The Compliance Routing Problem — A Practitioner-Built Ontology for Multi-Agency EHS Navigation

Adam J. Bick

2026-04-09

Abstract

Environmental Health and Safety (EHS) practitioners perform a complex intellectual task that no existing ontology has formalized: routing a single workplace event through the overlapping jurisdictions of multiple independent federal agencies, where hazard type, operational context, and contextual conditions determine which regulatory frameworks activate and which obligations apply. This paper presents a formal OWL ontology for EHS compliance routing that captures this logic as a deterministic system. The ontology is structured around the Employee-Hazard bipolar model, mediated by the 5 E's of Safety and assessed through the ARECC decision-making framework. Its central contribution is the treatment of hazard type classification as a compliance activation mechanism rather than a descriptive label. A three-axis routing model (HazardType × ActionContext × ContextualCondition) produces deterministic regulatory obligation sets from classified inputs. CompoundHazardProfiles formalize the additive nature of compliance when multiple hazard types co-occur. Four regulatory modules extend the core model into EPCRA chemical inventory and TRI reporting, Clean Air Act Title V permitting, OSHA 300 recordkeeping, and incident management, wired together through an Establishment class that anchors facility-level obligations. Nine worked scenarios validate the routing logic across chemical, biological, electrical, mechanical, physical, ergonomic, and psychosocial hazard types. A designed Geo-Compliance Extension adds a fourth routing axis, FacilityJurisdiction, that layers state, county, and municipal regulatory obligations on top of the federal baseline without modifying the core ontology. The ontology provides a formal, transferable structure for multi-agency EHS compliance knowledge that currently exists only in the professional judgment of experienced practitioners.

1 Introduction

The traditional mandate of an Environment, Health, and Safety (EHS) professional is often distilled into two core functions: ensuring worker safety by enforcing applicable industry regulations, and maintaining corporate compliance by managing standard frameworks and reporting to the relevant authorities. However, this sterile definition collapses upon contact with the physical reality of the job — specifically, the loading dock.

Consider a routine scenario where a worker accidentally punctures a 55-gallon drum of solvent during offloading. In an instant, that single physical event fractures into a complex, multi-agency compliance cascade. The Occupational Safety and Health Administration (OSHA) dictates the worker's emergency response, exposure limits, and required personal protective equipment. Simultaneously, the Environmental Protection Agency (EPA) governs the hazardous waste containment, spill threshold reporting, and environmental impact. Meanwhile, the Department of Transportation (DOT) regulates the now-compromised shipping container and its subsequent movement. None of these regulatory agencies coordinate their requirements or share a unified compliance framework.

Consequently, the EHS practitioner is forced to act as the unmodeled integration layer between these disparate, siloed entities — the manual routing engine for the organization, tasked with deciphering how a single physical hazard triggers compounding, overlapping, and occasionally contradictory regulatory pathways. Despite the critical nature of this multi-agency routing process, the domain knowledge required to execute it remains largely informal and undocumented.

This paper proposes that EHS compliance routing possesses a distinct, formal architectural structure that has not yet been adequately captured or formalized in the existing literature. It presents a formal ontology for EHS compliance routing, grounded in established industrial hygiene frameworks and validated against real-world multi-agency scenarios. The model captures the structural relationships between hazard classification, operational context, and regulatory activation — relationships that practitioners navigate intuitively but that have never been structured for computational logic.

2 Literature Review

Ontological approaches to occupational safety and health have developed along two independent tracks. The first models workplace safety as a knowledge domain, formalizing hazard types, incidents, and risk relationships. The second models regulatory compliance as a rule-checking problem, encoding specific legal texts into machine-readable formats. Neither track has addressed the central challenge facing EHS practitioners: routing a single workplace event through the overlapping jurisdictions of multiple independent regulatory agencies.

