 0.26$ µm, direct culture exposure, 10 and 1000 µg/cm <sup>2</sup> for 24 h                   | BEAS-2B human bronchial epithelial cells          | Oxidative, inflammatory, and barrier-related endpoints                           | Observation-level particle/exposure/model context, model-specific interpretation with human causal inference limited |
| Liang et al. [19], 2022                | Animal toxicology, primary              | 50-nm PS, gavage, 0.25-250 mg/kg/day for 28 days                                                           | Adult male C57BL/6J mice                          | BBB, mitochondrial/ATP, neuronal, and behavioral endpoints                       | Multiscale CNS response, mechanistic/model-specific evidence and negative disease-model findings retained            |
| Aghilinasrollahabadi et al. [20], 2021 | Environmental fate / process, primary   | LDPE and PET films, weathering and Pb/Zn sorption, no organismal exposure                                  | Not applicable                                    | Surface properties and metal-adsorption kinetics                                 | Abiotic fate/transformation represented without requiring a biological host or health endpoint                       |
| Marfella et al. [4], 2024              | Human clinical / observational, primary | PE/PVC detected in carotid plaque, internal burden, external route and duration unknown                    | Patients undergoing carotid endarterectomy        | Association with MI, stroke, or all-cause death, plaque inflammatory markers     | Prospective human association represented separately from causal inference                                           |

| Citation                       | Evidence type                        | Particle / exposure context                                             | Biological model                                        | Outcome type                                                            | Framework feature demonstrated                                                                     |
|--------------------------------|--------------------------------------|-------------------------------------------------------------------------|---------------------------------------------------------|-------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| Hwang et al. [21], 2019        | In vitro toxicology, primary         | PP size classes, powder/DMSO preparation, assay-specific concentrations | Human and rodent cell systems, sheep RBCs for hemolysis | Viability, cytokines, histamine, ROS, macrophage markers, and hemolysis | Particle-size/preparation dependence and cross-model cell-assay context                            |
| Kannan & Vimalkumar [22], 2021 | Narrative review, secondary evidence | Literature synthesis, no new exposure experiment                        | Not applicable                                          | Secondary synthesis / hypothesis, no independent primary outcome        | Secondary evidence kept separate from primary-study Observation and provenance-completeness counts |

Abbreviations: BBB, blood-brain barrier, EROD, ethoxyresorufin-O-deethylase, LDPE, low-density polyethylene, MI, myocardial infarction, PE, polyethylene, PET, polyethylene terephthalate, PP, polypropylene, PS, polystyrene, PVC, polyvinyl chloride, RBC, red blood cell, ROS, reactive oxygen species.

### 3.3 Provenance-aware evidence representation

The framework maintained a distinction between source-reported Evidence Statements and normalized Structured Observations. All nine primary-study Observations were linked to an Evidence Statement and retained a complete Publication à Study à Evidence Statement à Observation provenance chain (9/9, 100%). The narrative review was represented as secondary evidence and was not included in the primary-Observation denominator.

A worked example from Dong *et al.* [18] illustrates the structure. The source Evidence Statement reported that ROS accumulation increased significantly only in BEAS-2B cells exposed to PS microplastics at 1000 µg/cm<sup>2</sup>. The corresponding Structured Observation retained the 24-hour exposure duration, the 10 and 1000 µg/cm<sup>2</sup> test concentrations, the BEAS-2B model, relevant particle characteristics, and the absence of a significant ROS increase at 10 µg/cm<sup>2</sup>. The framework therefore supports the interpretation that high dose PS exposure produced a cellular oxidative stress response under the tested in vitro conditions, while not supporting causal human respiratory disease inference or equivalence to real world inhalation exposure.

### 3.4 Knowledge-graph proof-of-concept implementation

The finalized pilot records were instantiated in the schema-conformant explorer graph. The graph links provenance entities with particle, exposure, biological-model, tissue, mechanism, biomarker, and outcome concepts while retaining evidence-relevance context. After final schema reconciliation, the explorer and dashboard used aligned provenance terminology and publication/study records.

The purpose of the implementation was to demonstrate that the semantic framework could be instantiated and queried while preserving study-specific context, rather than to optimize graph size or network topology. Supplementary Table S2 summarizes the schema inventory, pilot coverage, provenance completeness, and final competency-question results.

