

A Conceptual Framework for Structural Interoperability in Environmental IoT: Towards a Heterogeneous Interaction Paradigm for Industrial Olfactory Monitoring

Jacques Mudumbi, Nadine Nibigira, Ebedon Mufind, Julius Niyongabo, Et Jérémie Ndikumagenge

Doctoral School Science, Center For Research In Infrastructure, Environment And Technologies, University Of Burundi

Abstract - This paper critically reviews research connecting environmental IoT, industrial odor control, and data interoperability, then introduces a clear conceptual framework and a practical benchmark dataset modeled after a working brewery in Rwanda, where the scent of malt hangs in the air. The dataset blends four core streams Source for process states and events, Emission for gas and odor sensors, Vector for meteorology, and Impact for crowdsourced complaints each kept at its own natural rhythm, whether measured every second, minute, hour, or triggered by an event. We present a four-layer architecture Physical & Capture, Edge & Preprocessing, Mediation & Orchestration, and Analytics & Decision that handles observations as typed entities, each clearly encoding when and where it was recorded, how certain it is, and where it came from, like mapping a sensor's brief flash of light to its precise location and time. Each data transformation acts as a contract-bound morphism that spells out distortion, latency, energy cost, and validity terms so engineers can optimize processing paths with precision instead of patching together ad hoc conversions like quick tape fixes. A lean metadata profile covering time, space, units, temporal range, uncertainty, and provenance along with simple semantic links, drives both structural and semantic interoperability, enables multi-resolution data fusion, and powers provenance-aware analytics; imagine data layers aligning like transparent maps under a clear light. With the brewery testbed, we show how the composite dataset lets us realistically test algorithms and architectures as they face calibration drift, timing slips, network lag, and tight energy limits the way steam hisses from a valve just a beat too late. We wrap up with a few practical tips to make E-IoT odor-monitoring setups in tight-budget environments more transferable, reproducible, and accountable so even a small sensor on a dusty window ledge can deliver reliable data.

Keywords : *Environmental IoT ; industrial odor monitoring ; data interoperability ; multimodal benchmark dataset ; provenance-aware analytics.*

I. INTRODUCTION

Industrial odor pollution has become a worldwide problem, connecting the sharp, lingering smells people notice in their neighborhoods with clear health, social, and economic harm. Across Europe, systematic reviews of industrial odour exposure published since 2000 have identified at least 30 epidemiological studies, showing that people living near these sources face 10–30% higher risks of headaches and breathing problems than those in unaffected areas, where the air smells clean and sharp after rain [1]. Across North America, studies of livestock odors found that when farms sat too close to homes, nearby property values dropped more than farmers gained families lost equity faster than producers earned, especially under the setback rules common in the early 2000s [2] [3]. Cities across Asia and Africa, from dense downtowns to fast-growing outskirts, reveal the same strain between breakneck industrial growth and weak oversight; residents often report sharp, sour smells near large animal farms and waste plants, prompting new experiments in local air-quality governance and citizen involvement [4] [5] [6].

As EIoT has grown at breakneck speed, it's created a patchwork of platforms each tuned to a particular use, technology, or setting, like tools laid out for different jobs on a crowded workbench. Different vendors' middleware, custom data schemas, and mismatched metadata practices splinter observation and analysis into isolated micro-ecosystems even when every sensor is watching the same stretch of sky [7] [8]. In the field of scent, fragmentation grows worse because there are so many different ways to sense an aroma like how one device measures molecules while another tracks faint air currents. Electronic noses fire off complex signals in milliseconds, while gas analysers, air-quality monitors, and process instruments usually report minute- or hour-averaged readings; citizen reports and incident logs trickle in as scattered, event-based notes ; meanwhile, meteorological stations follow their own rhythm of sampling and aggregation [9] [10].

This mix isn't just about communication protocols or file formats those are mostly handled through standard interfaces but it reaches deeper, shaping the very structure and meaning, like the way two systems might label the same data in completely different ways [11]. Observation types, units,

calibration states, quality flags, and data origins are recorded inconsistently, while activities, emission sources, and receptors are sorted with local vocabularies that rarely come with documentation like labels scribbled in a notebook no one else can read. Without a clear interoperability layer spelling out shared data schemas, timing rules, and reference ontologies, fusing results across platforms turns messy hard to repeat, tough to trust, like trying to match mismatched puzzle pieces [12] [13].

This silo effect creates ripples that are technical, institutional, and even about how we understand things like engineers speaking one language while policy teams speak another. In practice, when multi-rate, out-of-sync data streams can't be linked, it throws off any attempt to trace a sharp odor back to a single source or equipment glitch, and it makes comparing control methods from one site to another nearly impossible. Earlier studies on environmental and smart-grid monitoring show that diverse sensor readings can be brought together through formal interaction models and layered aggregation but only when shared through standardized, machine-readable interfaces [14] [7]. Within institutions, patchy data systems deepen the imbalance between operators, regulators, and communities each peering at the same issue through a fogged window, their perspectives incomplete and impossible to align [15]. When the meaning of terms isn't shared and the process isn't clearly mapped out, it blurs transparency and makes it hard to move from a single case like one city's traffic model to wider regulatory or design guidelines [16].

In industrial scent monitoring, these limits make it tough to reuse data from one site to another, to gauge how instruments really perform in the field, and to build predictive models that blend chemical traces, weather shifts, and human scent impressions [9] [17]. The real challenge is pushing past simple connections and syntax, aiming for structural interoperability that lets diverse EIoT nodes share, interpret, and combine scent data seamlessly like tracing the faint aroma of coffee across different sensors.

This article tackles one key question: how can we design a structural interoperability framework that lets different olfactory and environmental sensor networks work together to deliver clear, traceable, and practically useful data in industrial spaces something that goes beyond just matching communication protocols? Building on earlier research in semantic standards, environmental ontologies, and interoperable sensor networks [13] [11] [12], we argue that the framework should clearly define syntactic, structural, and semantic elements while addressing temporal precision and the limits of distributed processing that shape real-world EIoT systems like sensors scattered across a humid greenhouse.

Our idea introduces a mixed interaction model for tracking scents in industrial settings, built on a simple, rule-

based way of observing like noting the sharp tang of metal in the air. First, we define a unified data model that ties together observations and metadata the sensor's range and calibration, its maintenance state, uncertainty markers, and provenance all shaped by work to standardize models across health and environmental sciences [18] [7]. Second, we introduce ontological vocabularies to describe emission sources, process activities, environmental settings, and perceptual events, so data from different platforms and case studies align and can be found easily for instance, linking a factory's exhaust record with nearby air quality readings [11] [19]. Third, we formalize temporal structures so multi-rate time series and event logs line up and combine under clear, transparent rules not the hidden, platform-specific quirks that used to govern them, like clocks ticking out of sync in a noisy server room [12].

To show how temporal fusion works, imagine two observation streams $x(t)$ and $y(t)$ each ticking along at its own pace, one every Δt_x seconds and the other every Δt_y . We define the cadence-aware continuous-time reconstruction $\hat{u}(t)$ follows, its rhythm echoing each pulse like footsteps on a quiet floor.

$$\hat{u}(t) = \frac{\sum_i K_x(t - t_i)x_i + \sum_j K_y(t - t_j)y_j}{\sum_i K_x(t - t_i) + \sum_j K_y(t - t_j)}, \quad (1)$$

Here, K_x and K_y are kernels shaped by sensor response times and uncertainty models, while t_k marks each observation's timestamp, like a quiet tick in a steady sequence. This setup clearly shows how uneven sampling rates and varying measurement quality shape a common hidden signal, and it can stretch to handle vector-valued data and probabilistic forms as well. At a more abstract level, event tokens e_k each marked with a timestamp, like a sudden complaint or a late-night maintenance fix are mapped into the observation space through a labeling function.

$$L(e_k) = e_k \rightarrow (c_k, W(T_k)), \quad (2)$$

Here, c_k represents an ontology class that defines the event type, while T_k acts as a window of influence on the reconstructed signal like a brief pulse of light used for causal attribution, supervised learning, or rule-based reasoning. We test the framework by building a rich, mixed dataset that mirrors a real Bralirwa brewery rows of cooling vats and sacks of grain helping ground the model in everyday detail. Four native-rate data streams feed into the system: second-level scent traces from odor monitors, minute-by-minute environmental readings, hourly process and air-quality metrics, and event-driven logs capturing each incident or complaint. The testbed serves to run alignment procedures, fuse data over time, mediate semantics, and coordinate services across platforms much like the latest semantic IoT frameworks built for

environmental systems where sensors hum quietly in the field [7] [12].

The work blends a precise but stripped-down grammar for describing scent with a rich, mixed dataset like notes of citrus and smoke to push forward both theory and practice. In theory, it builds on current models of semantic and structural interoperability, pushing them into the little-studied world of tracking environmental odours like the faint scent of smoke drifting through city air [11] [9]. In practice, it serves as a consistent benchmark for testing how interoperability tools, data-fusion methods, and governance systems work together to make information exchange clear and shared awareness possible across factory floors, regulatory offices, and local communities [15].

II. COMPREHENSIVE LITERATURE REVIEW

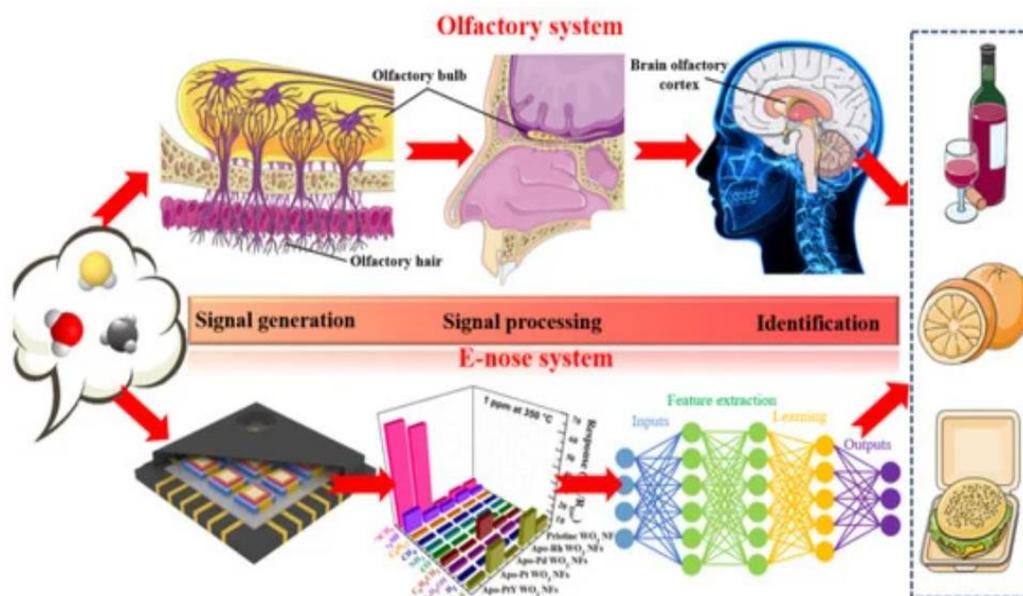


Figure 1. Electronic Noses and Gas Sensing [25]

Recent studies at the sensor level highlight how physical factors temperature, humidity, even the rush of gas can disrupt response stability and selectivity, while also reviewing hardware and algorithmic ways to correct for them [26]. High-speed designs now use nanostructured materials and streamlined flow paths to hit millisecond-level response times, letting robots trace shifting plumes and track rapid environmental changes like catching a faint whiff of smoke before it disappears [27] [28].