2.1 Safety Domain Ontologies

The most structurally mature attempt at a core occupational safety ontology is OSHDO-Core, proposed by (Ławrynowicz and Ławniczak 2016). Built on OHSAS 18001 (now superseded by ISO 45001), OSHDO-Core defines classes for OccupationalHazard, HazardousEvent, Risk, Worker, Workplace, OccupationalExposure, ProtectiveMeasure, and SafetyDataSheet. It is a vocabulary ontology: it names the entities and relationships within occupational safety but does not model what an organization is required to do about them. No regulatory framework appears in the model. No activation logic connects a hazard classification to a compliance obligation. OSHDO-Core maps the terrain without mapping the rules that govern it.

(Onut Badea et al. 2024) take a different approach with HSM-Onto, an ontology for occupational safety management tailored to workers with disabilities. HSM-Onto organizes its knowledge around HazardousSituation as a central class, branching into Types, Consequences, RiskManagement, Incident, OccupationalDisease, and Source/WorkProcess. It automates job hazard analysis and interfaces with real-time Internet of Things (IoT) monitoring to personalize mitigation strategies based on individual worker vulnerabilities. This work demonstrates an important insight: the worker is not a generic variable but a specific individual with physiological and cognitive constraints that shape what constitutes adequate protection. However, HSM-Onto models safety as a domain of knowledge, not as a domain of legal obligation. No regulatory agency appears in the ontology. The question of which external framework governs a given hazard, and which agency enforces it, lies entirely outside its scope.

(Sanchez-Pi et al. 2015) build an ontology for the offshore oil industry that fuses multi-sensor data with spatiotemporal constraints to predict accident risks. Their system tracks anomaly types (distinguishing "accident with injury" from "incident") and uses association rule mining to identify patterns that precede failures. The ontology serves as a backend data model for contextual reasoning about past incidents within a single industrial domain. It does not classify hazards by type in a way that connects to regulatory frameworks, does not model compliance obligations, and is not generalizable beyond the Petrobras operational environment for which it was designed.

(Single et al. 2020) digitize the Hazard and Operability (HAZOP) method using formal ontologies to model equipment, processes, and deviations in process safety engineering. Their system automates risk assessment by using semantic reasoners to identify causes, consequences, and safeguards for specific equipment configurations. This represents a sophisticated application of ontological reasoning to one component of EHS practice (engineering controls). It stops at the equipment level. The regulatory context that determines which standards apply, which agencies have jurisdiction, and which obligations activate based on hazard type is absent from the model.

Across this body of work, a consistent pattern emerges. Safety domain ontologies model what can go wrong and what might prevent it, but they do not model who requires what response. The regulatory dimension of occupational safety, the layer where classification becomes obligation, remains unformalized.

2.2 Hazard Classification Ontologies

A parallel stream of work has focused on formalizing hazard classification itself, particularly for chemical hazards. This work provides the taxonomic foundation for hazard identification but stops precisely where compliance routing begins.

(Pascazio et al. 2023) develop OntoSpecies, a Web Ontology Language (OWL) ontology that models chemical species with their physical properties, identifiers, and Globally Harmonized System (GHS) hazard classifications. OntoSpecies links chemical data to The World Avatar knowledge graph via SPARQL Protocol and RDF Query Language (SPARQL) endpoints, enabling complex queries such as solvent selection based on multiple chemical properties. The ontology classifies chemicals according to GHS criteria with precision. It does not model what those classifications require. A chemical classified as GHS Category 1 Flammable Liquid receives its hazard class, but the ontology contains no mechanism to route that classification to OSHA flammable storage requirements under 29 CFR 1910.106, DOT shipping obligations under 49 CFR 173, or EPA disposal requirements under 40 CFR 262.

(Medici et al. 2022) construct PestOn, an RDF/OWL ontology for pesticide information that tracks active ingredients, roles, hazard profiles, and regulatory authorization statuses. PestOn models the GHS classification of pesticide hazards exhaustively. Like OntoSpecies, it treats classification as the terminal output. The question of how a GHS-classified hazard activates obligations across multiple independent agencies does not arise within the model.