### 3.5 Evidence relevance and translational interpretation

The pilot corpus demonstrated that study type alone was insufficient to determine the translational meaning of an Observation. Across the nine primary studies, six experimental biological studies supported mechanistic and/or model-specific biological conclusions. The two human studies supported different inference classes: Nihart *et al.* [5] provided human tissue-occurrence/internal-burden evidence, whereas Marfella *et al.* [4] provided prospective human association evidence. The environmental fate/process study supported particle-transformation and metal-sorption interpretation without providing biological or human-health evidence. None of the nine primary studies were classified as supporting causal human-health inference.

Domain review further showed why a single global quality score would be inadequate. Important study-specific considerations included incomplete or method-dependent particle confirmation, uncertainty in environmental exposure comparability, differences between pristine and aged or field-collected particles, potential effects of co-contaminants or leachates, non-monotonic dose-

response patterns, and design limitations affecting human interpretation (Supplementary Table S1). These considerations did not make the studies scientifically unusable, instead, they changed the scientific use for which an Observation could appropriately contribute evidence.

For example, Dong et al. [18] supported a model-specific cellular oxidative response at a high in vitro exposure but not human respiratory-disease causation. Li et al. [15] demonstrated differences between virgin and photodegraded PS in zebrafish, supporting particle-aging context and model-specific effects without establishing direct human-health inference. Nihart et al. [5] supported human internal occurrence and tissue burden, whereas Marfella et al. [4] supported an association between MNP-positive carotid plaques and subsequent clinical events, neither design established causation.

**3.6 Competency-question evaluation and iterative refinement**

Competency-question testing was used not only as a final check but as an iterative mechanism for identifying semantic representation errors. During development, missing or overly broad relationships affected several queries. Two scientifically consequential problems were identified: particle-specific mechanism attribution and endpoint/tissue contextualization.

For CQ1, mechanisms could initially be retrieved from studies containing PS even when the mechanism was speculative or arose from a mixed-particle exposure. The representation was revised so that a mechanism remained linked to the same Observation/Evidence Statement and particle context that supported it, speculative mechanisms and mixed-exposure findings were not returned as PS-specific evidence. For CQ3, a blood-based finding had been conflated with a liver endpoint and a null liver result had been omitted. The representation was revised so that tissue/sample, endpoint, exposure context, and result remained co-scoped within the same Observation, with null findings retained explicitly.

Following these refinements and final schema reconciliation, all five competency questions passed. CQ1 retrieved mechanisms supported specifically for PS exposures, CQ2 retrieved oxidative-stress biomarkers, CQ3 retained the liver-specific null EROD result after oral particle exposure without conflating a blood assay, CQ4 distinguished human occurrence/association evidence from causal inference, and CQ5 reconstructed the complete provenance chain for the Dong cellular conclusion.

**Table 3. Competency-question evaluation and iterative framework refinement.**

| Competency question                                                                    | Capability tested                            | Representation issue identified during development / refinement                                                                                                                                                                                                                          | Final result                                                                                                                |
|----------------------------------------------------------------------------------------|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| CQ1. What mechanisms are associated with polystyrene particles?                        | Particle-specific mechanism retrieval        | Speculative mechanisms and findings from mixed-particle exposures could be returned as PS-specific. Mechanism links were re-scoped to the supporting Observation/Evidence Statement and particle context, speculative and mixed-exposure findings were no longer treated as PS-specific. | <b>Pass.</b> Retrieved mechanisms were supported specifically for PS exposures.                                             |
| CQ2. What biomarkers are linked to oxidative stress?                                   | Biomarker retrieval                          | Biomarker representation was incomplete during schema development and was completed/refined before the final rerun.                                                                                                                                                                      | <b>Pass.</b> Oxidative-stress biomarkers were retrieved with their study context.                                           |
| CQ3. What liver effects were reported after oral particle exposure?                    | Tissue / endpoint / result contextualization | A blood-based finding had been conflated with a liver endpoint and a null liver EROD result had been omitted. Tissue/sample, endpoint, exposure context, and result were co-scoped within the same Observation, and null findings were retained.                                         | <b>Pass.</b> The liver-specific null EROD result was retained without conflating the unrelated blood assay.                 |
| CQ4. What human evidence supports a health outcome?                                    | Human evidence type and inferential boundary | The Publication/provenance layer and evidence-type distinctions were refined so that human tissue occurrence and observational association remained distinguishable from causal inference.                                                                                               | <b>Pass.</b> Human occurrence/association evidence was retrieved without being represented as causal evidence.              |
| CQ5. What is the full provenance chain supporting the Dong et al. cellular conclusion? | Source provenance tracing                    | The Publication layer was completed so that the full source chain could be reconstructed.                                                                                                                                                                                                | <b>Pass.</b> Publication à Study à Evidence Statement à Observation was reconstructed for the Dong et al. cellular finding. |