These breakthroughs rest on clear, detailed models that track how sensors move and react like a lens adjusting focus in shifting light. A single sensor usually records data that looks like the true concentration field blurred a bit by its impulse response and sprinkled with noise,

A. State-of-the-Art in Electronic Noses and Gas Sensing

Electronic noses have grown from basic gas detectors that sensed just one compound into complex arrays of sensors that mimic how living creatures catch a scent like tracing the faint smell of rain before a storm, as shown in Figure 1. Modern platforms usually blend a mix of chemically distinct transducers metal oxide semiconductors, conducting polymers, quartz crystal microbalances, and surface acoustic wave devices often paired with tiny chromatographic or preconcentration stages to sharpen selectivity and widen the dynamic range [20] [21]. In food and postharvest quality control, these sensor arrays, paired with pattern recognition, already match panel tests in telling apart complex aroma profiles the subtle shift from citrus to honey is no longer lost on them [22] [23] [24].

$$s(t) = (h * c) + n(t) \int_{-\infty}^{\infty} h(\tau)c(t - \tau)d_{\tau} + n(t), (3)$$

This drives the use of deconvolution, drift correction, and transfer learning methods to calibrate data across different sites and instruments, much like adjusting a lens to sharpen a blurred image [29] [30]. Further downstream, feature extraction and gas recognition have moved beyond simple linear projections to kernel methods and deep models that work straight on raw time-series data, often drawing on the full response from the first sharp rise to the steady equilibrium and slow recovery of each cycle [31] [32]. From breath analysis to food quality checks and toxic-gas monitoring, experts agree that sensor arrays paired with modern learning algorithms can reliably tell apart complex mixtures in the lab but taking that success into the field still

demands clear modeling of shifting conditions, context factors, and uncertainty [33] [34].

B. The Interoperability Crisis : Structural vs. Temporal Barriers in IoT

In environmental IoT, cheap gas sensors have spread faster than the shared data models meant to organize them, like a swarm of tiny monitors blinking on every street corner. True digital transformation now depends on interoperability, especially in data-heavy fields like medicine and environmental monitoring, where different systems have to share, interpret, and reuse what they observe instead of just passing along packets of data. Here, it helps to tell structural and temporal interoperability apart the difference is like comparing the frame of a clock to the way its hands move through time [35] [36].

Structural interoperability means keeping data models, metadata, and ontologies aligned so their syntax and meaning fit together cleanly like puzzle pieces snapping into place. Even if devices use the same communication protocols, their data can't really be compared one might log volts, another counts, each with its own limits, calibration quirks, quality tags, and traces of where the reading came from. Clinical registry data show that when domain experts agree on shared information models, multi-level modeling and archetype reuse can greatly boost semantic interoperability much like fitting precise puzzle pieces into place [37]. In the environmental IoT world, observation grammars are still scattered, and ontologies tend to gloss over sensor abilities, event types, or uncertainty details like ignoring how a humidity sensor drifts after rain which makes it hard to discover resources, automate service composition, or ensure solid traceability [38].

Temporal interoperability adds another layer of constraints, like threads that tighten each moving piece in place. Different nodes generate asynchronous, multi-rate streams $x_i(t)$, each sampled at its own native cadence Δt_i and shaped by a unique response kernel $h_i(t)$ with its own latency like instruments striking slightly out of sync in a dense rhythm. If you lump data into broad time windows, you can blur quick bursts of scent, and careless resampling can twist the signal creating echoes and false links that were never there. So in formal treatments, latent fields are built through kernel-weighted fusion imagine blending signals with a smooth taper, for example,

$$\hat{u}(t) = \frac{\sum_i \sum_k K_i(t - t_{ik}) x_{ik}}{\sum_i \sum_k K_i(t - t_{ik})}, \quad (4)$$

where each kernel K_i captures both the sensor's physical response dynamics and the uncertainty tied to sensor i , like the faint buzz you hear when it starts collecting data. Still, you can't reliably set up those operators unless you've got metadata detailing cadence, timestamps, response times, and confidence intervals like notes tracking each beat of a metronome. Research on indexing frameworks for multi-source IoT streams

likewise reveals that temporal query accuracy and speed hinge on clearly defining stream rates, window sizes, and alignment rules for instance, when sensors report data every few seconds instead of minutes [39]. We're facing a socio-technical interoperability mess structural standards overlook how time actually behaves, while time-series systems rarely carry the deeper meaning of the sensors or where their data truly comes from.

C. Limitations of Existing Predictive Models for Odor Nuisance

Predictive models of odor nuisance usually break down into three main groups, like categories you can almost smell separating in the air. First, physics-based dispersion models from steady Gaussian plumes to shifting Lagrangian particles and detailed computational fluid dynamics setups calculate how concentrations spread, using emission rates and weather data, much like tracking smoke twisting in a light breeze. For a point source with *height* H , wind *speed* U , and dispersion coefficients σ_y and σ_z that depend on stability, the classical steady model expresses the mean concentration as,

$$C(x, y, z) = \frac{Q}{2\pi U \sigma_y \sigma_z} \exp\left(-\frac{y^2}{2\sigma_y^2}\right) \left[\exp\left(-\frac{(z-H)^2}{2\sigma_z^2}\right) + \exp\left(-\frac{(z+H)^2}{2\sigma_z^2}\right) \right]. \quad (5)$$

These models help with regulatory assessments, but they depend on precise, time-resolved emission data and often stumble when faced with sporadic, multi-source industrial sites or the tangled sub-grid details surrounding a facility's smokestacks [20].

Second, data-driven methods tap into real-world links among sensor readings, weather patterns, and what people actually observe like temperature spikes before a storm. They span everything from regression and time-series models to machine-learning classifiers that connect gas sensor arrays and surrounding conditions to odour indices or the likelihood of complaints for instance, when a sharp sulfur smell drifts through the air [21]. Recent research suggests composite odor air-quality indices that merge intensity, hedonic tone, and concentration into one measure, yet they still depend on site-specific calibration and often assume the data arrive as neat, evenly spaced readings rather than uneven streams that lag or drift like sensors catching faint traces of smoke in shifting air [21][6].

Third, hybrid models blend mechanistic rules with statistical learning for instance, they might feed dispersion results into complaint prediction as structured covariates, like threads woven into a larger pattern. Across all three families, uncertainty quantification often feels improvised people report point predictions or class probabilities but rarely carry forward the effects of measurement error, calibration drift, model structure gaps, or noisy human labels that blur the data like

smudged ink. Most published studies rely on small, site-specific datasets where complaint records, operational logs, and sensor data rarely line up like clocks set a few minutes apart making it hard to compare results across sites or build solid benchmarks. A steady gap remains between the sensor's refined signal work each pulse sliced and filtered with precision and the predictive models that can produce transparent, uncertainty-aware forecasts ready for real-world decision-making and debate among stakeholders.

D. Identification of the Research Gap : The Need for Interaction Paradigms

All the research circles back to a single paradox, sharp as a line drawn in chalk. On one hand, breakthroughs in materials science, sensor design, and machine learning have sharply boosted how sensitive, selective, and fast electronic noses are across many fields for instance, they can now pick up a faint whiff of coffee in seconds [20] [31] [27]. Meanwhile, the IoT systems that host these devices are still scattered a patchwork of half-finished fixes for semantic alignment, makeshift timing syncs, and only thin support for traceable data fusion and service composition [35] [37] [39]. Most current standards and ontologies handle structural details, timing aspects, and processing history separately like listing each on its own page instead of linking them together.

What's still missing is a clear way for environmental IoT systems to interact in odour management a lean but expressive grammar of observation that ties together structural semantics (like observable types, units, detection limits, calibration, and provenance), temporal semantics (cadence, response kernels, latency, windowing policies), and computational contracts that define where and how fusion, filtering, and prediction happen, and under what quality constraints. This approach would let fusion operators like convolutional reconstructions and kernel-based multi-sensor estimators draw their parameters straight from metadata, creating reproducible mappings that link real events such as complaints, process upsets, or regulatory exceedances into a shared observation space that feels as clear as watching colored lights sync on a single monitor. Turning these interaction rules into first-class, machine-ready tools would let environmental IoT systems move past one-off, isolated setups and build connected, reliable networks for managing industrial odours like mapping scent patterns drifting across a factory yard.

III. THEORETICAL FOUNDATIONS OF THE PARADIGM

A. Defining the Heterogeneous Interaction Paradigm

The heterogeneous interaction paradigm defines a clear, machine-readable contract that directs how varied sensing agents, human observers, and computational services share, interpret, and combine olfactory data like the sharp scent of oil in a factory floor to manage industrial environments.

Instead of handling sensor readings as plain time series, this approach treats every measurement as a structured artifact a tuple $o = (y, \sigma, \theta, \tau, \pi)$, where y is the reported quantity (say, concentration or the sharpness of an odor), σ defines structural semantics like type, units, and detection limits, θ describes sensor state and behavior (calibration settings, drift patterns, response kernel), τ conveys temporal semantics such as cadence, timestamp precision, latency, or windowing rules, and π traces provenance, noting the device's ID, firmware version, processing history, and confidence scores. Treating σ , θ , τ , and π as first-class components fits with semantic interoperability methods in diverse IoT networks, where raw device payloads like a sensor's burst of temperature data are normalized into canonical forms before reasoning or analytics take place [40] [41].

In this model, algorithms don't handle hidden device feeds anymore they work with clear observation artifacts, each carrying its own defined meaning, like data tagged with a bright label. Fusion operators, timing policies, and predictive models pull their settings from metadata instead of site-specific code, separating how algorithms are built from vendor quirks and local preprocessing steps like skipping the custom script a technician might tweak on a chilly server room morning. This design builds on cross-domain mediation gateways that bridge clashing IoT standards through a tight set of shared concepts and event types [42], yet it goes further linking structural, temporal, and provenance semantics within one cohesive observation grammar, like threads woven into a single patterned fabric. The model takes shape through three layers that constantly interact, like ripples overlapping on a pond. The perception layer gathers raw signals and human inputs, tagging each with fine-grained notes on sensor behavior and quality just as measurement-driven IoT designs aim to shift intelligence out to the edge [43]. The mediation layer uses a lean yet expressive grammar to translate varied schemas into standard observation models, then offers clear computational contracts for alignment, fusion, and quality checks much like the steady pulse of link-layer normalization in mixed sensor networks [40]. The orchestration layer builds higher-level services detection, attribution, forecasting while enforcing strict quality limits and provenance-aware governance, tapping blockchain-based data systems where smart contracts capture trust and auditing rules for sensor networks, as precise as a timestamped log flickering on a monitor [44] [45].