At the federal infrastructure level, (Mortensen et al. 2024) document the US National Nanotechnology Initiative consortium's effort to standardize nanoEHS data using Findable, Accessible, Interoperable, and Reusable (FAIR) principles, Semantic Web technologies, and tools like the EPA's OntoSearcher. This work maps nanomaterial stressors to Molecular Initiating Events and Adverse Outcome Pathways, building a shared data infrastructure across federal agencies. The consortium demonstrates that US federal agencies are actively adopting Resource Description Framework (RDF) and OWL for hazard data standardization. Critically, this effort standardizes data sharing between agencies without modeling the compliance routing logic that practitioners must execute when those agencies' requirements overlap. The infrastructure stops at interoperability. The practitioner's problem of determining which obligations activate for a given hazard in a given context remains unaddressed.

The GHS framework itself [United Nations (2023); OSHA 29 CFR 1910.1200] represents the most widely adopted hazard classification standard, defining pictograms, signal words, hazard statements, and Safety Data Sheet formats. GHS standardizes how hazards are communicated. It does not standardize how those hazards route through regulatory systems. A GHS-classified flammable liquid simultaneously triggers obligations under OSHA (storage and handling), DOT (transport), and EPA (disposal), but the GHS framework contains no model of this multi-agency activation. The world has standardized hazard classification. It has not standardized hazard routing.

2.3 Compliance Ontologies and Automated Rule Checking

A third body of work directly addresses regulatory compliance through ontological modeling, but each effort operates within a single regulatory domain or jurisdiction without addressing the multi-agency intersection problem.

(Park and Shin 2025) construct a domain ontology for fall hazards in construction that maps OSHA safety requirements (29 CFR 1926.500) into machine-readable Semantic Web Rule Language (SWRL) rules linked to Building Information Modeling (BIM) elements. Their architecture separates the physical environment (Construction Class), the operational context (Conduct Class), and the regulatory logic (Safety Control Class). This tripartite separation correctly identifies that compliance requires linking physical conditions to operational activities to regulatory rules. However, the Safety Control Class is hardcoded to a single agency (OSHA) and a single hazard type (falls). If a roofing task involves asbestos removal, triggering EPA requirements under 40 CFR 61 and DOT requirements under 49 CFR 171 alongside OSHA fall protection, the model has no architecture to represent this overlap. Park and Shin demonstrate the structural pattern that compliance ontologies require. They do not solve the multi-agency routing problem.

(Dimyadi et al. 2016) propose a compliance auditing architecture for Building Information Modeling that separates the Building Compliance Model (physical BIM data) from the Regulatory Knowledge Model (legal rules) and bridges them through Compliant Design Procedure workflows. Their approach uses Legal Rule Markup Language (LegalRuleML) and executable Business Process Model and Notation (BPMN) workflows to manually encode fire safety regulations into machine-readable rules checked against BIM objects. Zhang et al. (2026) advance this foundation by automating the rule generation process itself, combining Bidirectional Encoder Representations from Transformers (BERT) semantic extraction with context-free grammar to translate Chinese building codes into executable Shapes Constraint Language (SHACL) constraints through a human-in-the-loop workflow, achieving 95.8% translation accuracy. Together, these papers track a decade of progress in transforming regulatory text into machine-readable compliance logic and prove that regulatory language possesses a capturable formal structure. Both efforts, however, operate within a single regulatory domain (building and fire codes) and evaluate static physical BIM elements rather than dynamic worker-hazard interactions. Neither models multi-agency jurisdiction or the compound hazard scenarios that arise when a single workplace event triggers obligations across independent agencies.

(Oyasiji et al. 2024) propose an AI architecture combining natural language processing (NLP), machine learning, and legal knowledge graphs (Legal Knowledge Interchange Format Core [LKIF-Core], LegalRuleML) to automate compliance checking for the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and anti-money-laundering regulations. This work acknowledges the existence of multi-jurisdictional compliance (GDPR versus the California Consumer Privacy Act [CCPA], for instance) and demonstrates the efficacy of combining NLP with ontological mapping for regulatory texts. However, it addresses information privacy and financial regulations, not physical workplace hazards. The compliance routing problem in EHS differs fundamentally: it involves physical hazard types that simultaneously activate independent agencies governing different aspects of the same material or event.