All five competency questions passed after iterative framework refinement. These results represent task-specific proof-of-concept performance and are not intended as formal ontology validation or as a general retrieval-accuracy estimate.

### 3.7 Interactive explorer and dashboard

An interactive web interface was developed as a proof-of-concept means of browsing the pilot corpus, Evidence Statements, Observations, scientific concepts, provenance paths, and graph views. The explorer is publicly accessible and identifies the frozen 10-paper pilot corpus used for this study. The dashboard was regenerated after schema reconciliation so that the displayed provenance terminology and publication/study records were consistent with the finalized framework.

The interface is presented as a demonstration of inspectability rather than as an independently evaluated user-facing system. No claims about usability, retrieval accuracy, or comparative interface performance are made in the present study.

## 4. Discussion

### 4.1 Principal findings and contribution

This study presents a provenance-aware semantic framework and evidence-aware knowledge-graph implementation designed for a specific challenge in MNP environmental-health research which is that findings that appear to concern the same particle, mechanism, or outcome may differ substantially in the conditions under which they were generated and therefore in the conclusions they can support. The framework represents both the scientific relationship and the evidence context that gives that relationship meaning.

The 10-publication proof of concept demonstrated three features central to the framework. First, heterogeneous environmental, experimental, and human evidence could be represented within a common structure without requiring missing information to be inferred. Second, all primary-study Observations retained complete source provenance. Third, the evidence-relevance

representation distinguished mechanistic/model-specific findings, human tissue occurrence, human association, and causal human-health inference rather than treating these as equivalent forms of evidence.

The competency-question results further illustrate the value of testing semantic models with scientifically meaningful queries. The initial representation could connect relevant entities while still producing scientifically incorrect scoping, such as attributing a mechanism to the wrong particle context or conflating a blood measurement with a liver endpoint. These errors became visible only when the framework was asked questions requiring contextual precision. Their correction therefore represents part of the framework-development process rather than merely a query-interface adjustment.

#### **4.2 Relationship to existing MNP evidence-assessment approaches**

The framework builds on prior MNP study-quality and relevance approaches that emphasize particle characterization, analytical methods, experimental design, exposure relevance, contamination control, and risk-assessment applicability [7–10]. These approaches are essential for critical appraisal, but they were designed primarily to help scientists evaluate study reliability and relevance rather than to preserve observation-specific inferential limits in a machine-readable structure during cross-study retrieval and synthesis.

The present framework is therefore complementary to, rather than a replacement for, existing quality-assessment systems. It does not generate a quantitative hazard score or global weight-of-evidence ranking. Instead, it attaches relevant particle, exposure, biological-model, credibility, and provenance context to each Observation so that a finding can be retrieved and compared without losing the conditions that constrain its interpretation.

#### **4.3 Provenance is necessary but not sufficient**

The explicit separation of Evidence Statements from Structured Observations is a second major contribution. A normalized assertion can be traced back to the supporting source passage, improving auditability and reducing the risk that structured knowledge becomes detached from its documentary basis. This is particularly important in AI-assisted workflows, where apparently precise structured output can obscure how much of an assertion was actually present in the source [14].

Provenance alone, however, does not establish that an assertion is analytically reliable, environmentally comparable, biologically relevant to a target question, or appropriate for causal interpretation. An answer may be perfectly traceable to a published paper and still overgeneralize that paper's findings. Separating provenance completeness from evidence relevance is therefore essential to the framework's purpose.

#### **4.4 Environmental-health translation and semantic representation**

The framework extends existing exposure-representation approaches by emphasizing the connection among particle state, environmental fate or transformation, exposure context, biological model, and evidentiary support. These distinctions are especially important for MNPs because particle physicochemical characteristics, environmental transformation, co-contaminants, and the realism of experimental exposure conditions can materially influence biological responses and their relevance to human health [1,6,9].

