From Analytical Chemistry to Predictive Toxicology: A Review of Data Pipelines Linking Environmental Measurement to Health Prediction

[1]Olatunde Adesoji Aladegba, [2]Moshood Ayinde
[1]Raven Energy, Nigeria; oaladegba@gmail.com

[2]Western Illinois University – Department of Biology, IL, USA; moshoodayinde35@gmail.com

Corresponding Author; oaladegba@gmail.com

Abstract

Predictions of environmentally induced disease depend on a chain that begins with instrument signals and ends with a consequential interpretation. Errors often arise at the joints between stages, where units, censoring decisions, quality flags, provenance, and uncertainty are lost. This critical integrative review treats the environmental-health data pipeline as the object of study. It organizes the chain into five stages: measurement and quantification, quality control and harmonization, feature engineering and integration, predictive modeling, and interpretation and decision support. Polychlorinated biphenyl analysis by GC-ECD and GC-MS provides a worked example of how calibration, recovery correction, blank subtraction, detection-limit handling, panel mismatch, and congener aggregation propagate into downstream predictions. Evidence is assessed according to whether it preserves measurement meaning, provenance, uncertainty, and decision validity across stage boundaries. The review does not claim systematic coverage or estimate pooled effects. It uses direct environmental-health evidence for toxicological claims and clearly separates lessons drawn by analogy from diagnostic laboratories, data governance, and asset-integrity engineering. The analysis identifies four requirements for trustworthy pipelines: machine-actionable provenance, explicit interface contracts, versioned and tested workflows, and end-to-end uncertainty propagation. It further argues that reproducibility should be evaluated at every joint rather than inferred from agreement in final model outputs. The practical research agenda is to conduct multi-laboratory pipeline challenges using shared materials and raw files, then benchmark reproducibility, transportability, auditability, calibration, and decision consequences. The central claim is that strengthening the connective infrastructure of the pipeline will often improve predictive trustworthiness more than incremental refinement of the final algorithm.

Keywords

polychlorinated biphenyls, GC-ECD, GC-MS, data pipeline, predictive toxicology, FAIR principles, provenance, exposome, oxidative stress, congener analysis, reproducibility, uncertainty quantification

DOI: 10.21590/10.03.15

1. Introduction

Environmental health science has quietly become a computational discipline. A claim that a population exposed to a class of persistent organic pollutants faces an elevated burden of oxidative stress, DNA damage or cancer is now, in most contemporary studies, the terminal output of a long computational chain rather than a direct observation. The mechanistic plausibility of such claims is itself well established, as persistent organic pollutants are increasingly understood through a framework of environmental insults that converge on oxidative and genotoxic pathways (Peters et al., 2021). Somewhere near the beginning of that chain, a physical sample, blood serum, adipose tissue, house dust, sediment, breast milk, is introduced to an instrument, and a detector produces a signal. Somewhere near the end, a model returns a number, a probability, a hazard ratio, a risk category, that is interpreted as evidence about human health. Between those two points lies a sequence of transformations, each of which is a scientific act with its own assumptions, its own error structure and its own opportunities for silent corruption. The thesis of this review is straightforward and, we contend, underappreciated: a prediction of environmentally induced disease is only as trustworthy as the pipeline that produced it.

This framing deserves emphasis because the field's attention is distributed unevenly across the chain. The modeling stage attracts intense methodological scrutiny. New algorithms are proposed, benchmarked, cross-validated and debated; questions of overfitting, feature selection, external validation and interpretability are the subject of a mature and growing literature. By contrast, the stages that precede modeling, the analytical measurement, the quality-control regime, the harmonization of heterogeneous datasets, the engineering of features, are frequently compressed into a single terse paragraph of a methods section, or delegated to a supplementary file, or omitted entirely. The implicit assumption is that these are solved, routine, uninteresting matters of laboratory hygiene. We argue the opposite: that the largest and least-examined risks to predictive validity live precisely in these upstream stages and, above all, in the transitions between them.

1.1 The chain from instrument to health prediction

Consider the canonical example that recurs throughout this review: the estimation of health risk from exposure to polychlorinated biphenyls. PCBs are a family of 209 structurally distinct congeners, differing in the number and position of chlorine substituents on the biphenyl backbone. They are persistent, bioaccumulative and, for a subset of congeners, mechanistically implicated in oxidative stress and genotoxicity through pathways involving reactive metabolites, quinone formation and the generation of reactive oxygen species. To predict a health outcome from PCB

exposure, one must first *measure* PCBs, and measurement here is not a single number but a profile across congeners, obtained by gas chromatography and quantified either by electron-capture detection or by mass spectrometry.

The instrument does not report concentrations. It reports peak areas, retention times and, in the case of GC-MS, mass spectra. Converting a peak area into a concentration in nanograms per gram of lipid requires a calibration curve, an internal-standard correction, a recovery adjustment based on surrogate spikes, a blank subtraction and a decision about how to treat signals that fall below the limit of detection. Each of these is a modeling choice in miniature. The resulting concentration is then combined with concentrations of other congeners, perhaps aggregated into a summed measure or a toxic-equivalency-weighted total, perhaps expressed as a ratio designed to serve as an indicator of a particular exposure source or metabolic process. That engineered feature is joined, often across separate data systems, to biomarkers of effect, an oxidative-stress marker such as a measure of lipid peroxidation or oxidized guanine, a comet-assay readout of DNA strand breaks, or a clinical outcome. Only then does a predictive model consume the assembled matrix and return a health prediction, which a human being must finally interpret and act upon.

At every arrow in that sentence, a quantity crosses a boundary. It crosses from the instrument's native units into a reported unit; from a laboratory's internal congener nomenclature into a shared panel; from one study's non-detect convention into another's; from a raw measurement into a derived indicator; from a data table into a model's feature vector; from a model's output into a sentence a clinician or regulator reads. These crossings are the joints of the pipeline, and joints are where load concentrates.

1.2 Why the pipeline, not the model, is the object of study

There is a seductive but mistaken mental model in which analytical chemistry produces clean, objective numbers that are then handed off to statisticians, who bear responsibility for everything that follows. This division of labour encourages each party to treat the interface as someone else's problem. The chemist certifies that the measurement is fit for the chemist's purpose, typically regulatory compliance or method validation, without knowing that it will later be differenced against another laboratory's measurement to compute a ratio. The modeller receives a spreadsheet of concentrations and treats them as ground truth, unaware that half the values are substituted imputations for non-detects, or that recovery correction was applied inconsistently across batches. Neither party is negligent within their own frame; the failure is structural, a failure of the joint.

Treating the pipeline as the object of study means asking a different set of questions. Not merely *is this measurement accurate?* but *does the meaning of this measurement survive its journey to the model?* Not merely *does this model predict well on held-out data?* but *would the same input data, re-processed by the documented procedure, yield the same features, and would the same features, fed to the documented model, yield the same prediction?* These are questions of reproducibility, transferability and provenance, and they are properties of the pipeline as a whole, not of any single stage.

1.3 Motivation and scope

The motivation for this reframing is both scientific and practical. Scientifically, the reproducibility crisis that has unsettled many empirical fields has an acute form in environmental health, where datasets are heterogeneous, measurement is expensive and irreproducible, and the causal distance between exposure and outcome is large. A predictive claim that cannot be reproduced from documented inputs and procedures is not a scientific result in the strong sense; it is an anecdote with a confidence interval. Practically, environmental-health predictions increasingly inform regulation, remediation prioritization, clinical counselling and litigation. Decisions with real consequences are being made on the output of pipelines whose internal joints have never been audited.

This review is organized around the five-stage pipeline. Section 2 introduces the stages and, at its outset, a summary table mapping each stage to its characteristic reproducibility risks; the subsections then examine each stage and, crucially, the joint that follows it. Section 3 develops the infrastructural response, FAIR principles, pipeline-as-software, provenance and containerized workflows. Section 4 articulates the properties of a transferable pipeline and connects them to the modernization of United States environmental-health data infrastructure. Section 5 draws brief comparative lessons from diagnostic-laboratory informatics and from materials and asset-integrity measurement. Section 6 concludes. Throughout, we reference established ideas, the FAIR principles (Wilkinson, 2016), computational-toxicology blueprints (Thomas, 2019), the exposome (Rappaport & Smith, 2010; Escher, 2017), generically and without inventing quantitative claims; our contribution is architectural and conceptual rather than a meta-analysis of effect sizes.

1.4 Review design and evidence-selection logic

This article uses a critical integrative review design rather than a systematic review or meta-analysis. Its purpose is to identify recurrent failure mechanisms across the measurement-to-decision chain and to synthesize design requirements that are dispersed across analytical chemistry, environmental epidemiology, computational toxicology, data stewardship, and workflow engineering. The unit of analysis is therefore the pipeline joint, defined as a transfer of data, metadata, assumptions, or uncertainty from one stage to the next.

Evidence was retained when it informed at least one of four analytic functions: measurement validity, harmonization and semantics, reproducible computation and provenance, or interpretation of predictions in consequential settings. Direct environmental-health evidence was given priority for claims about exposure and toxicology. Literature from diagnostic laboratories, enterprise data governance, and asset-integrity engineering was used only to identify transferable control practices and is labeled as domain analogy rather than direct toxicological evidence. Sources with no clear relationship to a pipeline failure, control, or decision were not used to support causal claims.

The review does not claim exhaustive coverage, estimate pooled effect sizes, or rank algorithms by predictive performance. Its rigor rests on a transparent analytic framework applied consistently across stages: identify the artifact transferred at each joint, specify the information required to

preserve its meaning, describe the failure produced when that information is lost, and state a verifiable control. This structure also exposes the review to falsification, because each proposed control can be tested against cross-laboratory reproducibility, transportability, and audit outcomes.

1.5 Evidence base and boundaries of cross-domain synthesis

The pipeline framing developed here sits at the intersection of several applied literatures that have, in their own domains, confronted the problem of moving measurements reliably from an instrument to a consequential decision. We gather that adjacent work here into four thematic groups before returning, in Section 2, to the environmental-health pipeline itself. These bodies of work are not toxicological, but each illuminates a facet of the connective-tissue problem, laboratory infrastructure, data governance, physical-measurement integrity, and process discipline, that recurs throughout this review.