Vigneron et al. (2013) build an ontology for occupational health and safety legal compliance knowledge management within the French regulatory system, the most direct European parallel to the compliance routing problem.

(Tan et al. 2012) formalize Process Safety Management (PSM) under 29 CFR 1910.119 as an OWL ontology in Protégé, representing the 14 PSM elements as a knowledge base for automating audit processes. This represents a single-standard ontology. It models one OSHA regulation without touching EPA, DOT, or the broader multi-agency ecosystem.

(Gallina et al. 2025) builds an ontology for EU Machinery Directive compliance, focusing on multi-concern validation (safety and cybersecurity) for machinery legislation. This work demonstrates that compliance ontology research is active but confirms the pattern: each effort addresses a single regulatory domain without modeling the intersection problem.

2.4 Commercial Systems

Commercial EHS software and industrial IoT platforms have approached the problem from the implementation side. The IndustryOS/ElixirClaw platform (Gill 2026) uses knowledge graphs built on the ISA-95 standard and AI agents to govern manufacturing environments, orchestrating safety workflows such as personal protective equipment (PPE) checks and fire detection through contextual logic and policy-as-code. These systems model internal policy enforcement and physical hazard detection effectively. They do not model external regulatory compliance routing. The platform can detect that a fire has occurred and trigger an internal response protocol. It cannot determine that the same event activates OSHA reporting under 29 CFR 1904, EPA release notification under 40 CFR 302, and DOT emergency response under 49 CFR 172. Commercial systems treat compliance as a checklist layered on top of operational workflows rather than as an interconnected system with its own formal structure.

2.5 Identified Gap in the Literature

The literature reviewed above converges on a consistent structural absence. Safety domain ontologies model workplace hazards without modeling regulatory obligations. Hazard classification ontologies formalize what a substance or condition is without formalizing what that classification requires. Compliance ontologies encode specific regulations without addressing how multiple independent agencies simultaneously govern the same physical event. Commercial systems detect hazards and enforce internal policies without modeling the external regulatory logic that determines legal obligations.

No existing ontology models the US multi-agency EHS compliance routing problem: the intersection where OSHA, EPA, DOT, the National Fire Protection Association (NFPA), ISO, and state agencies simultaneously govern the same workplace event, and where hazard type, operational context, and contextual conditions determine which frameworks activate and which obligations apply. No existing work formalizes EHS from the practitioner's perspective, where the core intellectual task is not hazard identification or single-regulation compliance but the mediation between a classified hazard and the full set of overlapping regulatory obligations it triggers.

This paper addresses that gap. The following sections present an ontology that models compliance routing as a formal system, using hazard type as a compliance activation classifier, operational context as a routing parameter, and contextual conditions as obligation modifiers to determine the complete set of multi-agency requirements triggered by a given workplace event.

3 Domain Analysis

The literature reviewed in Section 2 reveals a field that has formalized fragments of occupational safety without formalizing the domain's underlying architecture. Before introducing the ontology, this section maps that architecture as it exists in professional EHS practice. The structures described here are not inventions of this paper; they are established frameworks taught in EHS education, codified in professional standards, and executed daily by practitioners. The contribution of this paper is the recognition that these frameworks compose into a formal system with a distinct computational structure.

3.1 The Bipolar Model: Employee and Hazard

EHS practice is organized around two poles. The first pole is the Employee: a specific individual with a job role, training history, health status, and physiological characteristics that shape their vulnerability to workplace exposures. The second pole is the Hazard: a source or situation with potential for harm, characterized by its physical nature, its toxicological properties, and the conditions under which it produces injury or illness.

The entire domain of EHS operates in the space between these two poles. Every safety program, every regulatory standard, every compliance obligation exists to govern the relationship between a worker and a hazard. The Occupational Safety and Health Act of 1970 itself encodes this bipolarity: it requires employers to furnish "employment and a place of employment which are free from recognized hazards" that could cause death or serious harm to "employees" (29 USC 654, Section 5(a)(1)). The statute names both poles and locates the employer's obligation between them.