Representing these components together enables comparison across evidence types without collapsing them. Environmental fate studies can characterize particle transformation. Experimental studies can evaluate biological effects under specified conditions. Human biomonitoring can establish internal occurrence and observational studies can estimate associations with health outcomes. These forms of evidence may contribute to a translational chain, but each answers a different scientific question. A semantic framework that preserves those distinctions can make evidence gaps visible rather than implicitly filling them.

#### **4.5 Implications for AI-assisted evidence synthesis**

The distinction between evidence retrieval and evidence interpretation becomes increasingly important as generative AI and retrieval-augmented systems are used to summarize scientific literature. A retrieval system can return authentic, source-linked findings while still combining them into a conclusion stronger than the underlying studies support. Knowledge-graph connectivity likewise does not by itself establish that a multi-study reasoning path is scientifically justified.

The evidence-relevance annotations developed here provide a potential control layer between retrieval and interpretation. They make explicit which features of a finding should remain attached during synthesis and which inference classes are supported or not supported. In the present proof of concept, however, these annotations classify and flag evidence rather than computationally prevent an AI system from crossing an inference boundary. Formalizing the criteria into reproducible rules and prospectively testing evidence-gated AI synthesis are subsequent methodological steps.

#### **4.6 Limitations**

Several limitations should be considered. First, the 10-publication corpus was purposively selected to maximize contrast across evidence types and was not intended to represent the frequency, distribution, or overall strength of evidence in the broader MNP literature. The findings therefore should not be interpreted as a systematic review or quantitative summary of MNP health effects.

Second, the pilot contains one principal Structured Observation per primary study for framework evaluation. This was sufficient to test the provenance and evidence-relevance architecture but does not capture the full within-paper complexity of all reported outcomes. Larger implementations will require multiple Observation units per publication while maintaining correct within-study scoping.

Third, several evidence-relevance dimensions require expert scientific judgment. Although the primary corpus underwent domain review of designated relevance and interpretation fields, the complete corpus was not independently re-extracted by multiple reviewers, and formal inter-rater reliability was not estimated. Future work should operationalize the criteria more explicitly and test reproducibility across independent domain experts.

Fourth, the pilot was developed through an iterative human-directed, AI-assisted extraction workflow. Exact model and prompt provenance were not retained for every field of the frozen 10-paper dataset, so this study cannot be used to estimate the accuracy or reproducibility of AI extraction. The current evaluation concerns the final reviewed semantic representation, not model performance.

Fifth, the framework is a proof-of-concept semantic model rather than a formally validated ontology. The study did not undertake comprehensive logical-consistency testing, community ontology-governance procedures, or formal alignment across all classes with established external ontologies. Likewise, inference-boundary annotations were not implemented as computational hard gates.

Finally, the public interface was developed as an exploratory demonstration and was not independently evaluated for usability or retrieval performance.

#### **4.7 Future directions**

Future work should evaluate the framework across larger and more diverse collections of MNP literature and further assess the reproducibility of its evidence-relevance representations across independent studies and reviewers. Expansion will also require continued refinement of the semantic model as new particle characterization methods, exposure contexts, and evidence types emerge.

Additional work is needed to determine how provenance-aware and evidence-relevance information can be incorporated into AI-assisted literature synthesis while preserving study-specific context and limiting unsupported generalization. Formal evaluation of extraction, representation, and evidence-informed reasoning should therefore accompany future scaling of the framework.

Greater interoperability with established environmental-health terminologies and provenance standards will also be important. At larger scale, the framework may support systematic comparison of evidence across environmental, experimental, and human studies and help identify areas in which the available evidence remains incomplete or poorly connected across translational levels.

## **5. Conclusion**

We developed a provenance-aware semantic framework and proof-of-concept knowledge graph for representing heterogeneous MNP environmental-health evidence. The framework preserves the distinction among source evidence, normalized Observations,

exposure and biological context, and the scientific interpretations those Observations may support. In a 10-publication pilot, all primary-study Observations retained complete provenance, experimental and human evidence were represented under distinct inference classes, and competency-question testing exposed and helped correct context-specific representation errors. The present system is evidence-aware rather than an automated evidence-gating engine, but it provides a structured foundation for trustworthy AI-assisted synthesis in which scientific limits on interpretation can be made explicit, auditable, and computable.