1.5.1 Diagnostic-laboratory informatics and infrastructure

A substantial body of work has examined how the physical and informational infrastructure of diagnostic laboratories determines the reliability and reach of the measurements they produce, spanning laboratory spatial planning, regulatory-compliant molecular and pathology facility design, materials and energy strategy, risk-managed construction, blood-collection network design, lifecycle and resilience planning, digital and BIM-enabled coordination, healing-centred design, and the translation of infrastructure investment into measurable diagnostic outcomes (Ogbete et al., 2018; Ogbete et al., 2019; Ogbete et al., 2020; Ogbete et al., 2021; Ogbete et al., 2022; Ogbete et al., 2023; Ogbete & Aminu-Ibrahim, 2024; Aminu-Ibrahim & Ogbete, 2023; Aminu-Ibrahim et al., 2018; Aminu-Ibrahim et al., 2019; Aminu-Ibrahim et al., 2020; Aminu-Ibrahim et al., 2024). Complementary work on diagnostic workflow integration in infectious-disease management and on quantitative molecular assays such as droplet digital PCR speaks directly to how measurement quality and interpretive workflow interact at the bench (Asiedu, 2023a; Asiedu, 2024a). This literature is relevant because the diagnostic laboratory, like the environmental-health pipeline, must transmit the meaning of a measurement across institutional and instrumental boundaries without corruption.

1.5.2 Enterprise data governance and secure analytics

A second cluster addresses the governance, protection and auditability of data as it flows through enterprise systems, precisely the concerns that the FAIR and provenance discussion of Section 3 reframes for scientific pipelines. This work spans legal and ethical risk modeling in data-protection governance, comparative data-protection regulation and secure cloud implementation, enterprise log analytics and automated incident response, identity and access governance across hybrid and multi-cloud environments, infrastructure-as-code governance, data-sensitivity classification and regulatory traceability, data-lakehouse and data-loss-prevention governance, HIPAA-compliant secure-analytics architecture, continuous auditing and cloud-misconfiguration monitoring, API governance, business-intelligence and predictive-analytics risk architecture, and fraud-detection and due-diligence analytics for capital markets, together with the application of artificial

intelligence to software testing and regression control (Mbonu et al., 2018; Mbonu et al., 2019a; Mbonu et al., 2019b; Mbonu et al., 2020a; Mbonu et al., 2020b; Mbonu et al., 2020c; Mbonu et al., 2021a; Mbonu et al., 2021b; Mbonu et al., 2021c; Mbonu et al., 2022a; Mbonu et al., 2022b; Mbonu et al., 2022c; Aliliele et al., 2023a; Aliliele et al., 2023b; Aliliele et al., 2023c; Aliliele et al., 2024a; Aliliele et al., 2024b; Borah et al., 2022; Eboh & Aliliele, 2023; Eboh & Aliliele, 2024). The traceability, access-control and regulatory-compliance disciplines developed in this literature are the enterprise analogue of the machine-actionable provenance and enforced-interface contracts this review prescribes for the scientific joints. Earlier work on CI/CD controls, ETL integration, regulated-healthcare system governance, Git-based review, and multi-cloud integration offers transferable controls for versioning, lineage, interface testing, and accountable change management (Badmus et al., 2018; Badmus et al., 2019a; Badmus et al., 2019b; Badmus et al., 2022a; Badmus et al., 2022b).

1.5.3 Materials characterization and asset integrity

A third cluster concerns physical-measurement integrity in materials science and engineering asset management, where instrument-derived measurements of composition, degradation and structural soundness feed predictions of remaining useful life and failure risk. This work includes non-destructive testing and integrity management for refinery process equipment, NDT for weld integrity in power infrastructure, welding-inspection standards and quality assurance, green corrosion inhibition and environmental degradation of materials, and charge-transport, recombination and instability characterization in thin-film and optoelectronic devices (Atta et al., 2018; Atta et al., 2020; Atta et al., 2021; Atta et al., 2022; Atta et al., 2024). We return to this domain in Section 5.2 because its mature culture of measurement uncertainty and traceability offers a direct template for environmental-health practice.

1.5.4 Lean process and quality management

A fourth cluster examines process design, standardization and quality management in resource-constrained and small-enterprise settings, the organizational discipline that Section 4.3 argues must accompany any technical interface contract. This work spans multi-stakeholder governance alignment, translation of regulatory requirements into compliance workflows, Lean Six Sigma adaptation for small enterprises, standard operating procedures as strategic assets, user-acceptance testing and validation, CRM and workflow automation, data-driven process optimization and bottleneck identification, client-onboarding cycle-time reduction, change management in digital transformation, and integrated lean-digital scaling frameworks (Eyetsemitan et al., 2020; Eyetsemitan et al., 2021; Eyetsemitan et al., 2022; Eyetsemitan et al., 2023a; Eyetsemitan et al., 2023b; Eyetsemitan et al., 2024a; Eyetsemitan et al., 2024b; Ambali et al., 2021; Oyeleye et al., 2022). Its emphasis on documented, repeatable procedure as the substrate of reliability mirrors the pipeline-as-software argument of Section 3.2.

1.5.5 Distributed sensing and pipeline infrastructure

Distributed low-power sensing and energy-autonomous data-acquisition methods relevant to measurement-to-decision infrastructure include (Asiedu & Quainoo, 2023b; Asiedu & Quainoo, 2023c; Asiedu & Quainoo, 2024b; Quainoo et al., 2024). Machine-learning-enabled IoT energy management also illustrates how sensing, prediction, and resource constraints must be treated as one pipeline rather than separate engineering problems (Kumuyi et al., 2024).

2. The five-stage pipeline

We model the environmental-health data pipeline as five sequential stages. The decomposition is deliberately coarse; real workflows contain finer substeps and occasional loops. But the five-stage view has the virtue of making the joints explicit, because each stage hands a well-defined artifact to the next: the measurement stage hands quantified concentrations to the quality-control stage; quality control hands harmonized, flagged data to feature engineering; feature engineering hands an analysis-ready feature matrix to modeling; modeling hands predictions and uncertainties to interpretation. The reliability of the whole is the product, not the sum, of the reliability of these hand-offs.

The table below summarizes the five stages, the artifact each produces, the dominant reproducibility risks that arise within the stage, and, most importantly, the joint-level failure that arises when its output is passed downstream. The remainder of Section 2 elaborates each row.

Table 1. Stages of the environmental-measurement-to-prediction pipeline, with principal reproducibility risks.

Stage	Function	Principal reproducibility risk
1. Measurement & quantification	Instrument analysis; calibration; congener quantification	Instrument/method variation; calibration drift; detection-limit handling
2. QC & harmonization	Recovery, blanks, units; cross-lab comparability	Missing QC metadata; inconsistent units and congener panels
3. Feature engineering & integration	Derive features; link exposure, biomarker, outcome data	Ad-hoc, undocumented transforms; record-linkage errors
4. Predictive modelling	Map features to oxidative-stress / damage / risk endpoints	Leakage; overfitting; no external validation
5. Interpretation & decision support	Translate predictions into actionable outputs	Opaque models; poor uncertainty communication

2.1 Measurement and quantification

The pipeline begins with an analytical measurement, and for PCBs the two workhorse techniques are gas chromatography with electron-capture detection and gas chromatography coupled to mass spectrometry (Ballschmiter et al., 1992; Muir & Sverko, 2006). Both separate the congeners in time as they elute from a capillary column; they differ in how they detect and identify the separated

compounds. GC-ECD exploits the strong electron affinity of chlorinated compounds: the detector measures the reduction in a baseline current as electron-capturing analytes pass through, yielding a signal that is highly sensitive to organochlorines and comparatively inexpensive to operate. GC-MS instead fragments and ionizes each analyte and measures mass-to-charge ratios, providing structural confirmation and the ability to resolve co-eluting compounds by their spectra, at higher cost and, historically, somewhat lower sensitivity for some congeners, though modern instruments narrow that gap (Ballschmiter et al., 1992). The same mass-spectrometry data-processing burden, peak detection, retention-time alignment and identification, recurs in the untargeted chemical and metabolomic assays that populate the exposome, where dedicated processing workflows have become standard (Smith et al., 2006; Dunn et al., 2011).

The two techniques embody a classic trade-off with direct downstream consequences. GC-ECD is exquisitely sensitive but structurally blind: it reports that *something* electron-capturing eluted at a given retention time, and the analyst must assume, on the basis of retention-time matching against standards, that the something is the target congener. When two congeners co-elute, as many PCB congeners do on common column chemistries, GC-ECD cannot distinguish them and reports a lumped signal (Muir & Sverko, 2006). GC-MS, by selecting characteristic ions, can often deconvolve co-elutions and confirm identity, reducing the risk of misassignment but introducing its own quantification subtleties around ion-ratio confirmation and matrix effects. The choice between them is therefore not merely instrumental; it determines which congeners can be individually quantified and which are reported only as co-eluting groups, and this determines the granularity of every feature that can be engineered downstream.

2.1.1 Calibration

Neither detector reports concentration directly. Quantification rests on a calibration relationship between instrument response and known analyte amount, established by analyzing standards of certified concentration. The analyst must choose the calibration model, single-point, multi-point linear, weighted linear or quadratic, and the calibration range, and must decide whether response is normalized to an internal standard. These choices are consequential. A calibration curve that is linear over the range of the standards may be extrapolated, implicitly and without warning, when a sample's response exceeds the highest calibrator; the resulting concentration is a model-based guess dressed as a measurement. Weighting schemes matter disproportionately at the low end, where many environmental samples sit and where the health-relevant signal often lives. Two laboratories analyzing identical samples with identical instruments can report systematically different concentrations solely because one used unweighted and the other used $1/x^2$-weighted regression.

2.1.2 Non-detect handling

Environmental concentrations are frequently near the detection threshold, and a substantial fraction of measurements return as non-detects, signals indistinguishable from the blank. How these are handled is one of the most consequential and least-visible decisions in the entire pipeline.

Common conventions include substituting zero, substituting the limit of detection, substituting the limit of detection divided by two or by the square root of two, or treating the values as left-censored and estimating the underlying distribution by maximum likelihood or regression-on-order-statistics methods. These conventions can differ by an order of magnitude in the imputed value, and because non-detects cluster at the low-exposure end of the distribution, the choice directly shapes the exposure gradient against which a health effect is estimated. A study reporting a dose-response relationship for a lowly-concentrated congener may be reporting, in part, an artifact of its substitution rule.

2.1.3 Recovery and surrogate correction

Analytes are lost during extraction, cleanup and concentration. To correct for these losses, analysts spike samples before extraction with surrogate standards, isotopically labelled or otherwise resolvable analogues, and measure their recovery (Muir & Sverko, 2006). Reported concentrations are then divided by the fractional recovery to estimate the true amount. This correction is standard and defensible, but it introduces variability: recovery differs by congener, by matrix and by batch, and the decision of whether to apply congener-specific, class-averaged or no recovery correction is a fork in the pipeline. A dataset in which some batches were recovery-corrected and others were not, perhaps because they were generated years apart under different standard operating procedures, carries a batch effect that masquerades as biological signal.