The point where these poles meet is the Hazardous Exposure Situation: the operational moment when an employee encounters a hazard. This concept, adapted from (Onut Badea et al. 2024), is the central event of the domain. Every risk assessment evaluates it. Every control measure attempts to prevent or mitigate it. Every incident investigation reconstructs it.

The Employee pole is not a generic placeholder. As (Onut Badea et al. 2024) demonstrate with HSM-Onto, individual worker characteristics shape what constitutes adequate protection. A worker's job role determines which hazards they encounter. Their training status determines whether they can recognize and respond to those hazards. Their health status, including pre-existing conditions, age, and concurrent exposures, determines their susceptibility. EHS practitioners group workers with similar exposure profiles into Similar Exposure Groups (SEGs) to manage monitoring and protection strategies efficiently, but the underlying reality is that the Employee pole carries individual variation that the Hazard pole does not.

3.2 The Seven Industrial Hygiene Concepts: Describing the Hazard Pole

The Hazard pole is not a simple label. It is described through seven foundational concepts from industrial hygiene (IH) that provide the scientific vocabulary for understanding how hazards produce harm:

Toxicity is the inherent capacity of a substance to cause adverse biological effects, quantified by measures such as LD50 and LC50. The principle attributed to Paracelsus applies: the dose makes the poison. A substance is not inherently "safe" or "dangerous" in isolation.

Exposure is the contact between a worker and a hazardous agent, characterized by concentration, duration, and frequency. Exposure is the bridge between a hazard existing in the environment and a hazard causing harm to a person. Without exposure, toxicity alone does not produce risk.

The Dose-Response Relationship is the quantitative relationship between the amount of agent absorbed and the magnitude of the biological effect. This is the cornerstone of toxicology and the scientific basis for occupational exposure limits.

Routes of Entry describe the pathways by which a hazardous agent reaches the body: inhalation (the most common occupational pathway), dermal absorption, ingestion, and injection. Each route determines the appropriate monitoring strategy and control measures.

Occupational Exposure Limits (OELs) are regulatory and recommended concentration limits for hazardous agents in the workplace atmosphere. These include OSHA Permissible Exposure Limits (PELs), American Conference of Governmental Industrial Hygienists (ACGIH) Threshold Limit Values (TLVs), and National Institute for Occupational Safety and Health (NIOSH) Recommended Exposure Limits (RELs). OELs are derived from dose-response data with safety factors applied.

The Time-Concentration Relationship captures the interplay between exposure duration and agent concentration that determines total dose, expressed as Time-Weighted Average (TWA), Short-Term Exposure Limit (STEL), and Ceiling values. A critical insight of this concept is that low-dose, long-term exposures can be more dangerous than short, high-dose events.

Individual Susceptibility accounts for variation in individual response to hazardous exposures based on genetics, pre-existing health conditions, age, sex, nutritional status, and concurrent exposures. This concept explains why OELs cannot protect all workers equally and connects directly back to the Employee pole's individual characteristics.

These seven concepts are not a taxonomy of hazard types. They are the lenses through which any hazard is understood, measured, and managed. A chemical hazard is described by its toxicity, the routes through which workers are exposed, the dose-response curve that determines what exposure levels produce harm, and the OELs that translate that science into enforceable limits. The seven IH concepts provide the descriptive framework for the Hazard pole regardless of hazard type.

3.3 Hazard Type as a Classification System

Distinct from the IH concepts that describe a hazard's properties is the classification of hazards by type. The EHS domain recognizes seven primary hazard types: Physical (noise, vibration, radiation, temperature extremes), Mechanical (moving parts, struck-by and caught-in scenarios), Chemical (toxic substances, flammable liquids, corrosives), Biological (bloodborne pathogens, mold, infectious agents), Psychosocial (workplace stress, violence, harassment), Ergonomic (repetitive motion, awkward postures, heavy lifting), and Electrical (shock, arc flash, electrocution).