## Declarations

### Ethics approval and consent to participate

This study analyzed information from published scientific literature and did not involve new human-participant or animal research, institutional ethics approval was therefore not applicable.

### Consent for publication

Not applicable.

### Availability of data and materials

Data supporting the analyses presented in this study are included within the article and its supplementary materials. The structured 10-publication pilot dataset is provided as Supplementary Table S1. Additional implementation materials are not publicly available at this stage because they are undergoing institutional intellectual-property review. Requests for additional information may be directed to the corresponding author.

### Competing interests

The authors declare that they have no competing interests.

### Funding

This work was supported by institutional resources and research infrastructure provided by the University of Missouri-Kansas City (UMKC) School of Medicine. No external grant funding was received specifically for this study.

### Authors' contributions

SB conceived the study, developed the overall semantic and evidence-relevance framework, supervised the work, contributed to methodology and scientific interpretation, and drafted and revised the manuscript. AG performed data curation, corpus structuring, software and knowledge-graph implementation, visualization, competency-question testing, and contributed to manuscript preparation. SM contributed computational-methods expertise, review of the knowledge-graph/AI implementation, and manuscript revision. MS provided MNP domain-science review of particle characterization, exposure relevance, biological-model relevance, methodological limitations, and supported-versus-not-supported interpretations, and contributed to manuscript revision. All authors should review and approve the final contribution statement before submission.

### Acknowledgements

Not applicable.

### Use of generative AI and AI-assisted tools

AI-assisted tools were used during portions of the structured extraction and manuscript-development workflow. All AI-assisted outputs used in the study were subject to human review and correction, and scientific interpretation, domain review, final verification, and responsibility for the manuscript remain with the authors.