2.1.4 The measurement-to-QC joint

The joint that follows measurement is the point at which a set of quantification decisions, calibration model, non-detect rule, recovery convention, is frozen into a table of numbers and passed onward, typically stripped of the metadata that would allow those decisions to be reconstructed. The downstream analyst receives a column labelled, say, *PCB-153 (ng/g lipid)* and has no way, from the number alone, to know whether a value near the detection limit is a genuine measurement or an imputed substitution, whether it was recovery-corrected, or whether it represents PCB-153 alone or PCB-153 co-eluting with PCB-132. The joint fails not because the measurement is wrong but because its meaning is not transmitted with it. Everything Section 2.2 attempts to do is, in effect, an attempt to reconstruct meaning that should have travelled with the number in the first place.

2.2 Quality control and harmonization

Quality control is the stage at which the pipeline attempts to establish that the measurements are what they claim to be, and harmonization is the stage at which measurements from different sources are rendered comparable. In a single-laboratory, single-study context these functions are partly internal to the measurement stage; they become a distinct and demanding pipeline stage the moment data are pooled across laboratories, studies or time periods, which is precisely the situation in which predictive models are most valuable and most at risk. Real-time machine-learning dashboards developed for hospital supply-chain risk offer a transferable example of keeping operational risk signals visible to decision-makers (Filani et al., 2022).

2.2.1 Blanks and contamination control

Because PCBs are ubiquitous environmental contaminants, they are present in the very laboratory that measures them, in solvents, in glassware, in the ambient air. Method blanks, samples carried through the entire procedure without any environmental matrix, quantify this background. Reported concentrations must be interpreted against blank levels, and a common rule requires that a sample signal exceed some multiple of the mean blank before it is considered a true detection. The blank-handling rule is another silent fork: subtracting the mean blank, subtracting nothing, or censoring values below a blank-derived threshold each produce different low-end concentrations, and the low end is where the exposure contrast that drives a health prediction frequently resides.

2.2.2 Surrogates, duplicates and control charts

Beyond blanks, a mature quality-control regime tracks surrogate recoveries over time, analyzes duplicate samples to estimate precision, includes certified reference materials to establish accuracy, and maintains control charts that flag when the analytical system drifts out of statistical control (Muir & Sverko, 2006). These artifacts are the evidence base for asserting that a batch of measurements is fit for use. Critically, they are usually retained by the laboratory and *not* propagated with the data. When a modeller pools data from five studies, the surrogate-recovery distributions, the duplicate precisions and the control-chart histories of those five studies are almost never available, and so the modeller cannot weight or down-weight measurements by their demonstrated reliability. The information exists; the joint does not carry it.

2.2.3 Limits of detection and quantification

The limit of detection (LOD) and limit of quantification (LOQ) define the concentrations below which a signal cannot be reliably distinguished from noise and below which it cannot be reliably quantified, respectively. They are not physical constants; they depend on the instrument, the matrix, the sample mass, the extraction efficiency and the calculation convention. Different laboratories compute them differently, some from blank standard deviations, some from signal-to-noise ratios, some from calibration-curve statistics. When datasets with different LODs are pooled, a value reported as a non-detect in a high-LOD study might have been a confident detection in a low-LOD study, and vice versa. Harmonizing across heterogeneous detection limits is a genuine statistical problem, and treating it as a clerical one, by naively pooling substituted values, injects a source-dependent bias that no downstream model can distinguish from real exposure variation.

2.2.4 Units and lipid normalization

PCB concentrations are reported in several unit systems, and the conversions among them are not innocuous. A concentration can be expressed per unit of whole matrix, for example nanograms per gram of wet serum, or normalized to lipid content, nanograms per gram of lipid, because lipophilic compounds partition into the lipid fraction and lipid normalization improves comparability across individuals with different blood-lipid levels. But lipid normalization requires a lipid measurement, which itself has error and can be computed by different formulae. A dataset that mixes wet-weight and lipid-adjusted values, or that applies different lipid formulae, contains a units heterogeneity

that is invisible in a column of numbers labelled only *concentration*. Unit errors are among the most damaging joint failures precisely because they are arithmetic, deterministic and easy to make: a factor-of-a-thousand slip between nanograms and micrograms, or a failure to lipid-normalize one cohort, propagates undetected into every downstream feature.

2.2.5 Congener panels

Perhaps the most PCB-specific harmonization challenge is the congener panel. There are 209 congeners, but no study measures all of them; each measures a panel chosen for analytical tractability, toxicological relevance or historical continuity (Ballschmiter et al., 1992). One study reports a panel of a few dozen congeners centred on the most abundant and most toxic; another reports a different, partially overlapping panel; a third reports only summed measures or a handful of indicator congeners. To combine these datasets, the analyst must align panels, and alignment is fraught. Summing congeners across studies with different panels compares sums over different sets, an apples-to-oranges aggregation that systematically underestimates total PCBs in the study with the smaller panel. Congener-specific analyses are limited to the intersection of panels, discarding information. Reconstructing a common total from partial panels requires assumptions about the congener distribution that, if wrong, bias the result. The congener panel is thus a harmonization joint of the first importance, and one that is routinely mishandled by silent summation.

2.2.6 The QC-to-integration joint

The joint following quality control is where flags, censoring indicators and reliability metadata should be carried forward into the integrated dataset, and where they are most often dropped. An analysis-ready table that records only concentrations, discarding the detect/non-detect flags, the blank-qualified indicators, the recovery values and the panel membership, has thrown away exactly the information that feature engineering and modeling need to weight the data honestly. The failure is again one of transmission: the quality-control stage may have done excellent work, but if its conclusions do not travel across the joint in machine-actionable form, the downstream stages operate as if the work had never been done.

2.3 Feature engineering and data integration

Feature engineering converts quantified, quality-controlled measurements into the variables a model actually consumes, and data integration assembles those variables, together with outcome and covariate data, into a single analysis-ready matrix. This is the stage where domain knowledge is encoded, and where, correspondingly, the most consequential and most defensible transformations occur, alongside the most opportunities for provenance to evaporate.

2.3.1 Congener ratios and derived indicators

A raw congener concentration is rarely the most informative feature. Toxicologists and exposure scientists routinely derive indicators intended to capture mechanism or source. Ratios of congeners with differing metabolic susceptibility can serve as indicators of metabolic capacity or of time

since exposure; ratios of lower- to higher-chlorinated congeners can indicate the weathering of a source; toxic-equivalency-weighted sums aggregate dioxin-like congeners according to their relative potency; homolog-group sums aggregate congeners by chlorination level. Each derived indicator embeds a hypothesis about biology, and each is only as meaningful as the measurements it is built from. A ratio of two congeners, one of which is frequently a non-detect, inherits the full arbitrariness of the non-detect substitution rule and can swing wildly depending on the value imputed to its denominator. An indicator built from congeners drawn from a poorly harmonized panel inherits the panel mismatch. The virtue of derived indicators, that they encode mechanism, is inseparable from their vulnerability, that they concentrate and amplify upstream errors.

2.3.2 Linking exposure to effect and outcome

The scientific payload of the pipeline is the linkage between exposure features and health data: biomarkers of oxidative stress such as measures of lipid peroxidation, protein carbonyls or oxidized DNA bases; biomarkers of genotoxicity such as comet-assay tail moments, micronucleus frequencies or DNA-adduct levels; and clinical outcomes such as cancer diagnoses. Interpreting such linkages mechanistically is aided by conceptual scaffolds such as the adverse-outcome-pathway framework, which connects a molecular initiating event through intermediate key events to an apical outcome (Ankley et al., 2010). These data typically originate in different systems from the exposure measurements, a clinical laboratory, a pathology registry, a questionnaire, and must be joined by record linkage, matching records to the correct individual, sample and time point. Record linkage is a joint in the most literal sense, and it is error-prone: identifiers may be missing, inconsistent or deliberately obscured for privacy; time alignment between an exposure measurement and an effect measurement may be ambiguous; a single individual may appear under different keys in different systems. A linkage error does not merely add noise; it can systematically attenuate or distort an association by pairing the wrong exposure with the wrong outcome, and it is nearly undetectable downstream because the joined table looks perfectly clean.

2.3.3 The exposome context

The exposome, the totality of environmental exposures across the life course, provides the conceptual backdrop for integration (Rappaport & Smith, 2010). In its full ambition it calls for the joint analysis of chemical exposures alongside complementary measurements: targeted and untargeted chemical assays, and omics layers, transcriptomic, proteomic, metabolomic and epigenomic, that capture the biological response to exposure (Escher, 2017; Peters et al., 2021). The untargeted layers depend on their own dedicated analytical and data-processing workflows, whose peak-alignment, annotation and normalization choices constitute yet another set of joints (Smith et al., 2006; Dunn et al., 2011). Integrating these layers multiplies the harmonization challenge: each omics platform has its own normalization, batch-effect and missing-data conventions, and each joins to the chemical data through its own record-linkage joint. The exposome vision is scientifically compelling precisely because it promises to connect exposure to mechanism to outcome; it is operationally demanding precisely because it multiplies the number of joints, and therefore the number of places where meaning can be lost.

2.3.4 Provenance and versioning

Feature engineering is where provenance most urgently needs to be recorded and most often is not. Every derived feature is the output of a transformation applied to identifiable inputs under a specific rule: this ratio was computed from these two congener columns using this non-detect convention; this summed total aggregated these congeners; this record was linked to this outcome by this key with this confidence. If that lineage is recorded, machine-actionably, the feature can be recomputed, audited and corrected. If it is not, the feature is an orphan: a number whose derivation lives only in the memory of the analyst who produced it, or in an undocumented script that has since been edited. Versioning compounds the problem. Datasets are revised, laboratories re-issue corrected values, panels are extended, linkage keys are updated. Without dataset and code versioning, an analysis cannot even state *which* version of the data it consumed, and two analyses that disagree cannot determine whether they disagree about method or about data.

2.3.5 The integration-to-modeling joint

The joint between integration and modeling is the analysis-ready matrix itself, and its central hazard is that the matrix presents every value with identical apparent authority. A confidently measured, recovery-corrected, panel-aligned congener concentration and a substituted non-detect from a mismatched panel occupy adjacent cells and look identical. The model cannot see the difference unless the difference is encoded as a feature, a censoring indicator, a reliability weight, a provenance flag. In the overwhelmingly common case where the matrix is a bare table of numbers, the joint has flattened a rich, heterogeneous reliability structure into uniform false confidence, and the modeling stage inherits that false confidence without warning.

2.4 Predictive modeling

Only at the fourth stage does the pipeline reach the activity that receives the lion's share of methodological attention. Predictive models, ranging from logistic and Cox regression through penalized regression to random forests, gradient-boosted trees and neural networks, consume the analysis-ready matrix and return predictions of oxidative-stress burden, DNA-damage level or cancer risk. The large-scale computational-toxicology programs that have industrialized this stage, high-throughput in vitro screening and the predictive models built upon it, illustrate both its promise and its dependence on upstream data quality (Dix et al., 2007; Judson et al., 2010; Kavlock et al., 2012; Tice et al., 2013; Richard et al., 2016). Our argument is emphatically not that modeling is unimportant; it is that the modeling stage is bounded from above by the integrity of everything that precedes it, and that this boundedness is routinely ignored.