This taxonomy, adapted from (Onut Badea et al. 2024), is typically treated in the literature as a descriptive label. A hazard is classified as "chemical" or "ergonomic," and the classification serves organizational and communication purposes. What the literature has not recognized is that hazard type classification is functionally a compliance activation mechanism. Each hazard type activates a specific set of regulatory frameworks. A Chemical hazard activates OSHA Hazard Communication (29 CFR 1910.1200), EPA Resource Conservation and Recovery Act (RCRA), GHS classification requirements, and ACGIH TLV standards. An Electrical hazard activates OSHA electrical standards (29 CFR 1910 Subpart S), NFPA 70E, and NFPA

70 (National Electrical Code). The classification is not merely descriptive. It is a routing mechanism that determines which bodies of regulation apply.

This insight has a critical corollary: workplace hazards rarely present as pure single-type events. An employee manually lifting and tilting a 55-gallon drum of solvent to pour into a machine confronts a Chemical hazard (inhalation of vapors, dermal contact) and an Ergonomic hazard (heavy lifting, awkward posture) simultaneously. Both compliance pathways activate. OSHA Hazard Communication applies to the chemical exposure. The NIOSH Lifting Equation and OSHA General Duty Clause apply to the ergonomic risk. Neither cancels the other. Compliance is additive: when multiple hazard types co-occur, the resulting regulatory obligation is the union of all constituent type activations. This is what practitioners manage daily, and no existing ontology captures it.

3.4 The 5 E's: The Mediation Layer

Between the Employee pole and the Hazard pole sits a mediation layer that practitioners call the 5 E's of Safety. These are not hazard descriptions or regulatory citations. They are the five categories of intervention process through which organizations manage the Employee-Hazard relationship:

Education provides workers with knowledge of hazards, safe work practices, and risk management skills. It includes role-specific instruction, hands-on demonstrations, hazard recognition exercises, and scenario-based learning.

Encouragement motivates and reinforces safe behaviors through positive recognition, incentive programs, peer-to-peer safety coaching, and leadership engagement in safety activities. It builds the cultural foundation that sustains safety performance.

Engineering encompasses the design, development, and maintenance of physical controls that eliminate or reduce hazards at the source: machine guarding, ventilation systems, ergonomic workstation design, automated shutoffs, and containment systems.

Enforcement ensures consistent adherence to safety rules through documented disciplinary procedures, supervisor accountability, and corrective action processes. Enforcement is not purely punitive. It establishes accountability and operates in conjunction with education and encouragement.

Evaluation provides continuous measurement and improvement of safety systems using both lagging indicators (Total Recordable Incident Rate, Days Away/Restricted/Transferred rate) and leading indicators (near-miss frequency, audit completion rates, corrective action closure times, training compliance metrics).

The 5 E's mediate between the poles in both directions. They protect the Employee from the Hazard (engineering controls reduce exposure; education enables hazard recognition) and they mitigate the Hazard's potential to cause harm (evaluation identifies emerging risks before they produce incidents; enforcement maintains the integrity of protective systems). Every safety program, every regulatory compliance activity, and every corrective action maps to one or more of these five processes.

3.5 The ARECC Framework: Structuring the Decision Process

While the 5 E's describe the intervention categories, practitioners need a decision-making framework to structure how they move through a hazardous situation. The American Industrial Hygiene Association (AIHA) provides this with the ARECC framework: Anticipate, Recognize, Evaluate, Control, and Confirm.

Anticipate is the proactive identification of potential hazards during design, planning, or before the introduction of new processes, chemicals, or equipment. Added to the IH framework in 1994 by AIHA president Harry Ettinger, this step shifts the discipline from reactive hazard management to proactive risk prevention.

Recognize is the identification of existing hazards through inspection, inventory review, Safety Data Sheet (SDS) analysis, workplace observation, employee interviews, and review of incident and sampling data.