When response kernels and cadence markers are built into the perception artifacts, temporal fusion works directly on kernels and timestamps instead of clumsy, resampled arrays. It allows kernel-weighted reconstructions and uncertainty propagation that stay true to each sensor's unique behavior, while provenance terms make sure every fused estimate can be tracked back to its sources, calibration state, and full processing chain right down to the hum of the sensor itself. The idea is that

an observation matters only when its numeric value carries along the setting in which it was measured a principle that keeps analytics reproducible, decisions traceable, and policy checks grounded, much like noting the air's sharp tang before logging an odour reading.

B. Taxonomy of Data Archetypes in Industrial Environments

Industrial smell monitoring operates inside a data network that's anything but uniform its sensors, timing, reliability, and oversight each follow their own rhythm, like valves clicking open and shut in uneven bursts. A tight framework of data types forms the spine of the interaction grammar, shaping how each source gets encoded, merged, and stored like sorting vivid bits of color into one clear mosaic. The first archetype takes shape through primary gas-sensing streams, using fixed chemical sensors and electronic noses that catch faint traces of smoke in the air. These devices produce time-series data sometimes steady, sometimes in quick bursts whose behavior can be described as $x_i(t) = (h_i * u)(t) + n_i(t)$, represents the hidden odor field, $h_i(t)$ the sensor's response pattern, and $n_i(t)$ the faint hiss of measurement noise. Because high-frequency behavior, cross-sensitivities, and drift all come into play, the calibration record, detection limits, and kernel settings have to live inside the observation data itself much like advanced gas-sensor fusion and soft-sensing systems that lean on detailed sensor models to keep predictions steady [46] [47].

The second archetype includes support sensors that track conditions like wind speed and direction, temperature, humidity, and key process details for instance, the sharp drop in temperature before dawn. These series adjust how dispersion and sensor transfer functions behave, often running on their own sampling rhythms and reliability patterns like instruments tuned to slightly different beats. Research in environmental IoT shows that ignoring shifts over time like sudden drops in soil moisture after rain can skew results and waste resources [48] [43]. In this framework, auxiliary sources hold exact units, sampling meanings, and trust markers like crisp timestamps on a sensor readout so fusion operators can lock in proper alignment windows and maintain strict quality gates.

The third type comes from people themselves things like organized complaint logs, an inspector's handwritten notes, even crowdsourced odor reports that describe a sharp chemical smell hanging in the air. These observations come in late, feel scattered, and depend on personal judgment like notes scribbled after the moment's already passed. Studies on trust- and privacy-aware IoT data sharing stress that human input should be seen as a probabilistic, bias-tinged signal like a sensor flickering under bright light rather than a flawless, noise-free label [49] [50]. So the grammar encodes how severe a complaint feels, its emotional tone, and when it occurred as probabilistic labels with clear measures of uncertainty and bias

letting later models carry that human doubt forward instead of clinging to one shaky story.

A fourth archetype covers event and administrative records things like maintenance logs, emission inventories, and incident reports that track every fix and flare. These records often serve as the ground truth for attribution work, yet they come in bursts patchy, uneven, and rarely in sync with time, like scattered notes on a shifting timeline. Industrial monitoring studies reveal that taking administrative data at face value can hide unusual patterns or wrongly pin blame, like mistaking a flickering gauge for operator error. In this interaction model, the records include completeness and confidence details, tied through clear provenance links to both physical readings like temperature logs and human observations.

Finally, a fifth archetype emerges in the form of derivative artefacts dispersion maps, rebuilt scent fields, anomaly markers, and risk scores that flicker like warning lights. These outputs feed into higher-level fusion and governance decisions, so they have to keep clear records of their origins and uncertainty much like an entropy-based fusion model that handles every intermediate layer as a signal carrying meaning, not as leftover noise.

When we map these archetypes into the grammar, clear design rules emerge like patterns snapping neatly into place. Fast sensor data flows into hot, entropy-aware buffers, while admin and provenance records stay locked on immutable or blockchain-based ledgers; human reports come through trust-filtered channels that guard privacy and build in attack resilience, like a sealed conduit humming quietly in the background. This taxonomy turns diverse data sources into connected, auditable pieces of one contract-driven information space, like fitting mismatched puzzle pieces into a smooth, unified frame.

C. Information Entropy and Preservation in Multi-Scale Systems

Beyond simply getting systems to talk to each other, this approach runs up against the need to keep every bit of meaningful detail intact as data are sampled, squeezed down, and blended across scales like preserving the sharp glint of a sensor's reading even after heavy compression. Industrial smell-detection systems work across scales quick bursts of scent flicker through individual sensors, slow waves of dispersion drift through the air, and only later come a few scattered human complaints. Entropy-based studies in smart sensing and process monitoring show that simple averaging or heavy compression can wipe out the very intricate signal patterns the flicker of a sensor spike that best reveal anomalies and faults [51] [52]. From an information-theory angle, picture the generative model again like tracing how a spark of data spreads through its layers $x_i(t) = (h_i * u)(t) + n_i(t)$, with

a hidden field $u(t)$ and a mix of uneven kernels h_i , each shifting like ripples across water.

As h_i spreads, the Fisher information about the high-frequency parts of u fades ; temporal resampling that overlooks the kernel's shape dulls those sharp, fleeting transients even more. Let X stand for the full set of observation streams, and let $F(X)$ capture them in a single fused image, like threads woven into one bright strand. In an entropy-aware design, the goal is to pick fusion mappings that squeeze out the most mutual information, $I(u; F(X))$, while staying within limits on bandwidth, energy, and storage. Information-theoretic fusion frameworks for multi-source industrial data show that conditional entropy and mutual information can guide dynamic fusion, emphasizing the sensors and attributes that sharpen discrimination while shutting out those that merely echo background noise [53] [54]. In wireless sensor and IoT networks, compressed sensing provides a fresh angle it taps into sparsity in the right basis so signals can be captured at rates well below the Nyquist limit, yet still hold enough detail to rebuild, like tracing a full picture from a few bright pixels [55]. For smell monitoring, this means we can cut communication costs by sending task-focused projections or symbolic snapshots of how the sensors change so long as the encoding reflects the right dimensions and timing, like capturing the brief puff of scent after a door opens.

Entropy plays a key role in shaping how we sample and store data like adjusting the rate of snapshots to match the system's shifting patterns. Studies on complexity-entropy sampling show that by tweaking the sampling period to match the estimated statistical complexity, a system can keep its natural rhythm like a pendulum's steady swing while cutting data volume dramatically [56]. In this new model, moments of high entropy spikes the system flags through local entropy or complexity checks stay in sharper detail and for longer stretches, while the calmer, low-entropy background gets quickly compressed or reduced to a brief summary. Studies on lossless and near-lossless compression in industrial monitoring show that entropy limits set a solid benchmark for how much data can be compressed without losing any value for later analysis for instance, sensor readings stay as sharp as the original feed [57] [58].

Most importantly, the observation records and their resulting products carry forward clear traces of entropy and uncertainty, like faint fingerprints that show exactly how much we still don't know. Process-monitoring frameworks that choose modeling data through entropy-based rules show how these budgets can anchor reliable detection and diagnosis, even amid the shifting noise and heat of complex industrial settings [52] [59]. When it comes to smell regulation, rules don't have to focus only on concentration limits they can also set minimum information standards, like requiring that blended scent data keep a defined level of shared information with incoming

reports or logged complaints that trace when and where odors spike [53] [54].

By treating entropy and information preservation as core priorities right alongside structure, timing, and provenance the mixed-interaction model ties the technical framework to the demand for traceable, uncertainty-smart decisions in industrial odor control, where sensors catch faint chemical notes in the air [44] [60]. Information-theoretic metrics serve as both blueprints for fusion algorithms and binding terms in data-governance agreements [50] [61], making sure the layers of abstraction needed for scale don't quietly strip away the raw, evidential weight of the data that underpins every regulatory or civic decision.

IV. THE PROPOSED CONCEPTUAL FRAMEWORK : A CAUSAL-CENTRIC ARCHITECTURE

The proposed framework organises heterogeneous urban-industrial signals into an event-driven, causal-centric pipeline that maps perceived odor episodes into auditable governance actions. Conceptually, it couples structural causal reasoning with time-series analytics, in line with recent causal discovery schemes for observational and interventional temporal data [62] and with work contrasting purely predictive against genuinely causal models [63]. At its core, an odor event $e_{odor}(t)$ triggers an operator

$$P: (S, E, V, I) \times e_{odor}(t) \rightarrow A_t, \quad (6)$$

which transforms SEVI sources SCADA (S), IoT MOS sensors (E), meteorology (V) and complaints (I) into a set of recommended actions A_t under quantified uncertainty and provenance constraints.

Causal tags and provenance tokens are introduced at ingestion and preserved through normalization, mediation and decision layers, echoing label-aware causal feature design [64] and fine-grained provenance management in linked-data systems [65] and big-data pipelines [66]. Time-lag correction defines an effective time $t_{eff} = t_{meas} - \Delta t$, ensuring that inferred relationships respect temporal precedence.

The resulting unified stream feeds ensemble predictors and counterfactual evaluation, enabling industrial actors to query expectations of the form $E[Y | \text{mid}do(A = a)]$ for odor severity Y under candidate interventions A . In contrast to static warehouses, the architecture behaves as an adaptive, causal governance engine for odor-related nuisance. **Figure 2.** The SEVI Interaction Framework Architecture : event-driven pipeline from SEVI sources through normalization and mediation to causally annotated governance outputs.

A. Layered Interaction Architecture and the SEVI Mapping

The architecture is structured as a four-layer interaction stack operated in an event-driven manner, where the primary orchestration token is a detected odor episode. When

$e_{odor}(t)$ occurs either via MOS sensors or citizen complaints only the corresponding temporal neighbourhood is propagated through the stack, following event-driven sensing and computing paradigms that minimise redundant sampling and processing [67].

Zone 1 (Physical Layer) hosts SEVI sources as stochastic processes $\{X_k(t)\}_{k \in \{S,E,V,I\}}$. Their semantic roles mechanistic drivers (S), exposure indicators (E), atmospheric modifiers (V), and human observations (I) mirror causal role differentiation in modern causal feature selection and diagnostic schemes [63, 64].

Zone 2 (Normalization Layer) applies a lean operator N that maps each raw record $r_{k,t}$ to an enriched tuple

$$\tilde{r}_{k,t} = (x_{k,t}^*, u_k, \sigma_{k,t}, c_{k,t}, \pi_{k,t}), \quad (7)$$

injecting harmonised units, uncertainty descriptors, causal tags and provenance, in line with lightweight annotation strategies

in soft sensing and provenance-aware data management [65, 66].

Zone 3 (Mediation Layer) estimates source-specific lags Δt_k , computes

$$t_{eff,k} = t_{meas,k} - \Delta t_k, \quad (8)$$

and resamples all streams onto a common grid, producing fused vectors z_j with covariance Σ_j and aggregated causal/provenance metadata. This reflects best practice in causal time-series modeling where lag structure is explicitly inferred rather than assumed [62].