2.4.1 Bounded by upstream integrity

A predictive model is a function that maps inputs to outputs; it cannot manufacture information that its inputs do not contain, and it cannot detect corruption that its inputs do not disclose. If the exposure gradient in the training data is an artifact of a non-detect substitution rule, the model will learn to predict the artifact. If a batch effect from inconsistent recovery correction correlates with outcome by chance, the model will exploit it as if it were signal. If record-linkage errors have

attenuated the true association, the model will report a weaker relationship and its confidence intervals will not reflect the linkage uncertainty, because that uncertainty was never passed across the joint. Cross-validation, the standard defence against overfitting, protects against learning noise within a dataset; it offers no protection against a systematic upstream bias that is present identically in every fold. A model can be internally validated to any desired rigor and still be predicting an artifact of the pipeline that produced its inputs. This concern is acute precisely for the high-throughput screening data that increasingly train such models, where thousands of assay endpoints are generated under conditions that must themselves be dosimetrically anchored to be interpretable (Rotroff et al., 2010; Wetmore, 2015).

2.4.2 Data leakage across the joints

A particularly insidious failure at this stage is data leakage, in which information from outside the legitimate training context contaminates the model. In environmental-health pipelines, leakage frequently originates upstream. If harmonization parameters, imputation models or normalization constants are estimated using the entire pooled dataset, including samples that will later serve as the test set, then test-set information has leaked into training through the preprocessing joint. The model's apparent performance is inflated, and the inflation is invisible unless the preprocessing is itself folded inside the cross-validation loop. This is a joint failure par excellence: the modeling stage is executed correctly, but the joint that feeds it has quietly violated the separation between training and evaluation.

2.4.3 Distribution shift and transportability

A model trained on one population, measured by one laboratory under one set of conventions, is deployed, implicitly or explicitly, on another. If the target population's exposures were measured by a different laboratory with a different congener panel, different detection limits and different non-detect conventions, then the model is receiving inputs drawn from a different distribution than it was trained on, even when the underlying biology is identical. This distribution shift is generated not by nature but by the pipeline, by inconsistency at the measurement and harmonization joints, and it undermines the transportability that makes predictive models valuable in the first place. A model that cannot be trusted on data processed by a different pipeline is a model of limited practical reach, and the remedy lies upstream, in harmonization, not in the model.

2.4.4 What the modeling stage can and cannot fix

Some upstream problems can be partially addressed at the modeling stage if, and only if, the relevant information has crossed the joints. Censoring can be modelled explicitly rather than imputed, but only if censoring indicators were preserved. Batch effects can be adjusted for, but only if batch identity was recorded. Measurement uncertainty can be propagated into predictive uncertainty, but only if per-measurement reliability was transmitted. In every case, the modeling stage's ability to compensate depends on the joints having done their job. Where the joints have flattened the data into bare numbers, the modeling stage is powerless to reconstruct what was discarded, and the honest response is not a cleverer model but a repaired pipeline.

2.5 Interpretation and decision support

The final stage returns the model's output to human judgement. A risk score, a probability, a classification is interpreted by a scientist, clinician, regulator or affected individual and translated into a decision: to counsel a patient, to prioritize a site for remediation, to set an exposure guideline, to pursue further study. This stage is where the pipeline meets consequence, and its two central requirements are interpretability and calibrated uncertainty. Human-centered automated annotation work reinforces the need to keep expert review visible when generative AI assists interpretation (Ladapo et al., 2024).

2.5.1 Interpretability

A prediction that cannot be explained cannot be responsibly acted upon in a high-stakes setting. Interpretability serves several functions: it allows a domain expert to check whether the model's reasoning is mechanistically plausible, to detect when the model is relying on an artifact, and to communicate the basis of a decision to those affected by it. Mechanistic frameworks such as adverse outcome pathways give this plausibility check a formal vocabulary, linking a statistically important indicator to a hypothesized chain of biological events (Ankley et al., 2010). In environmental health, interpretability also serves an audit function unique to the pipeline framing: an interpretable model that attributes its prediction to a particular congener indicator invites the question of how that indicator was derived, which in turn invites scrutiny of the upstream joints. Interpretability, in this sense, is not merely a property of the model but a window through which the entire pipeline can be examined. An opaque model forecloses that scrutiny and thereby conceals whatever joint failures lie behind its inputs.

2.5.2 Calibrated uncertainty

A prediction offered without an honest statement of its uncertainty is a claim to knowledge the pipeline does not possess. Calibration, the property that predicted probabilities match observed frequencies, is necessary but not sufficient; the uncertainty must also reflect the accumulated imprecision of the entire pipeline, not merely the sampling variability of the final model. In practice, the reported uncertainty of environmental-health predictions typically captures only the last stage's statistical error and omits the propagated contributions of calibration uncertainty, recovery variability, non-detect imputation, harmonization mismatch and linkage error. The result is systematic overconfidence: intervals that are too narrow because they account for one of many error sources. Honest uncertainty quantification requires that each stage's contribution be propagated across the joints, which is impossible when the joints discard the information needed to characterize it.

2.5.3 The interpretation-to-decision joint

The final joint carries the prediction into a decision, and it is where framing, communication and threshold-setting introduce their own distortions. A continuous risk estimate is dichotomized at a threshold whose choice encodes a value judgement about the relative costs of false positives and false negatives; a probability is verbalized as *likely* or *elevated* with attendant ambiguity; an

uncertainty band is dropped in the translation to a headline number. These are joints between the quantitative pipeline and human institutions, and while they lie partly outside the technical scope of this review, they inherit and can amplify every upstream weakness. A decision-maker who receives an overconfident prediction, stripped of its provenance and its true uncertainty, is the final victim of joints that failed far upstream.

3. Data standards and reproducibility

If the joints are the weakest links, then the response must be infrastructural: the pipeline must be engineered so that meaning, provenance and uncertainty travel across its joints by construction rather than by the diligence of individual analysts. This section develops that infrastructural response around four mutually reinforcing ideas: the FAIR principles, the treatment of the pipeline as software, machine-actionable provenance, and containerized workflows.

3.1 FAIR principles as pipeline infrastructure

The FAIR principles hold that data should be Findable, Accessible, Interoperable and Reusable, and that these properties should hold for machines as well as for humans (Wilkinson, 2016; Mons et al., 2017). Originally articulated as guidance for scholarly data stewardship, they map with striking precision onto the failure modes catalogued in Section 2, and they are best understood, in the pipeline framing, not as a data-archiving checklist but as requirements on the joints. Community efforts to catalogue standards, repositories and policies, and to make them themselves findable, provide the connective registry through which such requirements become operational (Sansone, 2019).

Findability requires that every dataset, and ideally every version of every dataset, carries a persistent, globally unique identifier and rich metadata. In pipeline terms, findability is what allows a downstream analysis to state unambiguously which data it consumed, defeating the versioning ambiguity of Section 2.3. A prediction that cites a versioned, identified dataset can be traced back through its inputs; one that cites *the data we received* cannot.

Accessibility requires that data can be retrieved by their identifier through an open, standard protocol, with authentication and authorization where needed. For sensitive health data, accessibility is necessarily mediated by governance, but the principle stands: the pathway to the data must be defined and durable, so that reproduction does not depend on emailing a colleague who may have moved on.

Interoperability is the FAIR principle most directly addressed to the joints. It requires that data use shared, formal vocabularies and representations so that they can be combined without bespoke translation. Almost every harmonization failure in Section 2.2, congener-panel mismatch, unit heterogeneity, inconsistent non-detect conventions, is a failure of interoperability. A controlled

vocabulary for congener identity, a mandatory and explicit unit representation, a standard encoding for censoring status: these are the concrete interoperability commitments that would let two datasets meet at a joint without silent corruption. Established chemical ontologies that assign persistent identifiers and structured relationships to chemical entities offer a concrete substrate for such a vocabulary (Degtyarenko et al., 2008), as do the minimum-reporting-standard initiatives developed for chemical and metabolomic analysis (Sumner et al., 2007).

Reusability requires rich, accurate metadata and clear provenance and licensing, so that data can be legitimately combined and reused in contexts beyond their origin. Reusability is where FAIR meets the transmission failures we have emphasized: the recovery values, the blank-qualification flags, the LOD calculations, the panel membership, all of which the joints tend to drop, are exactly the metadata that reusability demands accompany the data. Minimum-reporting-standard frameworks make this concrete by specifying which experimental and processing details must travel with a measurement for it to be reusable (Sumner et al., 2007). A FAIR-aligned pipeline is one in which this metadata is not an optional supplement but an inseparable part of the data object that crosses each joint.

3.2 The pipeline as software

The second infrastructural commitment is to treat the pipeline as software, and to hold it to the standards that software engineering has developed for managing complexity and ensuring reproducibility (Sandve et al., 2013; Stodden et al., 2018). In much environmental-health practice, the pipeline exists as a loose assemblage: a spreadsheet manipulated by hand, a statistical script edited in place, a sequence of manual steps recorded, if at all, in a methods narrative. This mode of working is irreproducible almost by definition, because the operative artifact, the exact sequence of transformations, is never captured.

Treating the pipeline as software means, first, that every transformation is expressed as code rather than as manual manipulation, so that the pipeline is executable and its behaviour is fully specified. It means, second, that the code is placed under version control, so that the exact state of the pipeline at the time of any analysis can be recovered and so that changes are tracked, attributable and reversible (Sandve et al., 2013). It means, third, that the pipeline is decomposed into modular, testable components, so that the non-detect handler, the unit converter, the panel aligner and the recovery corrector are each independently verifiable, ideally against known inputs and expected outputs. It means, fourth, that the pipeline is accompanied by automated tests, including tests targeted specifically at the joints: a test that fails if a unit column is missing, a test that fails if two datasets with incompatible panels are summed, a test that fails if non-detects are present without accompanying censoring flags. Such tests convert the silent joint failures of Section 2 into loud, early errors.

The discipline of pipeline-as-software also reframes the deliverable of an environmental-health study. Conventionally, the deliverable is a paper reporting a result. In the pipeline framing, the deliverable includes the executable pipeline itself, versioned, tested and documented, as a first-

class scientific artifact. Empirical studies of journal reproducibility policies suggest that requiring such artifacts materially improves the odds that a result can be re-executed (Stodden et al., 2018). A result that is reported without the pipeline that produced it is, from the standpoint of reproducibility, incomplete.