Evaluate is the qualitative and quantitative assessment of exposures through personal air sampling, direct-reading instruments, and biological monitoring, compared against OELs. This step synthesizes hazard assessment, exposure assessment, and population assessment to determine whether the current exposure level poses unacceptable risk.

Control is the implementation of the Hierarchy of Controls to reduce or eliminate exposures. The Hierarchy ranks interventions by effectiveness: elimination (removing the hazard entirely), substitution (replacing it with a less hazardous alternative), engineering controls (physical isolation), administrative controls (work practice changes, scheduling, rotation), and personal protective equipment (PPE, the last line of defense and least effective because it depends on consistent worker usage).

Confirm, added in 2011, verifies that controls are effective and desired outcomes are achieved. This step closes the loop. An incident investigation is fundamentally a Confirm-phase activity: it confirms whether the controls identified during earlier ARECC phases were adequate. Investigation findings feed back into the cycle as new hazards are anticipated, failure modes are recognized, residual risks are evaluated, and improved controls are implemented.

ARECC is not a checklist. It is a recursive process. Every significant workplace change, every new chemical introduced, every incident investigated cycles through ARECC. The framework provides the procedural architecture through which practitioners apply the 5 E's to manage the Employee-Hazard relationship.

3.6 Synthesis: The Integrated Architecture

These frameworks — the bipolar model, the seven IH concepts, the hazard type taxonomy, the 5 E's, ARECC, and the Hierarchy of Controls — are not isolated tools. They compose into an integrated architecture. The Employee and Hazard poles define the domain's scope. The seven IH concepts provide the scientific vocabulary for characterizing the Hazard pole. The hazard type taxonomy classifies hazards in a way that activates specific regulatory frameworks. The 5 E's supply the intervention categories that mediate between the poles. ARECC provides the decision process through which practitioners apply those interventions. The Hierarchy of Controls ranks intervention effectiveness within the Control phase.

This architecture is implicit in professional practice. Practitioners learn these frameworks separately and integrate them through experience. No existing formalization captures how they compose. When a senior EHS practitioner retires, the integrated understanding of how hazard type activates regulatory frameworks, how action context modifies that activation, and how contextual conditions refine the output leaves with them. The architecture exists in professional judgment rather than in any formal, transferable model.

Section 4 formalizes this architecture as an ontology.

4 The Ontology: Formal Architecture

This section presents the EHS Ontology (v3.1), a Web Ontology Language (OWL) formalization of the domain architecture described in Section 3. The ontology is implemented in Turtle (TTL) syntax using standard Semantic Web vocabularies: RDF and RDFS for class and property definitions, OWL for formal semantics, SKOS for concept definitions, and Dublin Core Terms for metadata. The ontology is structured as a core model with four regulatory modules that extend it into specific compliance domains.

4.1 Core Architecture

The ontology encodes the bipolar model as two top-level OWL classes, `ehs:Employee` and `ehs:Hazard`, connected by the object property `ehs:isExposedTo`. This property captures the fundamental relationship of the EHS domain: a worker contacts a hazard.

```
ehs:Employee rdf:type owl:Class ;
    skos:definition "A worker or individual in a workplace who may be
    exposed to hazards. One of the two poles of the EHS bipolar model." .

ehs:Hazard rdf:type owl:Class ;
    skos:definition "A source or situation with potential for harm.
```

```
    The second pole of the EHS bipolar model." .

ehs:isExposedTo rdf:type owl:ObjectProperty ;
    rdfs:domain ehs:Employee ;
    rdfs:range ehs:Hazard .
```

The Employee pole is not a flat class. It carries substructure through `ehs:WorkerCharacteristic`, with subclasses for `ehs:JobRole`, `ehs:TrainingStatus`, and `ehs:HealthStatus`. Employees are linked to their characteristics via `ehs:hasCharacteristic` and grouped into `ehs:SimilarExposureGroup` instances via `ehs:belongsToSEG` for exposure monitoring purposes.