## References

1. Bhattacharyya S, Greer ML, Salehi M (2025) Impact of micro- and nanoplastics exposure on human health: focus on neurological effects from ingestion. *Front Public Health* 13:1681776. 10.3389/fpubh.2025.1681776
2. World Health Organization. Dietary and inhalation exposure to nano- and microplastic particles and potential implications for human health. Geneva: World Health Organization (2022) ISBN 978-92-4-005460-8
3. Leslie HA, van Velzen MJM, Brandsma SH, Vethaak AD, Garcia-Vallejo JJ, Lamoree MH (2022) Discovery and quantification of plastic particle pollution in human blood. *Environ Int* 163:107199. 10.1016/j.envint.2022.107199
4. Marfella R, Prattichizzo F, Sardu C et al (2024) Microplastics and nanoplastics in atheromas and cardiovascular events. *N Engl J Med* 390(10):900–910. 10.1056/NEJMoa2309822
5. Nihart AJ, Garcia MA, El Hayek E et al (2025) Bioaccumulation of microplastics in decedent human brains. *Nat Med* 31(4):1114–1119. 10.1038/s41591-024-03453-1
6. Gouin T, Ellis-Hutchings R, Pemberton M et al (2024) Addressing the relevance of polystyrene nano- and microplastic particles used to support exposure, toxicity and risk assessment: implications and recommendations. *Part Fibre Toxicol* 21:39. 10.1186/s12989-024-00599-1
7. de Ruijter VN, Redondo-Hasselerharm PE, Gouin T, Koelmans AA (2020) Quality criteria for microplastic effect studies in the context of risk assessment: a critical review. *Environ Sci Technol* 54(19):11692–11705. 10.1021/acs.est.0c03057
8. Gouin T, Ellis-Hutchings R, Thornton Hampton LM, Lemieux CL, Wright SL (2022) Screening and prioritization of nano- and microplastic particle toxicity studies for evaluating human health risks - development and application of a toxicity study assessment tool. *Microplast Nanoplast* 2:2. 10.1186/s43591-021-00023-x
9. Vogel A, Tentschert J, Pieters R et al (2024) Towards a risk assessment framework for micro- and nanoplastic particles for human health. *Part Fibre Toxicol* 21:48. 10.1186/s12989-024-00602-9
10. Lane T, de Ruijter VN, Wardani I, Barchiesi M, Koelmans AA (2026) A QA/QC screening framework to assess studies reporting microplastic presence in human organs, tissues, and bodily fluids. *Environ Int* 214:110429. 10.1016/j.envint.2026.110429
11. Wolffe TAM, Vidler J, Halsall C, Hunt N, Whaley P (2020) A survey of systematic evidence mapping practice and the case for knowledge graphs in environmental health and toxicology. *Toxicol Sci* 175(1):35–49. 10.1093/toxsci/kfaa025
12. Chan LE, Thessen AE, Duncan WD et al (2023) The Environmental Conditions, Treatments, and Exposures Ontology (ECTO): connecting toxicology and exposure to human health and beyond. *J Biomed Semant* 14:3. 10.1186/s13326-023-00283-x
13. Peng G, Lin Y, van Bavel B, Li D, Ni J, Song Y (2022) Aggregate exposure pathways for microplastics (mpAEP): an evidence-based framework to identify research and regulatory needs. *Water Res* 209:117873. 10.1016/j.watres.2021.117873
14. Chhetri TR, Halchenko YO, Jarecka D (2025) Bridging the scientific knowledge gap and reproducibility: a survey of provenance, assertion and evidence ontologies. *Proc Int World Wide Web Conf Companion* 2025:924–928. 10.1145/3701716.3715483
15. Li X, Zheng T, Zhang J, Chen H, Xiang C, Sun Y et al (2024) Photoaged polystyrene microplastics result in neurotoxicity associated with neurotransmission and neurodevelopment in zebrafish larvae (*Danio rerio*). *Environ Res* 250:118524. 10.1016/j.envres.2024.118524
16. Chen G, Xiong S, Jing Q, van Gestel CAM, van Straalen NM, Roelofs D et al (2023) Maternal exposure to polystyrene nanoparticles retarded fetal growth and triggered metabolic disorders of placenta and fetus in mice. *Sci Total Environ* 854:158666. 10.1016/j.scitotenv.2022.158666
17. Pannetier P, Morin B, Le Bihanic F, Dubreil L, Clérandeau C, Chouvellon F et al (2020) Environmental samples of microplastics induce significant toxic effects in fish larvae. *Environ Int* 134:105047. 10.1016/j.envint.2019.105047
18. Dong CD, Chen CW, Chen YC, Chen HH, Lee JS, Lin CH (2020) Polystyrene microplastic particles: in vitro pulmonary toxicity assessment. *J Hazard Mater* 385:121575. 10.1016/j.jhazmat.2019.121575
19. Liang B, Huang Y, Zhong Y, Li Z, Ye R, Wang B et al (2022) Brain single-nucleus transcriptomics highlights that polystyrene nanoplastics potentially induce Parkinson's disease-like neurodegeneration by causing energy metabolism disorders in mice. *J Hazard Mater* 430:128459. 10.1016/j.jhazmat.2022.128459
20. Aghilinasrollahabadi K, Salehi M, Fujiwara T (2021) Investigate the influence of microplastics weathering on their heavy metals uptake in stormwater. *J Hazard Mater* 408:124439. 10.1016/j.jhazmat.2020.124439

21. Hwang J, Choi D, Han S, Choi J, Hong J (2019) An assessment of the toxicity of polypropylene microplastics in human derived cells. *Sci Total Environ* 684:657–669. 10.1016/j.scitotenv.2019.05.071

22. Kannan K, Vimalkumar K (2021) A review of human exposure to microplastics and insights into microplastics as obesogens. *Front Endocrinol (Lausanne)* 12:724989. 10.3389/fendo.2021.724989

Figures




Figure 1

**Overall semantic framework for micro- and nanoplastic environmental-health evidence.**The framework organizes heterogeneous evidence across Environment, Agent, Host, Evidence, and Governance, with Exposure Context linking particle/environmental conditions to biological systems and provenance preserving the source basis of each Structured Observation.





Figure 2

**From source evidence to evidence-bounded scientific interpretation.** A verified Dong et al. [18] example illustrates the provenance chain and the contextual information used to distinguish interpretations supported by the study from conclusions that extend beyond the available evidence.

## A KNOWLEDGE GRAPH ARCHITECTURE



## B PROOF-OF-CONCEPT IMPLEMENTATION




Figure note: The scientific layer shows contextual relationships used for representation and retrieval, not a mandatory causal pathway.

Figure 3

Knowledge-graph architecture and proof-of-concept implementation of the semantic framework.




**Figure 4**

**Four contrasting evidence-relevance patterns from the pilot corpus illustrating how particle state, exposure context, biological model, and study design constrain scientific interpretation.**

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryTableS1.xlsx](#)
- [SupplementaryTableS2.docx](#)