3.3 Provenance

Provenance is the recorded history of how a data artifact came to be: which inputs, which transformations, which parameters, which versions, in which order, by whom or by what. Provenance is the direct antidote to the orphaned features of Section 2.3 and the flattened reliability of Section 2.4. Where provenance is machine-actionable, recorded in a structured, queryable form rather than in prose, any value in the analysis-ready matrix can be interrogated: what is this, where did it come from, how was it derived, how reliable is it (Sandve et al., 2013; Goble et al., 2020). Where provenance is absent, that interrogation is impossible, and the joints have effectively erased the pipeline's own memory.

Machine-actionable provenance has a further virtue: it makes the joints auditable by parties other than the original analyst. A regulator, a peer reviewer or a downstream researcher can, given the provenance record, reconstruct the derivation of a prediction and locate the point at which an assumption was introduced. This transforms reproducibility from an act of good faith into a verifiable property (Goble et al., 2020). It also enables targeted correction: when a laboratory re-issues corrected concentrations, a provenance-aware pipeline can identify exactly which downstream features and predictions depend on the corrected values and recompute only those, rather than forcing a wholesale and error-prone redo.

The recording of provenance should span all five stages and, especially, the joints between them. At the measurement joint, provenance records the calibration model, the non-detect rule and the recovery convention. At the QC joint, it records blank handling, LOD calculations and panel membership. At the integration joint, it records every derived-feature formula and every linkage key and confidence. At the modeling joint, it records the exact training data version, the preprocessing folded into cross-validation and the model's hyperparameters. At the interpretation joint, it records the threshold and the uncertainty components propagated. A pipeline whose provenance is complete in this sense is one whose every joint can answer for itself.

3.4 Containerized and workflow-managed execution

The final infrastructural commitment concerns the execution environment. A pipeline expressed as code still depends on the software environment in which that code runs: the interpreter version, the library versions, the operating-system libraries, the numerical back-ends. Subtle differences in these dependencies can change results, sometimes materially, and they are a notorious source of failure to reproduce an analysis on a different machine or at a later date. Containerization addresses this by packaging the pipeline together with its complete software environment into a portable, versioned image that runs identically wherever it is deployed. The container is, in effect, a guarantee that the *how* of execution travels with the *what* of the code and the *which* of the data.

Workflow-management systems complement containerization by orchestrating the stages of the pipeline explicitly: they represent the pipeline as a directed graph of steps with declared inputs and outputs, execute the steps in the correct order, cache intermediate results, and, crucially, record the execution as it happens, generating provenance automatically (Di Tommaso et al., 2017; Köster, 2012). A workflow-managed, containerized pipeline is one in which reproducibility is largely a byproduct of how the work is done rather than an additional burden imposed after the fact, and recent work has articulated how such workflows can themselves be made FAIR (Goble et al., 2020). The joints become explicit edges in the workflow graph, each with typed inputs and outputs, and a type mismatch at a joint, a dataset without units, a feature without provenance, becomes a detectable, preventable error rather than a silent corruption.

Together, FAIR data, pipeline-as-software, machine-actionable provenance and containerized, workflow-managed execution constitute the infrastructure that makes joints trustworthy. None is novel; each is established practice in adjacent computational fields (Sandve et al., 2013; Di Tommaso et al., 2017; Köster, 2012; Goble et al., 2020). The argument here is that their adoption in environmental health is not a matter of housekeeping but the central determinant of whether predictions of environmentally induced disease can be believed.

4. Toward a transferable pipeline

Reproducibility, the ability to obtain the same result from the same inputs and procedure, is necessary but not sufficient. The higher aspiration is transferability: the ability to move a pipeline, or a component of it, from the context in which it was built to a new context, a new laboratory, a new cohort, a new instrument, and have it continue to function correctly. Transferability is what turns a bespoke analysis into reusable infrastructure, and it is the property that would allow the field to accumulate reliable predictive capability rather than rebuilding it study by study. This section articulates the properties a transferable pipeline must possess and connects them to the modernization of United States environmental-health data infrastructure.

4.1 Properties of a transferable pipeline

4.1.1 Portability

A transferable pipeline runs correctly outside its birthplace. Portability is delivered largely by the containerization and workflow management of Section 3.4, which sever the pipeline's dependence on a particular machine's idiosyncratic environment. But portability also requires that the pipeline make no undocumented assumptions about the data's local conventions, no hidden expectation that concentrations arrive lipid-normalized, no implicit congener panel baked into a summation, no silent reliance on a particular non-detect encoding. A portable pipeline declares its input

requirements explicitly and validates them on receipt, so that a mismatch is caught at the boundary rather than propagating inward.

4.1.2 Modularity

A transferable pipeline is decomposed into components with well-defined interfaces, so that a laboratory adopting it can replace the components specific to its own instruments, its own calibration procedure, its own non-detect convention, while retaining the shared downstream logic. Modularity is what allows the pipeline to accommodate the legitimate heterogeneity of analytical practice without forcing every laboratory into an identical procedure, and without requiring a wholesale rewrite each time a component changes. The joints, in a modular pipeline, become the module interfaces, and they are precisely where the contracts, the declared types, units and metadata, are enforced.

4.1.3 Explicit, enforced interfaces at the joints

The recurring lesson of Section 2 is that joints fail silently. The remedy at the architectural level is to make joints explicit and to give them enforced contracts. An interface contract specifies what a stage requires of its input and guarantees about its output: that concentrations arrive with declared units, that non-detects arrive with censoring flags, that congeners arrive tagged with a controlled-vocabulary identifier, that derived features arrive with provenance. When such contracts are enforced automatically, a violation, the very kind of silent corruption catalogued throughout Section 2, is converted into an immediate, localized failure at the joint where it occurs, rather than a subtle distortion discovered, if ever, in the final prediction. This is the single most important architectural property for trustworthiness, because it targets the weakest links directly.

4.1.4 Provenance and versioning as first-class outputs

A transferable pipeline emits, alongside its predictions, a complete provenance record and explicit version identifiers for its data, its code and its execution environment. These are not optional artifacts produced on request but standard outputs of every run, so that any prediction can be traced, audited and reproduced by a party who was not present at its creation (Goble et al., 2020). Transferability without provenance is dangerous, because a pipeline moved to a new context without a record of its assumptions may run silently on data it was never designed for.

4.1.5 Calibrated, propagated uncertainty

Finally, a transferable pipeline carries uncertainty across its joints so that the final prediction's uncertainty reflects the whole chain. When a component is replaced, say, a laboratory substitutes its own measurement module, the uncertainty characteristics of the new component propagate through to the prediction, so that a prediction made from noisier inputs is honestly reported as more uncertain. A pipeline that propagates uncertainty is self-aware about the quality of its own inputs in a way that a point-estimate pipeline can never be.

4.2 Infrastructure implications for environmental-health data systems

These properties provide testable design requirements for environmental-health data systems. Large chemical-monitoring and biomonitoring programs often combine datasets produced under different conventions, software environments, and quality-control regimes. The resulting interoperability risk should be measured rather than assumed. Audits should quantify unit ambiguity, missing provenance, incompatible identifiers, undocumented transformations, and information loss at each handoff. Integrated visualization models for continuous performance monitoring illustrate one approach for displaying pipeline status, data quality, and downstream outcomes together (Ogbole et al., 2021).

Computational-toxicology blueprints developed within this ecosystem have already articulated a vision in which high-throughput chemical and biological data feed predictive models of hazard, and in which the connective infrastructure, standardized data formats, controlled vocabularies, shared identifiers for chemicals, and reproducible analysis workflows, is treated as essential rather than incidental (Thomas, 2019; Dix et al., 2007; Richard et al., 2016). The screening enterprise these blueprints coordinate, from in vitro assay batteries through the dosimetric extrapolation that makes their readouts interpretable, is itself a multi-stage pipeline whose value depends on the integrity of its joints (Rotroff et al., 2010; Wetmore, 2015; Tice et al., 2013), and shared chemical ontologies are among the concrete standards that make those joints interoperable (Degtyarenko et al., 2008). The pipeline framing of this review is complementary to those blueprints: it emphasizes that the return on investment in high-throughput measurement and in sophisticated modeling is capped by the quality of the joints between them, and that national infrastructure should therefore prioritize interoperability standards, persistent identifiers, provenance capture and reproducible, containerized workflows as the substrate on which measurement and modeling both depend.

A defensible modernization program would test four controls: persistent versioned identifiers for datasets and chemicals; machine-readable encoding of units, detection limits, censoring status, and congener panels; automated provenance capture; and executable workflows supplied with published results. Evaluation should compare these controls with current practice using reproducibility, error-detection rate, analyst time, external calibration, and uncertainty retention. The exposome requires this type of infrastructure because integration across chemical, clinical, and omics layers creates many opportunities for semantic and analytic mismatch (Escher, 2017; Rappaport & Smith, 2010).

4.3 The organizational dimension

It would be incomplete to present transferability as a purely technical achievement. The joints of the pipeline are also organizational: they lie between the chemist and the statistician, between the measuring laboratory and the modeling group, between the data-collecting agency and the analyzing researcher. Technical contracts at the joints must be matched by organizational agreements, shared standards adopted across laboratories, incentives that reward the production of FAIR, provenance-rich data rather than only the publication of results, and training that equips

analysts to treat the pipeline as an engineered artifact. Experience from process-improvement and quality-management practice in resource-constrained organizations underscores that documented, standardized procedure is itself an organizational asset rather than mere overhead (Ambali et al., 2021; Eyetsemitan et al., 2022). The most robust interface contract is undermined if the parties on either side of it do not agree on the vocabulary, and the most elegant containerized workflow is unused if the culture of the field does not value reproducibility as a deliverable. Transferable pipelines are therefore as much a matter of institutions and incentives as of software, and durable progress requires attention to both.

5. Discussion: cross-domain transferability and evidentiary limits

The pipeline requirements identified above should not be accepted solely because they appear reasonable. Their transferability depends on whether comparable measurement-to-decision failures occur in other high-stakes domains and whether the proposed controls address those failures. Diagnostic-laboratory informatics and engineering asset-integrity management offer bounded comparisons because both link calibrated measurements to consequential decisions. The discussion below treats these domains as sources of testable design hypotheses, not as direct evidence that an intervention effective in one field will produce the same result in environmental health.

5.1 Diagnostic-laboratory informatics

Clinical diagnostic laboratories run measurement-to-decision pipelines at enormous scale, and they operate under a regulatory and accreditation regime that has forced explicit attention to the joints. A clinical result, a serum analyte concentration, a genetic variant call, a pathology classification, must travel from instrument to electronic health record to clinician to patient, and errors at any joint can cause direct harm. The way diagnostic infrastructure and workflow are designed materially shapes this reliability, a relationship examined in depth in work on laboratory facility design, diagnostic-network planning and workflow integration (Ogbete et al., 2019; Aminu-Ibrahim & Ogbete, 2023; Asiedu, 2023a). The field's response is instructive on several counts.