Zone 4 (Application Layer) consumes the unified, annotated series in Random-Forest-based predictors and a decision-support system, closing an event-driven loop from odor perception to ranked interventions and counterfactual evaluation [63].

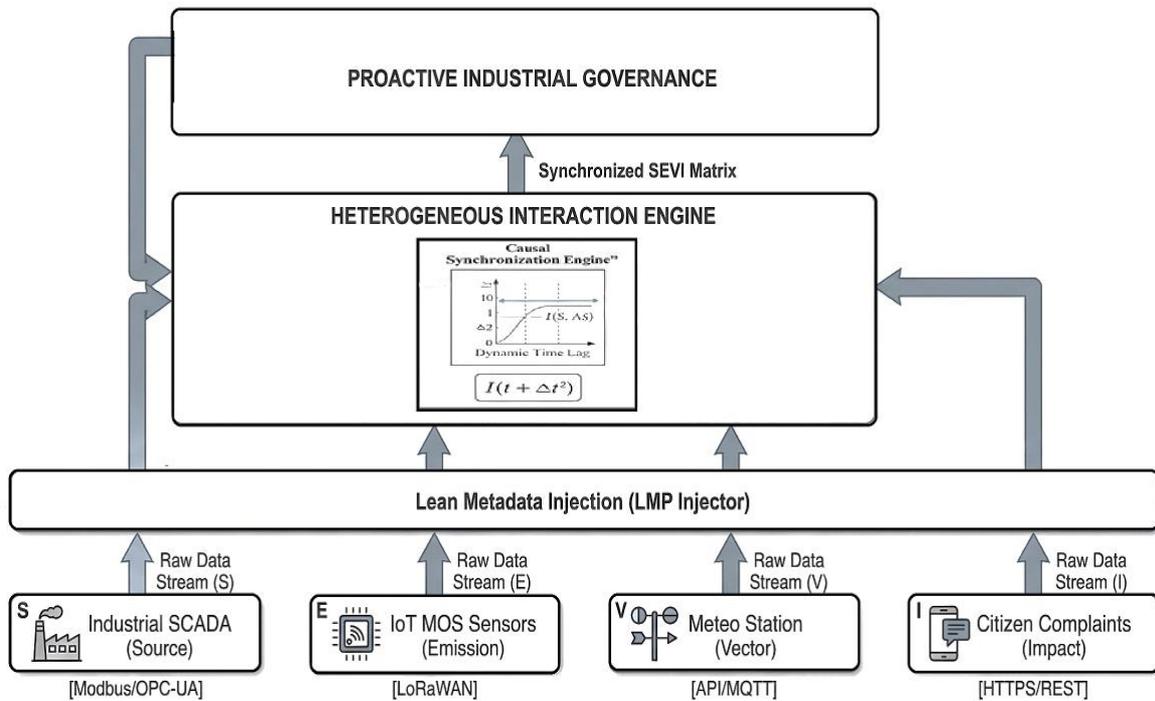


Figure 2 : The SEVI Interaction Framework Architecture

B. Formalizing the SEVI Interaction Paradigm

The SEVI paradigm treats each odor episode ω as a *reconstructed causal object* rather than a passive aggregation of records. For a given episode, the system instantiates $\mathcal{C}(\omega) = (G_\omega, \theta_\omega, U_\omega, \Pi_\omega)$, where G_ω is a directed causal graph over SEVI variables, θ_ω collects estimated causal effects, U_ω encodes multi-source uncertainty and Π_ω is a fine-grained provenance trace from graph edges to original SEVI records. This design echoes log-data-driven causal reasoning pipelines in manufacturing, where raw operational logs are transformed

into causal structures supporting root-cause analysis [68], and federated Bayesian causal inference frameworks that explicitly parameterise effects and uncertainties across distributed plants [69].

Temporal precedence is enforced through source-specific lag correction $t_{eff,k} = t_{meas,k} - \Delta t_k$, followed by alignment on a common grid; residual ambiguity is captured by confidence scores that weight downstream inference. Evidence from SCADA, MOS sensors, meteorology and complaints is then integrated into G_ω via reliability-aware

fusion, so that counterfactual queries of the form $E[Y|do(A = a), C(\omega)]$, can be evaluated for candidate interventions A . The scientific originality lies precisely in this *episode-level causal reconstruction*, which upgrades SEVI from archival storage to a provenance-preserving causal narrative, in line with recent calls for operational, decision-oriented causal models in industrial environments [68, 69].

C. The Lean Metadata Profile (LMP) for Causal Integrity

The Lean Metadata Profile (LMP) specifies a minimal yet expressive interaction contract that guarantees causal integrity across normalization and mediation. Each normalized SEVI record is represented as $\tilde{r}_{k,t} = (x_{k,t}^*, u_k, \sigma_{k,t}, c_{k,t}, \pi_{k,t})$, where $x_{k,t}^*$ is the unit-harmonised value, u_k its physical unit, $\sigma_{k,t}$ an uncertainty descriptor, $c_{k,t}$ a causal role tag, and $\pi_{k,t}$ a provenance token. Such compact, causality-aware metadata mirror industrial causal pipelines that enrich log streams with orientation rules and variable-selection steps to support structure learning and effect estimation [68, 69].

In event-driven operation, an incoming complaint I_i at time t_0 triggers a mediation query over a temporal window $W(t_0, \tau)$. The SCADA interaction contract is expressed as

$$S(W(t_0, \tau)) = \{\tilde{r}_{S,t}: t_{eff,S} \in W(t_0, \tau), \pi_{S,t} \in \Pi_{spatio-op}, c_{S,t} \in C_{driver}\}, \quad (9)$$

ensuring that only lag-corrected, spatially consistent, causally relevant SCADA records are fetched when a human-reported odor event arrives. Returned tuples preserve $\pi_{S,t}$, so each edge in G_ω remains auditable back to individual SCADA messages, aligning with traceable, operator-supporting causal reasoning for manufacturing systems [68, 69].

D. From Big Data to Proactive Governance

Proactive governance reframes big data as a continuously interpreted evidence stream rather than a static repository. Lean Metadata Profiles (LMP) operationalise this shift by enforcing six core attributes on every record. Time encodes observation and validity intervals, enabling temporal joins and trend detection in provenance graphs for supervision and security analytics [70]. Space captures geolocation and spatial resolution, supporting provenance-aware spatial infrastructures and fitness-for-use assessments in environmental and urban monitoring [71]. Unit standardises measurement semantics so that heterogeneous sensor and administrative streams can be fused and normalised, a prerequisite for reliable automated annotation and prediction [72]. Role (SEVI) specifies stakeholder function in the socio-technical system (source, enforcer, validator, impacted), making it possible to attribute responsibility and model incentive-compatible interventions in governance-by-data regimes [73].

Trust formalises uncertainty through a confidence score, for example $C =$

$f(q_{data}, q_{provenance}, a_{algorithm}, s_{source})$, so that automated actions can be thresholded and audited. Finally, *Provenance* records derivation chains and processing steps, enabling traceable decision pipelines and real-time anomaly detection over complex workflows [70]. Embedded into policy dashboards and regulatory workflows, LMPs turn dispersed big data into an adaptive, closed-loop governance fabric where each decision is both anticipatory and retrospectively accountable.

V. CASE STUDY CONTEXTUALISATION : THE BRALIRWA BREWERY

A. Characterization of the Bralirwa Brewery Testbed : Topography, Land Cover, and Industrial Dynamics

The Bralirwa brewery is located in the steep, dissected volcanic highlands of western Rwanda, where sharp gradients and deeply incised valleys accelerate runoff, sediment transport, and the downstream propagation of pollutants toward lacustrine receptors such as Lake Kivu [74]. In this setting, short, flashy catchments and narrow valley bottoms intensify the coupling between point discharges from industrial facilities and diffuse inputs from agriculture, heightening the vulnerability of aquatic ecosystems and communities along the lake margin [75]. At the national scale, high population density and the intensifying energy–water–food–land nexus amplify competition among domestic, agricultural, and industrial water uses, placing a premium on catchment-sensitive abstraction, storage, and reuse strategies for large water users such as breweries [76].

Land cover around Nyamyumba comprises a heterogeneous mosaic of smallholder cropland, short-rotation Eucalyptus woodlots, and agroforestry. Although recent measurements suggest that well-managed eucalypt stands may have moderate transpiration relative to rainfall, their density and rotation length can still alter baseflow, dry-season resilience, and hence the reliability of industrial water supply [77]. These biophysical dynamics intersect with broader regional air-quality concerns: observations from the Rwanda Climate Observatory show that local and transported combustion emissions already elevate background pollutant concentrations, pointing to a growing need to manage both greenhouse gases and short-lived pollutants in rapidly developing East African economies [78]. Figure 3 situates the brewery as an emission epicentre embedded within nearby residential receptors, access roads, and the local drainage network connecting plant infrastructure to adjacent slopes and the lake shore.

Institutionally, Rwanda's green-growth agenda and emerging regulatory instruments create an enabling framework for stringent effluent standards, yet empirical assessments reveal incomplete compliance and persistent discharge of inadequately treated industrial wastewater, underscoring the

importance of capacity building, monitoring, and adoption of best-available treatment and odour-control technologies [79].