First, clinical laboratories rely heavily on standardized vocabularies and identifiers for tests, analytes and results, so that a measurement produced by one laboratory's instrument is interpretable, and its units unambiguous, when it arrives in another institution's record system. This is interoperability enforced at the joint, the very remedy Section 3.1 prescribes for congener-panel and unit heterogeneity. The environmental-health analogue, a universally adopted controlled vocabulary for congener identity and a mandatory unit encoding, lags well behind clinical practice.

Second, clinical laboratories embed quality control as a continuous, documented and externally audited activity, with control charts, proficiency testing across laboratories and mandated corrective action when control is lost. The evidence of fitness-for-use is generated systematically and retained. The environmental-health field generates similar evidence, blanks, surrogates, duplicates, reference materials, but, as Section 2.2 noted, typically fails to propagate that evidence across the joint into the analytical dataset. The clinical model suggests that quality-control metadata should travel with the result as a matter of course. Precise molecular assays such as droplet digital PCR illustrate how tightly interpretive confidence is coupled to documented analytical performance (Asiedu, 2024a).

Third, clinical informatics has confronted record linkage, matching results to the correct patient across systems, as a first-order safety problem, with substantial investment in identifiers, matching algorithms and error detection. The environmental-health integration joint of Section 2.3, where exposure is linked to effect and outcome, would benefit from importing this seriousness about linkage as a source of systematic, not merely random, error. The lesson is not that clinical practice is flawless, it is not, but that treating the joints as safety-critical, and regulating them accordingly, produces markedly more reliable transmission of meaning than treating them as clerical.

5.2 Materials and asset-integrity measurement

A second, less obvious parallel comes from materials science and engineering asset-integrity management, where physical measurements, of corrosion, crack growth, material composition, wall thickness, are fed into predictive models of remaining useful life and failure risk for infrastructure such as pipelines, pressure vessels and structural components (Atta et al., 2020; Atta et al., 2021). The structural resemblance to the environmental-health pipeline is close: a chain runs from a physical measurement, often by an instrument with its own calibration and detection limits, through quality control, feature engineering and predictive modeling, to a high-stakes decision about safety and maintenance.

This domain has developed several practices worth importing. It has a mature culture of measurement uncertainty, in which every measurement is accompanied by a quantified uncertainty and in which that uncertainty is propagated through the predictive model to the final risk estimate, so that a remaining-life prediction carries an honest confidence band. This is exactly the propagated, calibrated uncertainty that Section 2.5 found so often missing in environmental-health predictions. The asset-integrity field also emphasizes traceability, the ability to trace a measurement back to a calibrated reference standard, which is a form of provenance, and it standardizes inspection and quality-assurance practice so that measurements made by different inspectors remain comparable (Atta et al., 2022; Atta et al., 2018). The digital-twin idea, a persistent, versioned, machine-actionable representation that integrates heterogeneous measurements over time, is a suggestive template for the integrated, provenance-aware exposome dataset that the environmental-health pipeline aspires to build, and analogous integrity concerns arise even in the characterization of advanced materials and devices whose performance degrades under environmental stress (Atta et al., 2024).

Equally instructive are the domain's failure analyses. Investigations of infrastructure failures repeatedly trace root causes not to the physics of the material but to breakdowns in the measurement-to-decision chain: a mis-calibrated instrument, an inspection whose results were mis-transcribed, an uncertainty that was ignored, a data joint at which a critical caveat was lost. These are joint failures in the sense of this review, and their consequences, in a domain where failure is catastrophic and visible, have driven a discipline of joint-level rigor that environmental health, where the consequences are diffuse and slow, has been slower to adopt.

5.3 Transferable hypotheses and limits of analogy

The comparison yields five hypotheses for environmental-health pipelines: shared vocabularies should reduce semantic mismatch; quality-control metadata should retain information needed for downstream interpretation; record-linkage checks should reduce systematic attribution error; versioned provenance should improve reproducibility; and joint-level audits should detect failures hidden by acceptable endpoint performance. These propositions remain conditional until tested with environmental samples, laboratories, and decisions. Cross-domain consistency strengthens their plausibility, but it does not establish effect size, implementation feasibility, or superiority over simpler controls.

5.4 Limitations and research priorities

Three limitations define the interpretation of this review. First, the PCB example provides analytical depth but does not represent every contaminant class, matrix, or instrument. Second, the integrative design supports mechanism-based synthesis but does not establish exhaustive coverage or comparative effect sizes. Third, lessons imported from clinical laboratories and engineering asset management are analogies whose value depends on empirical testing in environmental-health workflows. The review therefore supports design requirements and research hypotheses, not universal estimates of pipeline failure frequency.

The immediate research need is a set of public pipeline challenge studies. Multiple laboratories should process shared reference materials and harmonized raw files through independently implemented workflows while preserving units, censoring flags, calibration records, feature definitions, software versions, and uncertainty. Performance should be judged at every joint, not only by agreement in the final prediction. Such studies would reveal whether provenance coverage, interface tests, and containerized execution improve reproducibility under realistic variation.

A second priority is decision-centered benchmarking. Environmental-health pipelines should report calibration, transportability, uncertainty retention, and downstream decision consequences alongside conventional discrimination metrics. Prespecified failure thresholds should govern whether a pipeline is accepted, restricted, recalibrated, or retired. This change would convert reproducibility from a retrospective aspiration into an operational release criterion.

6. Conclusion

Predictive toxicology is limited by the integrity of the full measurement-to-decision pipeline. The five-stage analysis shows that many consequential errors originate at stage boundaries: calibration and censoring choices lose context, quality-control evidence disappears, features lose provenance, datasets are linked incorrectly, and uncertainty is stripped from predictions before interpretation. A high-performing model cannot recover information that the upstream pipeline failed to preserve.

The proposed remedy is operational. Data and metadata should move together through explicit interface contracts. Workflows should be versioned, tested, and executable in controlled environments. Provenance should identify the origin and transformation history of every decision-relevant value. Uncertainty should be propagated from measurement through interpretation. These requirements support reproducibility within a laboratory and transportability across laboratories, instruments, cohorts, and institutions.

The review is integrative rather than exhaustive, and its cross-domain lessons require empirical testing in environmental-health settings. Multi-laboratory challenge studies should therefore evaluate each joint using shared materials, raw files, and prespecified failure thresholds. Evaluation should include provenance coverage, unit and censoring integrity, feature reproducibility, external calibration, uncertainty retention, and downstream decision consequences. Treating these measures as release criteria would make the pipeline, rather than the final algorithm alone, accountable for the trustworthiness of environmental-health predictions.

References

Aliliele, C., Mbonu, I.S. & Iwuanyanwu, U. (2023a). Advances in Predictive Analytics Models for Student Retention and Institutional Risk Management Systems. International Journal of Advanced Multidisciplinary Research and Studies, 3(6), pp.2692–2711. https://doi.org/10.62225/2583049X.2023.3.6.5990.

Aliliele, C., Mbonu, I.S. & Iwuanyanwu, U. (2023b). A Conceptual Framework for Continuous Cloud Misconfiguration Monitoring and Enterprise Risk Mitigation Strategies. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 9(10), pp.373–394. https://doi.org/10.32628/CSEIT2361071.

Aliliele, C., Mbonu, I.S. & Iwuanyanwu, U. (2023c). A Review of API Governance and Risk Prioritization Frameworks in Modern Financial Institutions. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 9(10), pp.395–433. https://doi.org/10.32628/CSEIT2361072.

Aliliele, C., Mbonu, I.S. & Iwuanyanwu, U. (2024a). A Conceptual Framework for Enterprise Data Sensitivity Classification and Regulatory Traceability Mechanisms. International Journal of Advanced Multidisciplinary Research and Studies, 4(6), pp.3103–3124. https://doi.org/10.62225/2583049X.2024.4.6.5991.

Aliliele, C., Mbonu, I.S. & Iwuanyanwu, U. (2024b). Advances in HIPAA Compliant Data Architecture and Secure Analytics Frameworks for Community Healthcare Organizations. Shodhshauryam, International Scientific Refereed Research Journal, 7(2), pp.277–324. https://doi.org/10.32628/SHISRRJ2472163.

Ambali, K.B., Eyetsemitan, R.A., Oyeleye, A.O. & Fadayomi, O. (2021). Lean Six Sigma for Small Enterprises: A Systematic Review and Lite-DMAIC Adaptation Framework for Resource-Constrained Organizations. Iconic Research and Engineering Journals, 5(5), pp.562–583. https://doi.org/10.64388/IREV5I5-1716957.

Aminu-Ibrahim, A. Y., & Ogbete, J. C. (2023). Healthcare infrastructure as a public health intervention using evidence from large laboratory networks. Shodhshauryam, International Scientific Refereed Research Journal, 6(1), 256-286. https://doi.org/10.32628/SHISRRJ23678

Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2018). Developing sustainable diagnostic laboratory infrastructure models for emerging and resource constrained health systems. Iconic Research and Engineering Journals, 1(8), 118-132. https://doi.org/10.64388/IREV1I8-1713586

Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2019). Capital project delivery models for high risk healthcare infrastructure in developing national health systems. Iconic Research and Engineering Journals, 2(10), 626-649. https://doi.org/10.64388/IREV2I10-1713588

Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2020). Infrastructure driven expansion of diagnostic access across underserved and rural healthcare regions. International Journal of Multidisciplinary Research and Growth Evaluation, 1(5), 691-706. https://doi.org/10.54660/IJMRGE.2020.1.5.691-706

Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2024). Governance and accountability models for public private partnerships in healthcare infrastructure development. International Journal of Advanced Multidisciplinary Research and Studies, 4(6), 2943-2960. https://doi.org/10.62225/2583049X.2024.4.6.5699

Ankley, G. T., Bennett, R. S., Erickson, R. J., Hoff, D. J., Hornung, M. W., Johnson, R. D., Mount, D. R., Nichols, J. W., Russom, C. L., Schmieder, P. K., Serrrano, J. A., Tietge, J. E., & Villeneuve, D. L. (2010). Adverse outcome pathways: A conceptual framework to support ecotoxicology research and risk assessment. Environmental Toxicology and Chemistry, 29(3), 730-741. https://doi.org/10.1002/etc.34

Asiedu, C. S. (2023a). Diagnostic integration in infectious disease management: A review of microbiology, virology, and serology workflows. International Journal of Medical Evaluation and Physical Report, 7(4), 173-195. https://doi.org/10.56201/ijmepr.v7.no4.2023.pg173.195

Asiedu, C. S. (2024a). Droplet digital PCR for HIV viral reservoir quantification: A review of diagnostic accuracy and clinical utility. International Journal of Medical Evaluation and Physical Report, 8(6), 253-273. https://doi.org/10.56201/ijmepr.v8.no6.2024.pg253.273

Asiedu, W., & Quainoo, R. (2023b). How far can energy harvesting take us? A systematic review of radio frequency strategies for energy autonomous sensing. Shodhshauryam, International Scientific Refereed Research Journal, 6(1), 448-466.

Asiedu, W., & Quainoo, R. (2023c). Toward maintenance free wireless infrastructure: Simulating performance and reliability in large scale intermittently powered IoT networks. Gyanshauryam, International Scientific Refereed Research Journal, 6(1), 489-510.

Asiedu, W., & Quainoo, R. (2024b). Rethinking energy, reliability, and latency trade-offs in green communication for next generation IoT. International Journal of Multidisciplinary Research and Growth Evaluation, 5(6), 1987-1994.

Atta, S., Adjaklo, E., Asomani, E., & Adenuga, O. M. (2022). Advances in welding inspection standards and quality assurance practices in large-scale power generation projects. International Journal of Engineering and Modern Technology, 8(5), 101-124. https://doi.org/10.56201/ijemt.v8.no5.2022.pg101.124

Atta, S., Asomani, E., & Adenuga, O. M. (2018). A systematic review of green corrosion inhibitors from agricultural extracts for mild steel protection in acidic and alkaline media. International Journal of Engineering and Modern Technology, 4(1), 64-97. https://doi.org/10.56201/ijemt.vol.4.no1.2018.pg64.97

Atta, S., Onyishi, E. E., Appiah, S. A., & Adenuga, O. M. (2024). Theoretical analysis of charge transport and recombination losses in thin film solar cell architectures. International Journal of Engineering and Modern Technology, 10(10), 147-170. https://doi.org/10.56201/ijemt.v10.no10.2024.pg147.170

Atta, S., Tofah, P., & Adenuga, O. M. (2021). Advances in non-destructive testing methods for weld integrity evaluation in high-capacity power infrastructure. International Journal of Engineering and Modern Technology, 7(1), 20-47. https://doi.org/10.56201/ijemt.vol.7.no1.2021.pg20.47

Atta, S., Tofah, P., Kankam, M. A., & Adenuga, O. M. (2020). A critical review of NDT-based integrity management and corrosion risk assessment strategies for refinery process equipment. International Journal of Engineering and Modern Technology, 6(2), 19-46. https://doi.org/10.56201/ijemt.vol.6.no2.2020.pg19.46

Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. Iconic Research and Engineering Journals, 2(6).

Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2019a). A conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures. Iconic Research and Engineering Journals, 3(5).

Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2019b). A governance framework for Salesforce platform management in regulated healthcare environments. Iconic Research and Engineering Journals, 3(6).

Badmus, O., Jooda, D., Ozowara, D. E., & Anunagba, C. O. (2022a). A framework for Git-based version control and code review governance in Salesforce development teams. Shodhshauryam, International Scientific Refereed Research Journal, 5(2), 473-493.

Badmus, O., Jooda, D., Ozowara, D. E., & Anunagba, C. O. (2022b). A systematic review of MuleSoft ESB integration patterns in multi-cloud Salesforce architectures. Gyanshauryam, International Scientific Refereed Research Journal, 5(3), 462-481.

Ballschmiter, K., Bacher, R., Mennel, A., Fischer, R., Riehle, U. & Swerev, M. (1992). The determination of chlorinated biphenyls, chlorinated dibenzodioxins and chlorinated dibenzofurans by GC–MS. Journal of High Resolution Chromatography, 15, 260–270.

Borah, S., Aliliele, K.C., Rakshit, S. & Vajjhala, N.R. (2022). Applications of Artificial Intelligence in Software Testing. In Cognitive Informatics and Soft Computing, Lecture Notes in Networks and Systems, Vol. 375, pp.727–736. Springer. https://doi.org/10.1007/978-981-16-8763-1_60.

Degtyarenko, K., de Matos, P., Ennis, M., Hastings, J., Zbinden, M., McNaught, A., Alcántara, R., Darsow, M., Guedj, M., & Ashburner, M. (2008). ChEBI: A database and ontology for chemical entities of biological interest. Nucleic Acids Research, 36(Database issue), D344-D350. https://doi.org/10.1093/nar/gkm791

Di Tommaso, P., Chatzou, M., Floden, E. W., Barja, P. P., Palumbo, E., & Notredame, C. (2017). Nextflow enables reproducible computational workflows. Nature Biotechnology, 35(4), 316-319. https://doi.org/10.1038/nbt.3820

Dix, D. J., Houck, K. A., Martin, M. T., Richard, A. M., Setzer, R. W. & Kavlock, R. J. (2007). The ToxCast program for prioritizing toxicity testing of environmental chemicals. Toxicological Sciences, 95(1), 5–12.

Dunn, W. B., Broadhurst, D., Begley, P., Zelena, E., Francis-McIntyre, S., Anderson, N., Brown, M., Knowles, J. D., Halsall, A., Haselden, J. N., Nicholls, A. W., Wilson, I. D., Kell, D. B., & Goodacre, R. (2011). Procedures for large-scale metabolic profiling of serum and plasma using gas chromatography and liquid chromatography coupled to mass spectrometry. Nature Protocols, 6(7), 1060-1083. https://doi.org/10.1038/nprot.2011.335

Eboh, E.E. & Aliliele, C. (2023). Portfolio Risk Optimization Models for Venture Capital Investment Strategy and Fraud Exposure Analysis. Iconic Research and Engineering Journals, 7(4), pp.759–792. https://doi.org/10.64388/IREV7I4-1715305.

Eboh, E.E. & Aliliele, C. (2024). AI-Driven Data Analytics Framework for Risk Assessment and Detection of Venture Capital and Private Equity Investment Fraud in U.S. Capital Markets. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(6), pp.2624–2665. https://doi.org/10.32628/CSEIT2410787.

Escher, B. I., Hackermüller, J., et al. (2017). From the exposome to mechanistic understanding of chemical-induced adverse effects. Environment International, 99, 97–106.

Eyetsemitan, R.A., Ambali, K.B., Oyeleye, A.O. & Fadayomi, O. (2020). Multi-Stakeholder Governance Alignment in Joint Venture Operations: A Conceptual Framework for Coordinating Business Processes in Highly Regulated Environments. Iconic Research and Engineering Journals, 4(4), pp.418–441. https://doi.org/10.64388/IREV4I4-1716955.

Eyetsemitan, R.A., Ambali, K.B., Oyeleye, A.O. & Fadayomi, O. (2021). Translating Tax and Regulatory Requirements into SME Compliance Workflows: A Conceptual Framework for Implementing IAS 12, VAT, PAYE, and Withholding Tax. Iconic Research and Engineering Journals, 4(11), pp.621–641. https://doi.org/10.64388/IREV4I11-1716956.

Eyetsemitan, R.A., Ambali, K.B., Oyeleye, A.O. & Fadayomi, O. (2023a). User Acceptance Testing in Small Business Technology Deployment: A Structured Validation Framework for Lean Operational Environments. International Journal of Multidisciplinary Research and Growth Evaluation, 4(6), pp.1512–1531. https://doi.org/10.54660/IJMRGE.2023.4.6.1512-1531.

Eyetsemitan, R.A., Ambali, K.B., Oyeleye, A.O. & Fadayomi, O. (2023b). Change Management in Small Business Digital Transformation: A Systematic Review and Lean Change Adoption Framework. Gyanshauryam, International Scientific Refereed Research Journal, 6(6), pp.521–550. https://doi.org/10.32628/GISRRJ236648.

Eyetsemitan, R.A., Oyeleye, A.O., Ambali, K.B. & Fadayomi, O. (2022). Standard Operating Procedures as Strategic Assets in Small Business Operations: A Systematic Review and Implementation Framework. Gyanshauryam, International Scientific Refereed Research Journal, 5(2), pp.438–465. https://doi.org/10.32628/GISRRJ225356.

Eyetsemitan, R.A., Oyeleye, A.O., Ambali, K.B. & Fadayomi, O. (2024a). CRM and Workflow Automation in Small Healthcare Practices: A Process Efficiency Framework for Scalable Patient Engagement. International Journal of Multidisciplinary Research and Growth Evaluation, 5(6), pp.1931–1949. https://doi.org/10.54660/IJMRGE.2024.5.6.1931-1949.

Eyetsemitan, R.A., Oyeleye, A.O., Ambali, K.B. & Fadayomi, O. (2024b). Data-Driven Process Optimization in Micro-Enterprises: A Conceptual Framework for Funnel Analysis and Bottleneck

Identification. International Journal of Multidisciplinary Research and Growth Evaluation, 5(6), pp.1950–1968. https://doi.org/10.54660/IJMRGE.2024.5.6.1950-1968.
Filani, O. M., Nnabueze, S. B., Ike, P. N., & Wedraogo, L. (2022). Real-time risk assessment dashboards using machine learning in hospital supply chain management systems. International Journal of Multidisciplinary Evolutionary Research, 3(1), 65-76. https://doi.org/10.54660/IJMER.2022.3.1.65-76
Goble, C., Cohen-Boulakia, S., Soiland-Reyes, S., Garijo, D., Gil, Y., Crusoe, M. R., Peters, K., & Schober, D. (2020). FAIR computational workflows. Data Intelligence, 2(1-2), 108-121. https://doi.org/10.1162/dint_a_00033
Judson, R. S., Houck, K. A., Kavlock, R. J., Knudsen, T. B., Martin, M. T., Mortensen, H. M., et al. (2010). In vitro screening of environmental chemicals for targeted testing prioritization: The ToxCast project. Environmental Health Perspectives, 118(4), 485–492.
Kavlock, R., Chandler, K., Houck, K., Hunter, S., Judson, R., Kleinstreuer, N., et al. (2012). Update on EPA's ToxCast program: Providing high-throughput decision support tools for chemical risk management. Chemical Research in Toxicology, 25(7), 1287–1302.
Kumuyi, O., Uzoka, E., Akeju, B., & Ozowara, D. E. (2024). Architecture for machine learning-enabled predictive energy management using IoT sensor networks. International Journal of Advanced Multidisciplinary Research and Studies, 4(6), 3138-3149. https://doi.org/10.62225/2583049X.2024.4.6.6004
Köster, J., & Rahmann, S. (2012). Snakemake: a scalable bioinformatics workflow engine. Bioinformatics, 28(19), 2520-2522. https://doi.org/10.1093/bioinformatics/bts480
Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2024). Keeping humans in the loop: Human-centered automated annotation with generative AI. International Journal of Multidisciplinary Futuristic Development, 5(1), 81-95.
Mbonu, I.S., Aliliele, C., Iwuanyanwu, U. & Oluoha, O.M. (2018). A Conceptual Framework for Legal and Ethical Risk Modeling in Enterprise Data Protection Governance Systems. Iconic Research and Engineering Journals, 2(2), pp.207–226. https://doi.org/10.64388/IREV2I2-1714911.
Mbonu, I.S., Aliliele, C., Iwuanyanwu, U. & Uzoka, E. (2020b). A Review of Identity and Access Management Integration Strategies in Hybrid and Multi Cloud Environments. International Journal of Multidisciplinary Research and Growth Evaluation, 1(5), pp.795–810. https://doi.org/10.54660/IJMRGE.2020.1.5.795-810.
Mbonu, I.S., Aliliele, C., Iwuanyanwu, U. & Uzoka, E. (2021b). A Conceptual Framework for Risk Based Business Intelligence Architecture in Financial Technology Platforms. International Journal of Multidisciplinary Research and Growth Evaluation, 2(6), pp.731–746. https://doi.org/10.54660/IJMRGE.2021.2.6.731-746.
Mbonu, I.S., Aliliele, C., Iwuanyanwu, U. & Uzoka, E. (2021c). Advances in Artificial Intelligence Techniques for Secure Software Testing and Automated Regression Control Mechanisms. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 7(5), pp.468–496. https://doi.org/10.32628/CSEIT217565.
Mbonu, I.S., Aliliele, C., Iwuanyanwu, U. & Uzoka, E. (2022b). A Conceptual Framework for AI Enabled IT General Controls and SOX Audit Automation Processes. Gyanshauryam, International Scientific Refereed Research Journal, 5(5), pp.384–414. https://doi.org/10.32628/GISRRJ2256239.
Mbonu, I.S., Aliliele, C., Uzoka, E. & Oluoha, O.M. (2019a). A Review of Comparative Data Protection Regulations and Secure Cloud Implementation Strategies Across Jurisdictions. Iconic