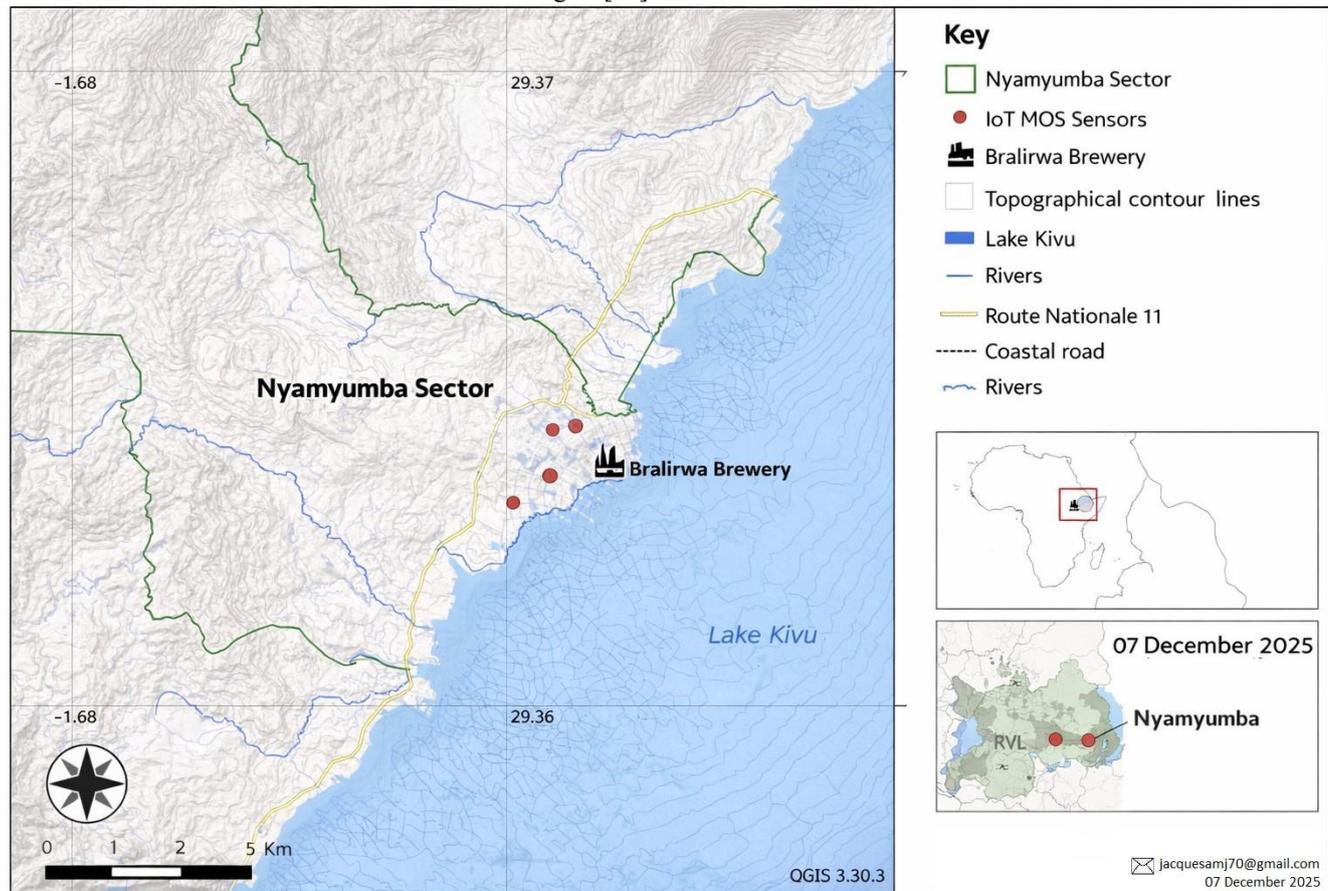


Figure 3 : Odor Pollution Monitoring and Management Site (Nyamyumba Sector)

B. Testbed Design : Simulating the 4 Streams in a Tropical Context

The experimental testbed models a mid-sized Bralirwa brewery, built to measure and fine-tune four linked resource streams water, energy, organic matter, and nutrients like tracking every drop and spark that keeps the place running. The process train runs from mashing and lautering to boiling the wort, then moves through fermentation, conditioning, and packaging, before ending at the wastewater treatment line, where solids are first separated, then flow through an anaerobic lagoon and into a constructed wetland that smells faintly of earth and hops. Brewers gather spent grain and trub in separate bins, sending them off for animal feed or into an anaerobic digester that turns the mix into biogas to power a compact heat-and-power unit humming in the corner. This system-level setup matches advanced brewery water networks and retrofit models, where process units and treatment stages sit inside one tightly linked optimization framework that hums like a well-tuned line of copper pipes [80] [81].

Boundary conditions weave local climate forcing into the system, like heat pooling on a sun-baked patch of asphalt. Ambient temperature, noted as T_a , factors into the energy

balance for heat transfer and fermentation, while rainfall's shifts in influent flow and makeup are captured by a time-varying function, $Q_r(t)$, that rises and dips like water spreading across wet pavement. We express mass continuity for a general stream s as follows:

$$\frac{dm_s}{dt} = \sum_{in} m_{in,s} - \sum_{out} m_{out,s} - r_s, \quad (10)$$

Here, m stands for the mass flow rate, while r_s captures the internal consumption or transformation like heat shifting through a metal pipe. The thermodynamics within a control volume of size V are determined by.

$$\rho c_p V \frac{dm_s}{dt} = Q_{in} - Q_{out} - Q_{loss}, \quad (11)$$

with ρ as density, c_p specific heat capacity, and Q heat fluxes. The biochemical methane potential of spent grain is approximated by

$$BMP = \sum_i C_i Y_i, \quad (12)$$

Here, C_i represents the concentrations of the organic compounds, while Y_i denotes the specific methane yields

measured from hands-on assays of brewery residues [82]. This model makes it possible to explore, step by step, how feedstock makeup and temperature both shaped by the thick, humid air of the tropics change the way gas is produced.

Operational control variables include the process water reuse fraction a_ω , the solids allocation fraction a_b , for digestion, and the nutrient recovery efficiency η_n , which measures how much nitrogen and phosphorus in the digestate and treated effluent make their way back to the soil like the sharp scent of damp earth after rain reminding you nutrients have returned. These design levers work like decision variables in the multi-objective optimization of industrial networks, where engineers choose recycle flows, storage tanks, and treatment capacities to cut both cost and resource use [81] [83]. Sensors track flow rates, chemical oxygen demand, biogas composition, electrical output, and nutrient levels data that help estimate key performance indicators as conditions shift in the humid tropical air and as operating policies change.

C. Theoretical Mapping : Applying the Framework to the Brewery Case

The theoretical framework conceptualises the brewery as a dynamical system in which four interacting stocks represent water, energy, organic material, and nutrients. Let the state vector be

$$X(t) = \begin{bmatrix} W(t) \\ E(t) \\ M(t) \\ N(t) \end{bmatrix}, \quad (13)$$

where W denotes water held in process and storage units, E the cumulative useful energy available from both the grid and on site generation, M the mass of organic solids such as spent grain and trub, and N the inventory of recoverable nutrients in aqueous and solid phases. The system evolution is described by $\frac{dX}{dt} = F(X, u, p, t)$. Here, u represents a vector of control variables that includes a_ω and a_b , while p is a set of parameters capturing climate effects, regulatory limits, and the texture of the feedstock think of how damp or coarse it feels.

The way the four streams interact shows up in the off-diagonal entries of the Jacobian matrix $\partial F/\partial X$, like faint cross-currents rippling through water. For instance, sending more spent grain to the digester boosts E through extra biogas, cuts the M left for sale or cheap disposal, and changes N as the digestate thickens with nutrients ready to return to the fields [82]. We estimate the energy gained from digestion by measuring how much heat the body releases, $E_{biogas} = \eta_{conv} LHV m_{bio}$, where m_{bio} is the mass of biodegradable volatile solids, LHV is the lower heating value, and η_{conv} the conversion efficiency of the combined digestion and CHP chain. Nutrient recovery potential is expressed as $N_{recov} = \eta_m (N_{solid} + N_{liquid})$, with N_{solid} and N_{liquid} the nutrient content of solids and treated effluents routed to land application.

Process water reuse is constrained by quality standards. If C_{in} is the contaminant concentration in incoming water, $C_{reclaimed}$ the concentration in reclaimed water, and $C_{\omega,max}$, the allowable limit for reuse, the mixed concentration at the point of use satisfies, $C_{out} = (1 - \alpha_\omega)C_{in} + \alpha_\omega C_{reclaimed} \leq C_{\omega,max}$.

This setup mirrors the integrated water-network optimization models used in breweries, where water reuse and regeneration are limited by quality standards and by how each unit interacts like pipes humming softly as they pass filtered water from one tank to the next [80]. We frame the decision problem as a multi-objective optimization across a set planning horizon, balancing key goals that stand in for real priorities.

$$\max f_1 = E_{net}, \min f_2 = W_{cons}, \max f_3 = N_{recov}, \quad (14)$$

Bound by material and energy limits, the capacity of the equipment, and the pull of environmental rules like keeping the heat exchangers humming without breaking a single code. Comparable multi-dimensional trade-offs shape how energy and emissions are planned in other resource-heavy industries, where analysts use linear or nonlinear programming models to map out scenario-specific paths like tracing the curve of a rising demand line across a bright monitor [84]. Here, Pareto-optimal results show that pushing for stricter water-saving goals can drive up the energy needed for advanced treatment or cut down the share of organics sent to high-value energy recovery like fewer rich, dark sludges feeding the biogas units.

Rwanda's tropical highlands shape which parts of this solution space actually work best, much like steep green hills guiding where water can flow. Warmer conditions in the mesophilic range speed up anaerobic reactions and boost methane output, often pushing the best performance toward heavier digestion loads and greater use of on-site biogas for heat and power, as tropical bioenergy research has shown [82]. At the same time, shifts in hydropower output and raw material supply caused by climate swings add new timing patterns to pp, which the model can handle through scenario-based or stochastic extensions.

In the end, the framework links plant-level optimization with regional policy, connecting the fine-tuned hum of a factory floor to the broader goals shaping an entire region. Changes in rules for effluent discharge, new perks for combined heat and power, and updated biosolid land-use guidelines all line up with outside shifts in the parameter vector pp like a sudden gust turning a weather vane. Building on lessons from African environmental governance, which favors flexible, cross-sector strategies [85] [86], the model lets us test how different rules and incentives shift the limits of water, energy, material, and nutrient performance like watching ripples spread across a clear pond after a single drop. The Bralirwa brewery case works as a practical testing ground, using hard numbers to connect how factories manage their

resources with the region's larger path toward sustainability like matching the rhythm of clinking glass bottles to the pulse of a growing economy.

VI. Discussion

The new interaction model highlights clear, machine-readable metadata and rhythm-sensitive fusion tools, allowing mixed scent data to blend smoothly without any makeshift prep work. By using metadata like native cadence, response kernels, latency, and confidence intervals to parameterize fusion kernels and alignment rules, the framework separates fusion logic from device-specific code, fitting naturally into measurement-driven, cross-platform IoT systems that aim to abstract sensing away from vendor hardware [87] [88]. This design makes it easy to scale sideways across big sensor fleets and several industrial sites, cutting back on custom resampling pipelines that now block reliable, repeatable deployments and the smooth handoff of models from one location to another [89].

In practice, making a system scale means running work across many machines like dozens of servers humming in sync. Capturing lightweight metadata right at the edge things like device descriptors, calibration state, and uncertainty estimates keeps fusion operators such as kernel-weighted reconstructions

$$\hat{u}(t) = \frac{\sum_i \sum_k K_i(t - t_{ik})x_{ik}}{\sum_i \sum_k K_i(t - t_{ik})}, \quad (15)$$

computable without heavy centralization, echoing the fog-and-edge workflows common in environmental monitoring and query processing [90] [89] [91]. A layered setup handling quick cleanup on local devices, gathering results midway, then indexing in the cloud helps even out delay, protect data, and cut computing costs, as shown in scalable IoT and distributed early-warning systems [92] [93].

Two complementary traits form the backbone of generalizability like sturdy threads woven through a single fabric. Semantic minimalism a lean, clearly documented grammar of observation that notes types, units, detection limits, provenance, and time echoes enterprise interoperability and big-data metamodel projects that prefer tight core vocabularies instead of sprawling schemas [94] [95]. This kind of grammar makes it easier to connect vendor-specific schemas with shared ontologies, cutting down on fragile point-to-point links that block smooth reuse across sites. Second, provenance-aware uncertainty propagation broadens how probabilistic methods handle data quality and completeness in IoT databases, carrying that rigor through the entire fusion pipeline like tracing each sensor's faint pulse of light to the final decision [96]. By bringing sensor response models and confidence metrics straight into the predictive stage, we can test how well results hold up across sites with different sensor setups, reporting

habits, and weather patterns doing it all under clear assumptions instead of buried preprocessing steps.

Even so, a few limits still stand, solid as a closed door. Because legacy instruments vary so wildly and the network drops out especially on those relay links in far-off monitoring sites we'll need sturdy fallback plans like imputation rules, tiered quality checks, and brief sync windows that open whenever the signal flickers back [97]. Different ways people report data and how citizens pitch in can spark doubts about trust and how well systems withstand attacks, mirroring worries in environmental sensing and trust management where even a single faulty air-quality reading can cause confusion [98] [99]. To truly measure how performance, uncertainty, and computational cost grow with network size and variety, we need broad, real-world testing across different industrial sites, backed by open benchmark datasets like running sensors in a dozen humming factories to see what the data really says.

By weaving structural and temporal meaning right into the data layer, odour monitoring shifts from tallying past complaints to building a living, proactive information system that guides industrial oversight almost like catching the scent before it drifts. When every fused estimate links clearly to its sensors, calibration logs, and alignment windows, auditors can test attribution claims and gauge how well mitigation steps stand up to measured uncertainty pushing provenance-focused methods from early-warning and citizen-based environmental tools into the realm of industrial odour control, where the sharp tang of ethanol in the air tells the story [93] [100]. Cadence-aware reconstructions preserve fleeting transients instead of blurring data into rough averages, enabling near-real-time alerts and focused interventions right where the process unfolds much like fog-based early-warning systems that flag other environmental risks [91].

This interaction model also makes it easy to build modular services that follow policy rules, like pieces snapping neatly into place. Standardized interaction agreements linking detection, attribution, and forecasting services can lock in quality safeguards like complete metadata, recent calibration, and trust scores for citizen data so high-stakes choices never rely on shaky evidence, much like the trust-management and secure data-fusion methods used across environmental and industrial IoT systems [99, 101]. This clear split of responsibilities makes layered access work smoothly : regulators and local communities can view summary metrics, exceedance odds, and data trails showing where numbers came from, while operators keep the raw telemetry balancing openness with commercial secrets and cybersecurity needs [102] [103].

From a policy angle, uncertainty-aware fused fields open room for finer, more flexible rules like adjusting a dial instead of flipping a switch. Rather than setting hard cutoffs on raw sensor data, define compliance through probabilistic

odor-impact scores or complaint-risk indices shaped by dispersion, exposure, and annoyance models like measuring how far a faint smell travels on a damp evening aligned with broader efforts toward data-governance frameworks in IoT-based systems [104]. These tools encourage constant fine-tuning, built-in backups, and open exchange of interaction rules, because better-qualified data cut down on regulatory risk and how often disputes flare up.

Still, turning those benefits into reality isn't just about mastering the technology it's more than wires and code humming under your fingers. It calls for tight coordination on how observation grammars and metadata profiles are standardized, clear governance for data lifecycle and quality control, and pilot runs that prove metadata-driven fusion can sharpen attribution, boost responsiveness, and build trust even in tough, high-noise industrial settings. Insights from smart-building and smart-city systems show that lasting growth happens only when ecosystems blend interoperable platforms, solid governance, and services shaped around their users like a dashboard that actually listens to the people behind it [95] [98]. The proposed framework maps out a clear path to bring this ecosystem logic into industrial scent monitoring and forward-looking odor management for instance, tracking the faint tang of chemicals before it drifts outside.

VII. Conclusion and Future Directions

This article introduces a way of interaction that brings together structural and temporal meaning for environmental IoT, helping factories manage odours more effectively like tracing a faint chemical scent before it spreads. Its main contribution is a concise grammar of observation that simultaneously captures observable types, units, detection limits, calibration state, provenance, and time-based descriptors native cadence, response kernels, and latency with the precision and economy found in metamodels for IoT big-data and interoperability research [105] [106]; it's the kind of design that feels as crisp as a measured pulse on an oscilloscope. By tuning cadence-aware fusion and alignment operators with this

metadata, the framework separates fusion logic from specific devices or platforms, allowing kernel-weighted reconstructions and rigorous uncertainty tracking across mixed sensor fleets much like current IoT data-quality research guiding environmental monitoring. Provenance-aware descriptors and layered quality gates build on blockchain and provenance systems for IoT data governance, creating decision pathways that stay traceable and auditable, like footprints pressed in wet sand. Together, these pieces link raw sensor signal modeling with system-level governance, forming a lean, machine-ready contract that guides olfactory monitoring like mapping the scent data from a single rose into a citywide control system.

The companion empirical study will put this paradigm to work and push it hard, running it through a step-by-step validation program that tests every layer, like tapping each rung of a ladder before climbing higher. In the first phase, we'll roll out observation grammar and cadence-aware fusion across mixed fleets of high-end gas sensors and budget electronic noses, building on fresh breakthroughs in rapid artificial olfaction and machine-learning-based gas sensing. We'll test edge-level metadata capture and online kernel estimation by comparing them with controlled scent releases in the lab. The next phase will roll out across multiple sites, using relay and backhaul setups modeled on remote environmental monitoring frameworks, to test how robust the system stays when connections falter, data must be imputed, and synchronization happens on the fly like sensors flickering in a storm. In the third phase, the interaction grammar connects with provenance stores and policy-aware data-sharing systems built on blockchain and reputation-based governance for IoT, testing how precisely it attributes data and predicts complaint risks against a curated ground truth like measuring the sharpness of a sensor's readout under a clear morning light. Open benchmark datasets, along with a testing-as-a-service setup for diverse IoT systems, will make it easier to compare results consistently and help the community adopt the new interaction approach like lining up different devices on one shared testing bench.

Bibliography

- [1] G.-F. Victor, M. De Sario, S. Vecchi, L. Bauleo, P. Michelozzi, M. Davoli et C. Ancona, «Industrial odour pollution and human health: a systematic review and meta-analysis,» *Environmental Health*, vol. 20, 2021.
- [2] B. Ernest F. et F. R., «An economic evaluation of livestock odor regulation distances,» *Journal of Environmental Quality*, vol. 33, pp. 1997-2006, 2004.
- [3] MINER, J. , Nuisance concerns and odor control, vol. 80, 1997, pp. 2667-72.
- [4] S. MUSHUI, «ADVANCE ON THE STUDY OF ODOR ASSESSMENT AND TREATMENT.,» *Environmental Chemistry*, 2011.
- [5] S. Yongdong et A. Anna Lisa, «Local environmental governance innovation in China: staging 'triangular dialogues' for industrial air pollution control,» *Journal of Chinese Governance*, vol. 3, pp. 351-369, 2018.
- [6] T. ZARRA , V. BELGIORNO et V. NADDEO, «Environmental Odour Nuisance Assessment in Urbanized Area: Analysis and Comparison of Different and Integrated,» *Atmosphere*, vol. 12, p. 690, 2021.
- [7] M. San Emeterio de la PARTE, J.-F. Martínez-Ortega, V. Hernández Díaz et N. Lucas Martínez, «Big Data and precision agriculture: a novel spatio-temporal semantic IoT data management framework for improved interoperability,» *Journal of Big Data*, vol. 10, pp. 1-32, 2023.
- [8] S. Mario, B. I. et C. I. , «Urban IoT ontologies for sharing and electric mobility,» *Semantic Web*, vol. 14, pp. 617-638, 2021.
- [9] C. D. et C. L. , «Evolution of Electronic Noses from Research Objects to Engineered Environmental Odour Monitoring Systems: A Review of Standardization Approaches,» *Biosensors*, vol. 9, 2019.