Research and Engineering Journals, 2(9), pp.482–501. https://doi.org/10.64388/IREV2I9-1714912.
Mbonu, I.S., Iwuanyanwu, U., Aliliele, C. & Uzoka, E. (2020a). A Conceptual Framework for Agile Supply Chain Digital Transformation with Embedded IT Risk and ISO Compliance Controls. Iconic Research and Engineering Journals, 3(11), pp.566–593. https://doi.org/10.64388/IREV3I11-1714916.
Mbonu, I.S., Iwuanyanwu, U., Aliliele, C. & Uzoka, E. (2020c). Advances in Infrastructure as Code Governance for Secure Terraform Based Enterprise Cloud Deployments. International Journal of Multidisciplinary Research and Growth Evaluation, 1(5), pp.811–828. https://doi.org/10.54660/IJMRGE.2020.1.5.811-828.
Mbonu, I.S., Iwuanyanwu, U., Aliliele, C. & Uzoka, E. (2021a). A Review of VoIP Forensic Analytics Models for Financial Fraud Detection and Regulatory Compliance Monitoring. International Journal of Multidisciplinary Research and Growth Evaluation, 2(6), pp.711–730. https://doi.org/10.54660/IJMRGE.2021.2.6.711-730.
Mbonu, I.S., Iwuanyanwu, U., Aliliele, C. & Uzoka, E. (2022a). A Review of Data Protection Impact Assessment Models in Multi Cloud Financial Infrastructure Systems. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 8(1), pp.589–623. https://doi.org/10.32628/CSEIT25442.
Mbonu, I.S., Iwuanyanwu, U., Aliliele, C. & Uzoka, E. (2022c). Advances in Cloud Identity and Access Governance Optimization in Large Scale AWS Enterprise Environments. Shodhshauryam, International Scientific Refereed Research Journal, 5(3), pp.403–438. https://doi.org/10.32628/SHISRRJ225490.
Mbonu, I.S., Iwuanyanwu, U., Uzoka, E. & Oluoha, O.M. (2019b). Advances in Enterprise Log Analytics and Automated Incident Response Architectures Using Python and SIEM Platforms. Iconic Research and Engineering Journals, 3(2), pp.1000–1019. https://doi.org/10.64388/IREV3I2-1714915.
Mons, B., Neylon, C., Velterop, J., Dumontier, M., da Silva Santos, L. O. B., & Wilkinson, M. D. (2017). Cloudy, increasingly FAIR; revisiting the FAIR Data guiding principles for the European Open Science Cloud. Information Services & Use, 37(1), 49-56. https://doi.org/10.3233/ISU-170824
Muir, D. & Sverko, E. (2006). Analytical methods for PCBs and organochlorine pesticides in environmental monitoring and surveillance. Analytical and Bioanalytical Chemistry, 386, 769–789.
Ogbete, J. C., & Aminu-Ibrahim, A. Y. (2024). Translating healthcare infrastructure investment into measurable population health and diagnostic outcomes. International Journal of Scientific Research in Humanities and Social Sciences, 1(2), 955-985. https://doi.org/10.32628/IJSRSSH242775
Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2018). Optimizing laboratory spatial planning strategies to improve diagnostic accuracy, safety, and clinical throughput. Iconic Research and Engineering Journals, 2(1), 87-113. https://doi.org/10.64388/IREV1I7-1713587
Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2019). Regulatory compliant design systems for molecular and pathology laboratories in highly controlled environments. Iconic Research and Engineering Journals, 3(4), 607-631. https://doi.org/10.64388/IREV3I4-1713589
Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2020). Sustainable materials selection and energy efficiency strategies for modern medical laboratory facilities. International Journal of

Multidisciplinary Research and Growth Evaluation, 1(5), 674-690. https://doi.org/10.54660/IJMRGE.2020.1.5.674-690

Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2021). Risk managed construction strategies for complex and highly regulated healthcare facility developments. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 7(5), 362-392. https://doi.org/10.32628/CSEIT217561

Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2022). Design standards and operational planning frameworks for scalable blood collection networks. Gyanshauryam, International Scientific Refereed Research Journal, 5(6), 267-300. https://doi.org/10.32628/GISRRJ225635

Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2023). Lifecycle performance evaluation of purpose built diagnostic laboratories supporting long term healthcare delivery. International Journal of Advanced Multidisciplinary Research and Studies, 3(6), 2605-2621. https://doi.org/10.62225/2583049X.2023.3.6.5698

Ogbole, J. I., Okoruwa, P. O., Babatope, O. M., & Oyewole, T. (2021). Developing an integrated data visualization model for continuous business performance monitoring and optimization. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 7(5), 441-467. https://doi.org/10.32628/IJSRCSEIT

Oyeleye, A.O., Eyetsemitan, R.A., Ambali, K.B. & Fadayomi, O. (2022). Reducing Client Onboarding Cycle Time in Small Professional Services Firms: A Lean Six Sigma Process Redesign Framework. Gyanshauryam, International Scientific Refereed Research Journal, 5(2), pp.467–498. https://doi.org/10.32628/GISRRJ225357.

Peters, A., Nawrot, T. S. & Baccarelli, A. A. (2021). Hallmarks of environmental insults. Cell, 184(6), 1455–1468.

Quainoo, R., Ogundapo, O., & Asiedu, W. A. (2024). Modeling the impact of impedance mismatch on wireless link performance: A framework for prototype development and system optimization. International Journal of Computer Science and Mathematical Theory, 10(2), 62-116. https://doi.org/10.56201/ijcsmt.v10.no2.2024.pg62.116

Rappaport, S. M., & Smith, M. T. (2010). Environment and disease risks. Science, 330(6003), 460-461. https://doi.org/10.1126/science.1192603

Richard, A. M., Judson, R. S., Houck, K. A., Grulke, C. M., Volarath, P., Thillainadarajah, I., Yang, C., Rathman, J., Martin, M. T., Wambaugh, J. F., Knudsen, T. B., Kancherla, J., Mansouri, K., Patlewicz, G., Williams, A. J., Little, S. B., Crofton, K. M., & Thomas, R. S. (2016). ToxCast chemical landscape: Paving the road to 21st century toxicology. Chemical Research in Toxicology, 29(8), 1225-1251. https://doi.org/10.1021/acs.chemrestox.6b00135

Rotroff, D. M., Wetmore, B. A., Dix, D. J., Ferguson, S. S., Clewell, H. J., Houck, K. A., LeCluyse, E. L., Andersen, M. E., Judson, R. S., Smith, C. M., Sochaski, M. A., Kavlock, R. J., Boellmann, F., Martin, M. T., Reif, D. M., Wambaugh, J. F., & Thomas, R. S. (2010). Incorporating human dosimetry and exposure into high-throughput in vitro toxicity screening. Toxicological Sciences, 117(2), 348-358. https://doi.org/10.1093/toxsci/kfq220

Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. (2013). Ten simple rules for reproducible computational research. PLOS Computational Biology, 9(10), e1003285. https://doi.org/10.1371/journal.pcbi.1003285

Sansone, S.-A., et al. (2019). FAIRsharing as a community approach to standards, repositories and policies. Nature Biotechnology, 37, 358–367.

Smith, C. A., Want, E. J., O'Maille, G., Abagyan, R., & Siuzdak, G. (2006). XCMS: Processing mass spectrometry data for metabolite profiling using nonlinear peak alignment, matching, and identification. Analytical Chemistry, 78(3), 779-787. https://doi.org/10.1021/ac051437y

Stodden, V., Seiler, J. & Ma, Z. (2018). An empirical analysis of journal policy effectiveness for computational reproducibility. Proceedings of the National Academy of Sciences, 115(11), 2584–2589.

Sumner, L. W., Amberg, A., Barrett, D., Beale, M. H., Beger, R., Daykin, C. A., Fan, T. W.-M., Fiehn, O., Goodacre, R., Griffin, J. L., Hankemeier, T., Hardy, N., Harnly, J., Higashi, R., Kopka, J., Lane, A. N., Lindon, J. C., Marriott, P., Nicholls, A. W., … Viant, M. R. (2007). Proposed minimum reporting standards for chemical analysis: Chemical Analysis Working Group (CAWG) Metabolomics Standards Initiative (MSI). Metabolomics, 3(3), 211-221. https://doi.org/10.1007/s11306-007-0082-2

Thomas, R. S., et al. (2019). The next generation blueprint of computational toxicology at the U.S. EPA. Toxicological Sciences, 169(2), 317–332.

Tice, R. R., Austin, C. P., Kavlock, R. J., & Bucher, J. R. (2013). Improving the human hazard characterization of chemicals: A Tox21 update. Environmental Health Perspectives, 121(7), 756-765. https://doi.org/10.1289/ehp.1205784

Wetmore, B. A. (2015). Quantitative in vitro-to-in vivo extrapolation in a high-throughput environment. Toxicology, 332, 94–101.

Wilkinson, M. D., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3, 160018.