- [10] S. QIU, P. Hou, J. Huang, W. Han et Z. Kang, «The Monitoring of Black-Odor River by Electronic Nose with Chemometrics for pH, COD, TN, and TP,» *Chemosensors*, 2021.
- [11] M. Cristina Di, M. Pulieri, D. Raho, A. N Muresan, T. Andrea, T. Jessica , E. Nestola, B. Alberto, M. Sabrina et R. Ilaria , «Assessing semantic interoperability in environmental sciences: variety of approaches and semantic artefacts,» *Scientific Data*, vol. 11, 2024.
- [12] S. M., T. D. , L. S. et B. Andrea , «Semantic and syntactic interoperability in online processing of big Earth observation data,» *International Journal of Digital Earth*, n° %111, pp. 95-112, 2017.
- [13] H. Jack , K. Garcia et R. S. , «Semantic Development and Integration of Standards for Adoption and Interoperability,» *COMPUTER*, vol. 25, pp. 0: 1-35, 2011.
- [14] S. E., F. G., K. B. Lee et G. E. , «A Methodology for Modeling Interoperability of Smart Sensors in Smart Grids,» *Ieee Transactions of on Smart Grid*, vol. 13, pp. 555-563, 2022.
- [15] A. F., H. M. et L. H., «Salmon, sensors, and translation: The agency of Big Data in environmental governance,» *Environment and Planning D Society and Space*, vol. 36, pp. 905-925, 2018.
- [16] V. Lars, P. Strömert, N. Matentzoglou, N. Karam, K. M., M. Prinz et R. Baum, «Suggestions for extending the FAIR Principles based on a linguistic perspective on semantic interoperability,» *Scientific Data*, vol. 12, 2024.
- [17] Z. T., B. V. et N. V. , «Environmental Odour Nuisance Assessment in Urbanized Area: Analysis and Comparison of Different and Integrated Approaches,» *Atmosphere*, vol. 12, p. 690, 2021.
- [18] R. E., D. Cora, G. B. , K. Saleh, T. Hauke , F. D. von Loewenich, T. Norbert, B. Claas , B. Martin , H. L. C. , P. D. Luis Alberto, B. Michael , I. Josef et T. Sylvia , «Towards interoperability in infection control: a standard data model for microbiology,» *Scientific Data*, vol. 10, 2023.
- [19] X. Xin, L. Hui et F. Chaoyang, «WMO: an ontology for the semantic enrichment of wetland monitoring data,» *International Journal of Digital Earth*, vol. 16, pp. 2189-2211, 2023.
- [20] T. M., D. Thi Thanh Le, L. Van Duy, H. N. , G. F. , N. Van Duy et B. F. , «Electronic noses based on metal oxide nanowires: A review,» *Nanotechnology Reviews*, vol. 11, pp. 897-925, 2022.
- [21] A.-S. Mehdi, P. M. , M. C., T. G., N. H. , X. W., B. Xu, Y.-D. Li et R. H. , «Silicon Nanowires for Gas Sensing: A Review,» *Nanomaterials*, vol. 10, 2020.
- [22] J. R., Ó. G., M. E. et G. Stefáns, «Flavor characterization of ripened cod roe by gas chromatography, sensory analysis, and electronic nose,» *Journal of Agricultural and Food Chemistry*, vol. 52 , pp. 6250-6, 2004.
- [23] A. Asgar, A. S. Mansol, A. Ashraf Khan, K. Muthoosamy et S. Y. , «Electronic nose as a tool for early detection of diseases and quality monitoring in fresh postharvest produce: A comprehensive review,» *Comprehensive Reviews in Food Science and Food Safety*, pp. 1541-4337, 2023.
- [24] G. S., F. J. Barba, J. Zhou, M. Wang et Z. Altintas, «Electronic Sensor Technologies in Monitoring Quality of Tea: A Review,» *Biosensors*, vol. 12, 2022.
- [25] Z. Zhai, Y. Liu, C. Li, D. Wang et H. Wu, «Electronic Noses: From Gas-Sensitive Components and Practical Applications to Data Processing,» *Sensors*, vol. 24, p. 4806, 24 July 2024.
- [26] R. Stefano, L. B. J., R. L. Dellacà et L. Capell, «Physical Confounding Factors Affecting Gas Sensors Response: A Review on Effects and Compensation Strategies for Electronic Nose Applications,» *Chemosensors*, 2023.
- [27] D. Nik, D. Drix, T. P. A. Warner, S. Rastogi, C. Della Casa, T. Ackels, A. T. Schaefer, A. van Schaik et M. Schmuker, «High-speed odor sensing using miniaturized electronic nose,» *Science Advances*, vol. 10, 2024.
- [28] W. C., S. K., C. Wang, Y. Yang,, F. Güder, K. T., C. V. et Z. M. , «Electronic nose for toxic gas detection based on photostimulated core-shell nanowires,» *Rsc Advances*, vol. 4, pp. 35084-35088, 2014.
- [29] R. A. et P. C., «Material Screening for Gas Sensing Using an Electronic Nose: Gas Sorption Thermodynamic and Kinetic Considerations,» *Acs Sensors*, 2021.
- [30] C. Hong, H. D. et Z. Jilin, «Gas Recognition in E-Nose System: A Review,» *Ieee Transactions of on Biomedical Circuits and Systems* , vol. 16, pp. 169-184, 2022.
- [31] F. Cong, H. Li, L. Li, H.-Y. Su, J. Tang, X. Bai et H. Liu, «Smart Electronic Nose Enabled by an All-Feature Olfactory Algorithm,» *Advanced Intelligent Systems*, vol. 4, 2022.
- [32] C. M. et O. M. , «Artificial Intelligence in Gas Sensing: A Review,» *Acs Sensors*, 2025.
- [33] F. S., C. Marzuoli, S. Pagliara, D. G. et L. Sangaletti, «Targeting biomarkers in the gas phase through a chemoresistive electronic nose based on graphene functionalized with metal phthalocyanines,» *Rsc Advances*, vol. 13, pp. 251-263, 2022.
- [34] H. Takeshi, A. Lin, V. C. Copa, L. P. Lopez, L. R. , B. L. , K. Y. , Y. Liu et S. A. , «An electronic nose using a single graphene FET and machine learning for water, methanol, and ethanol,» *Microsystems and Nanoengineering*, vol. 6, 2020.
- [35] L. M., S. J. , A. Essenwanger, J. Schepers et T. S. , «Why digital medicine depends on interoperability,» *Npj Digital Medicine*, vol. 2, 2019.
- [36] Z. K. et X. N., «Forming Interoperability Through Interorganizational Systems Standards,» *Journal of Management Information Systems*, vol. 30, pp. 269-298, 2014.
- [37] M. Lingtong, Q. Tian, X. Lu, J. An et D. H. , «An openEHR based approach to improve the semantic interoperability of clinical data registry,» *Bmc Medical Informatics and Decision Making*, vol. 18, 2018.
- [38] F. Y., Y. Yin, L. Xu, Y. Zeng et F. Wu, «IoT-Based Smart Rehabilitation System,» *Ieee Transactions of on Industrial Informatics*, vol. 10, pp. 1568-1577, 2014.
- [39] D. Quang-Tu, K. A. , W. Rahayu et K. Nguyen, «A Framework for IoT Streaming Data Indexing and Query Optimization,» *Ieee Sensors Journal*, vol. 22, pp. 14436-14447, 2022.
- [40] A. K., K. B. M. et P. K. S. V. , «Heterogeneous Interoperable Sensors Integrating Cognitive Knowledge for IOT-Based Cross-Domain Applications,» *Ieee Sensors Journal*, vol. 25, pp. 12119-12127, 2025.
- [41] F.-M. Daniel, B. J. , G.-A. J. et M. J. M. , «Towards dynamic and heterogeneous social IoT environments,» *Computing*, vol. 105, pp. 1141-1164, 2022.
- [42] T. Y., K. A., M. T. et C. D. , «Towards Interoperability of Entity-Based and Event-Based IoT Platforms: The Case of NGSI and EPCIS Standards,» *Ieee Access*, vol. 9, pp. 49868-49880, 2021.
- [43] J.-S. FU, Y. Liu,H., C. B. Bhargava et Z. Zhang, «Secure Data Storage and Searching for Industrial IoT by Integrating Fog Computing and Cloud Computing,» *Ieee Transactions of on Industrial Informatics*, vol. 14, pp. 4519-4528., 2018.
- [44] W. Qian et W. Zhuolin, «Blockchain-Enhanced IoT Sensor Data Management for Engineering Monitoring,» *Ieee Sensors Journal*, vol. 25, pp. 7363-7375, 2025.
- [45] A. Ali, K. MITRA et S. E. , «A blockchain-based SLA monitoring and compliance assessment for IoT ecosystems,» *Journal of Cloud Computing*, vol. 12, pp. 1-22, 2023.
- [46] Z. Junhua et G. Zhiqiang, «Collaborative Deep Learning and Information Fusion of Heterogeneous Latent Variable Models for Industrial Quality Prediction,» *Ieee Transactions of on Cybernetics*, vol. 55, pp. 1659-1672, 2025.
- [47] Z. Yuchen, Z. Liu, F. Yan, Y. C. , Y. Cao, Y. Bao, J. Mao, L. Siwei et H. Xiao, «Siamese Neural Network and Multimodal Data Fusion Approach for Small-Sample Learning in Industrial Soft Sensor Modeling,» *Ieee Sensors Journal*, vol. 24, pp. 33763-33777, 2024.
- [48] S.-S. TZENG , Y.-J. LIN et S.-W. WANG, «Age of Information in IoT Devices With Integrated Heterogeneous Sensors Under Slotted ALOHA,» *Ieee Sensors Journal*, vol. 25, pp. 20842-20853, 2025.

- [49] Z. Zaman , X. Wanli , G. Praveen , H. Wen , J. Jiaojiao et K. J. Sanjay , «Privacy-Preserving Probabilistic Data Encoding for IoT Data Analysis,» *Ieee Transactions of on Information Forensics and Security*, vol. 19, pp. 9173-9187, 2024.
- [50] L. MA, B. Duan, Z. Bo, L. Yang , F. Yingxun et M. Dongchao , «A trusted IoT data sharing method based on secure multi-party computation,» *Journal of Cloud Computing*, vol. 13, p. 138, 2024.
- [51] G. NASCIMENTO, L., G. S. F. Cristopher , R. O. et L. A. Andre L. , «Data Sampling Algorithm Based on Complexity-Entropy Plane for Smart Sensing Applications,» *Ieee Sensors Journal*, vol. 21, pp. 25831-25842, 2021.
- [52] W. SUN, Z. Zhuoteng, M. Fangyuan, W. Jingde et Cheng Ji , «Industrial Application of Data-Driven Process Monitoring with an Automatic Selection Strategy for Modeling Data,» *Processes*, 2023.
- [53] W. XU, P. Yanzhou, C. Xiu-Quan, D. Weiping et Q. Yuhua , «A Novel Dynamic Fusion Approach Using Information Entropy for Interval-Valued Ordered Datasets,» *Ieee Transactions of on Big Data*, vol. 9, pp. 845-859, 2023.
- [54] C. XU, C. Si-lu, K. Xiangjie, Z. Chi, Y. Guilin et F. Zaojun , «Robust Feature Selection by Removing Noise Entropy Within Mutual Information for Limited-Sample Industrial Data,» *Ieee Transactions of on Industrial Informatics*, vol. 21, pp. 3913-3923, 2025.
- [55] S. LI, L. XU et X. WANG, «Compressed Sensing Signal and Data Acquisition in Wireless Sensor Networks and Internet of Things,» *Ieee Transactions of on Industrial Informatics*, vol. 9, pp. 2177-2186, 2013.
- [56] N. Givanildo L., C. G. S. Freitas, R. O. et L. L. A. Andre , «Data Sampling Algorithm Based on Complexity-Entropy Plane for Smart Sensing Applications,» *Ieee Sensors Journal*, vol. 21, pp. 25831-25842, 2021.
- [57] H. LOUIE et A. MIGUEL, «Lossless Compression of Wind Plant Data,» *Ieee Transactions of on Sustainable Energy*, vol. 3, pp. 598-606, 2012.
- [58] K. J. BALAKRISHNAN et N. TOUBA, «Relationship Between Entropy and Test Data Compression,» *Ieee Transactions of on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, pp. 386-395, 2007.
- [59] A. KUMAR, H. Chung, T. Singhal, S. Yadav, G. Cody et G. Chang, «Industrial IOT Enabled SHM Systems,» *Structural Health Monitoring-An International Journal*, 2019.
- [60] A. Charu, M. Prashant Kumar, P. Pawan Kumar, S. B. Khan, A. Agarwal, T. Gadekallu et A. Malibari, «Preservation of Sensitive Data Using Multi-Level Blockchain-based Secured Framework for Edge Network Devices,» *Journal of Grid Computing*, vol. 21, 2023.
- [61] S. JAVAID, H. Afzal, M. Babar, F. Arif, Z. Simon Tan et M. Ahmad Jan, «ARCA-IoT: An Attack-Resilient Cloud-Assisted IoT System. *Ieee Access*,» vol. 7, pp. 19616-19630, 2019.
- [62] L. CASTRI, M. Sariah, H. Marc et B. Nicola , «CAnDOIT: Causal Discovery with Observational and Interventional Data from Time Series,» *Advanced Intelligent Systems*, vol. 6, 2024.
- [63] C. A. FURIA, R. TORKAR et R. FELDT, «Towards Causal Analysis of Empirical Software Engineering Data: The Impact of Programming Languages on Coding Competitions,» *Acm Transactions of on Software Engineering and Methodology*, vol. 33, pp. 1-35, 2023.
- [64] Z. LING, W. Jingxuan, Z. Yiwen, Z. Peng , W. Xingyu, Y. Kui et W. Xindong , «Label-Aware Causal Feature Selection,» *Ieee Transactions of on Knowledge and Data Engineering*, vol. 37, pp. 1268-1281, 2025.
- [65] M. WYLOT, P. Cudré-Mauroux, M. Hauswirth et G. Paul , «Storing, Tracking, and Querying Provenance in Linked Data,» *Ieee Transactions of on Knowledge and Data Engineering*, vol. 29, pp. 1751-1764, 2017.
- [66] R. WANG, W. S. Daniel, L. Guoqiang, R. Wong et C. Shipping , «Pipeline provenance for cloud-based big data analytics,» *Software Practice and Experience*, vol. 50, pp. 658-674, 2020.
- [67] L. Hanrui, D. KUMAR et N. EL-ATAB, «A Memristor-Based Event-Driven Reservoir Computing System for Cardiac Arrhythmia Detection,» *Advanced Intelligent Systems*, vol. 7, 2025.
- [68] C. HAGEDORN, J. HUEGLE et R. SCHLOSSER, «Understanding unforeseen production downtimes in manufacturing processes using log data-driven causal reasoning,» *Journal of Intelligent Manufacturing*, vol. 33, pp. 2027-2043, 2022.
- [69] X. XIAO, K. Alharbi, P. Zhang, Q. Hantang et Y. Xubo , «Bayesian federated causal inference and its application in manufacturing,» *Journal of Intelligent Manufacturing*, 2025.
- [70] Y. GAO, X.-y. CHEN et X. DU , «A Big Data Provenance Model for Data Security Supervision Based on PROV-DM Model,» *Ieee Access* , vol. 8, pp. 38742-38752, 2020.
- [71] J. FISCHER, L. Egli, J. Groth, B. Caterina, S. Ehrmann, H. Figgemeier, C. Henzen, C. Meyer, R. Müller-Pfefferkorn, A. Rümmler, M. Wagner, L. Bernard et R. Seppelt, «Approaches and tools for user-driven provenance and data quality information in spatial data infrastructures,» *International Journal of Digital Earth*, vol. 16, pp. 1510-1529, 2023.
- [72] F. MONTORI, L. Kewen, D. G. Matteo, P. Jayaraman, L. Bononi, T. Sellis et D. Georgakopoulos, «A Metadata-Assisted Cascading Ensemble Classification Framework for Automatic Annotation of Open IoT Data,» *Ieee Internet of Things Journal*, vol. 10, pp. 13401-13413, 2023.
- [73] N. ELKIN-KOREN et M. GAL, «The Chilling Effect of Governance-by-Data on Data Markets,» *University of Chicago Law Review*, vol. 6, p. 86, 2019.
- [74] J. d. D. NAMBAJIMANA, H. Xiu-bin, J. Zhou, M. F. Justine, L. Jin-lin , D. Khurram, R. Mind'je et G. Nsabimana, «Land Use Change Impacts on Water Erosion in Rwanda,» *Sustainability*, 2019.
- [75] J. NSANZABAGANWA, X. Chen , L. Tie , H. Egide , M. Richard , G. Aboubakar, F. Bakayisire , U. Adeline et S. Niyonsenga , «Impacts of Land Use and Land Cover Change on Non-Point Source Pollution in the Nyabarongo River Catchment,» *Rwanda. Water*, 2024.
- [76] K. MASIKU et E. NTAGWIRUMUGARA, «An impact analysis of population growth on energy-water-food-land nexus for ecological sustainable development in Rwanda,» *Food and Energy Security*, vol. 9, 2020.
- [77] C. MUGUNGA, D. Kool, M. Wijk, G. Mohren et K. Giller, «Water use by short rotation Eucalyptus woodlots in southern Rwanda,» *Agroforestry Systems*, vol. 89, pp. 1119-1139, 2015.
- [78] H. L. DEWITT, J. Gasore, M. Rupakheti, . E. P. Katherine, R. Prinn, J. d. D. Ndikubwimana, N. Julius et S. Bonfils , «Seasonal and diurnal variability in O3, black carbon, and CO measured at the Rwanda Climate Observatory,» *Atmospheric Chemistry and Physics*, 2019.
- [79] I. DUFATANYE, L. Youngjae, K. Hyungsoo et L. Sangyoun , «Industrial Wastewater Discharge and Compliance Investigation for Environmentally Resilient Rwanda,» *Water*, 2022.
- [80] T. H., Z. N. Pintarić, Y. Yang et Z. Kravanja, «Multilevel strategies for the retrofit of large-scale industrial water system: A brewery case study,» *Aiche Journal*, vol. 58, pp. 884-898, 2012.
- [81] Y. Gyeongbeom et R. G. , «Optimal Design of Batch-Storage Network with Recycle Streams. *Aiche Journal*, 2003, 49: 3084-3094».
- [82] S. N., S. C. et L. S., «Feasibility study for biohydrogen production from raw brewery wastewater,» *International Journal of Energy Research*, vol. 39, pp. 1769-1777, 2015.
- [83] K. P. et K. J. , «Coal Demand and Environmental Regulations: A Case Study of the Polish Power Sector,» *Energies*, 2020.
- [84] K. P. et K. J. , «Coal Demand and Environmental Regulations: A Case Study of the Polish Power Sector,» *Energies*, 2020.
- [85] N. A. KIMANI, «Collaborative Approach to Environmental Governance in East Africa,» *Journal of Environmental Law*, vol. 22, pp. 27-57, 2010.
- [86] W. ONZIVU, «Rethinking Transnational Environmental Health Governance in Africa: Can Adaptive Governance Help?,» *Review of European*

- Comparative and International Environmental Law*, vol. 25, pp. 107-122, 2016.
- [87] B. Riccardo, A. Kobeissi, F. Bellotti et A. D. Gloria, «Atmosphere, an Open Source Measurement-Oriented Data Framework for IoT,» *Ieee Transactions of on Industrial Informatics*, vol. 17, pp. 1927-1936, 2021.
- [88] B. A., S. Schmid, C. Schindhelm, A. Khelil, S. Käbisch, D. Kramer, D. Le-Phuoc, J. Mitic, D. Anicic et E. Teniente, «Enabling IoT Ecosystems through Platform Interoperability,» *Ieee Software*, vol. 34, pp. 54-61, 2017.
- [89] C. Zhipeng et S. Tuo, «Distributed Query Processing in the Edge-Assisted IoT Data Monitoring System,» *Ieee Internet of Things Journal*, vol. 8, pp. 12679-12693, 2021.
- [90] W. Wendong, C. Feng, B. Zhang et H. Gao, «Environmental Monitoring Based on Fog Computing Paradigm and Internet of Things,» *Ieee Access*, vol. 7, pp. 127154-127165, 2019.
- [91] T. Athanasios, A. Papamichail, I. Angelis, G. Koufoudakis, G. Tsoumanis et K. Oikonomou, «An Alertness-Adjustable Cloud/Fog IoT Solution for Timely Environmental Monitoring Based on Wildfire Risk Forecasting,» *Energies*, 2020.
- [92] S. Chayan, S. Nambi, R. Venkatesha Prasad, A. Biswas, R. Neisse et G. Baldini, «DIAT: A Scalable Distributed Architecture for IoT,» *Ieee Internet of Things Journal*, vol. 2, pp. 230-239, 2015.
- [93] P. S., S. Middleton, F. Chaves, R. Tao, Ö. Necmioğlu et U. Bügel, «A Semantic IoT Early Warning System for Natural Environment Crisis Management,» *Ieee Transactions of on Emerging Topics in Computing*, vol. 3, pp. 246-257, 2015.
- [94] J.-G. R., G. A. , C. Agostinho, F. Lampathaki et Y. Charalabidis, «Systematisation of Interoperability Body of Knowledge: the foundation for Enterprise Interoperability as a science,» *Enterprise Information Systems*, vol. 7, pp. 32-7, 2013.
- [95] B. M. R., A. Gill, G. Beydoun et B. Mccusker, «Big Data Management and Analytics Metamodel for IoT-Enabled Smart Buildings,» *Ieee Access*, vol. 8, pp. 169740-169758, 2020.
- [96] K. Mathias, L. Moestue, A. Obermeier et T. Widmann, «Assessing Completeness of IoT Data: A Novel Probabilistic Approach,» *Business and Information Systems Engineering*, vol. 67, pp. 797-814, 2024.
- [97] Z. Minghu et L. Xin, «Drone-Enabled Internet-of-Things Relay for Environmental Monitoring in Remote Areas Without Public Networks,» *Ieee Internet of Things Journal*, vol. 7, pp. 7648-7662, 2020.
- [98] A.-F. Ashraf E., F. Al-turjman, W. Alsalih et H. Hassanein, «A Priced Public Sensing Framework for Heterogeneous IoT Architectures,» *Ieee Transactions of on Emerging Topics in Computing*, vol. 1, pp. 133-147, 2013.
- [99] S. B., H. Vesal et A. A. , «A trust management scheme for IoT-enabled environmental health/accessibility monitoring services,» *International Journal of Information Security*, vol. 19, pp. 93-110, 2018.
- [100] M. Sachit , «Design and development of an open-source framework for citizen-centric environmental monitoring and data analysis,» *Scientific Reports*, vol. 12, 2022.
- [101] L. Hui, J. Hu, X. Wang, M. F. Alhamid et M. Jalil Piran, «Toward Secure Data Fusion in Industrial IoT Using Transfer Learning,» *Ieee Transactions of on Industrial Informatics*, vol. 17, pp. 7114-7122, 2020.
- [102] J. Sabeen, H. Afzal, M. Babar, F. Arif, Z. Simon Tan et M. Ahmad Jan, «ARCA-IoT: An Attack-Resilient Cloud-Assisted IoT System,» *Ieee Access*, vol. 7, pp. 19616-19630, 2019.
- [103] S. Shaila, S. Huda, J. Abawajy, W. N. Ismail et M. M. Hassan, «Malware Threats and Detection for Industrial Mobile-IoT Networks,» *Ieee Access*, vol. 6, pp. 15941-15957, 2018.
- [104] G. Jing, Y. Sun, R. Rameezdeen et C. Chow, «Understanding data governance requirements in IoT adoption for smart ports – a gap analysis,» *Maritime Policy and Management*, vol. 51, pp. 617-630, 2022.
- [105] B. M. R., G. A. , B. G. et M. Brad , «Big Data Management and Analytics Metamodel for IoT-Enabled Smart Buildings,» *Ieee Access*, vol. 8, pp. 169740-169758, 2020.
- [106] H. Daniel et H. A. , «Interoperability in the era of digital innovation: An information systems research agenda,» *Journal of Information Technology*, vol. 37, pp. 407-427, 2022.
- [107] M.HOLUB, R. Brandenburg, H. Grosch, S. Weinmann et B. Hansel, «Plasma Supported Odour Removal from Waste Air in Water Treatment Plants: An Industrial Case Study,» *Aerosol and Air Quality Research*, vol. 14, pp. 697-707, 2014.
- [108] B. J. , B. Nicolas D., B. A. C. , C. A., D. F. , K. Dewan, H. T. , J. M. , K. R. , L. G., M. M. , P. A. , S.-A. M. , S. C. et H. M. , «The role of digital data entry in participatory environmental monitoring,» *Conservation Biology*, vol. 30, 2016.
- [109] J. Cassiani, M. Dugarte, C. Arteta et S. Kessler, «Durability assessment of a tunnel structure with two-sided chloride ingress—A case study located in a tropical environment,» *Structural Concrete*, vol. 23, pp. 739-752, 2021.
- [110] C. S. et T. K., «Assessment of climate-based daylight performance in tropical office buildings: a case study,» *International Journal of Low-Carbon Technologies*, vol. 9, pp. 100-108, 2014.