The Hazard pole is described through the seven Industrial Hygiene (IH) concepts formalized as subclasses of `ehs:IndustrialHygieneConcept`: Toxicity, Exposure, DoseResponseRelationship, RouteOfEntry, OccupationalExposureLimit, TimeConcentrationRelationship, and IndividualSusceptibility. The object property `ehs:characterizedBy` links any Hazard instance to the IH concepts that describe it. RouteOfEntry is further subclassed into Inhalation, DermalAbsorption, Ingestion, and Injection.

The nexus where the two poles meet is `ehs:HazardousExposureSituation`, adapted from the central class in (Onut Badea et al. 2024). This class connects to both poles via `ehs:involvesEmployee` and `ehs:involvesHazard`, and links outward to the mediation layer via `ehs:mediatedBy`, to the ARECC decision process via `ehs:assessedBy`, and to outcomes via `ehs:developsInto`.

The 5 E's are formalized as subclasses of `ehs:SafetyProcess`: Education, Encouragement, Engineering, Enforcement, and Evaluation. Two directional properties capture the mediation role: `ehs:protects` (domain: SafetyProcess, range: Employee) and `ehs:mitigates` (domain: SafetyProcess, range: Hazard).

The ARECC framework is formalized as subclasses of `ehs:ARECCProcess`: Anticipate, Recognize, EvaluateARECC, ControlARECC, and Confirm. The Hierarchy of Controls is encoded as subclasses of `ehs:ControlMeasure` with a `ehs:hierarchyRank` datatype property (integer, 1 through 5) that preserves the effectiveness ordering: Elimination (1), Substitution (2), EngineeringControl (3), AdministrativeControl (4), and PPE (5).

4.2 Hazard Type as Compliance Activation Classifier

The central contribution of this ontology is the treatment of hazard type classification as a compliance activation mechanism rather than a descriptive label. Section 3.3 established this insight at the domain level. Here it is formalized.

`ehs:HazardType` is an OWL class with seven subclasses: PhysicalHazard, MechanicalHazard, ChemicalHazard, BiologicalHazard, PsychosocialHazard, ErgonomicHazard, and ElectricalHazard. Each Hazard instance is linked to its type via `ehs:hasHazardType`. Regulatory frameworks are formalized as subclasses of `ehs:RegulatoryFramework`, including OSHA_Framework (29 CFR 1910/1926), EPA_Framework (40 CFR), DOT_Framework (49 CFR), NFPA_Framework, NRC_Framework (10 CFR), MSHA_Framework (30 CFR), ANSI_Framework, GHS_Framework, ACGIH_Framework, NIOSH_Framework, and USCG_Framework (33/46 CFR).

The routing mechanism is the object property `ehs:activatesRegulation` (domain: HazardType, range: RegulatoryFramework). This property is instantiated through `ehs:ComplianceActivation` records that map each hazard type to its activated frameworks:

```
ehs:ChemicalActivation rdf:type ehs:ComplianceActivation ;
    ehs:triggeredByType ehs:ChemicalHazard ;
    ehs:activatesFramework ehs:OSHA_Framework , ehs:EPA_Framework ,
        ehs:GHS_Framework , ehs:ACGIH_Framework , ehs:DOT_Framework .

ehs:ElectricalActivation rdf:type ehs:ComplianceActivation ;
    ehs:triggeredByType ehs:ElectricalHazard ;
```

```
    ehs:activatesFramework ehs:OSHA_Framework , ehs:NFPA_Framework ,
        ehs:ANSI_Framework .
```

This pattern is repeated for all seven hazard types. The result is a complete mapping from hazard classification to regulatory activation. When a Hazard instance is classified as Chemical, the ontology produces the set {OSHA, EPA, GHS, ACGIH, DOT} as the activated frameworks. When it is classified as Electrical, the set is {OSHA, NFPA, ANSI}. The classification is not an annotation. It is a routing instruction.

4.3 Three-Axis Routing

Hazard type alone does not determine the full set of applicable regulations. A chemical spill activates different frameworks depending on whether the chemical was being transported, stored, processed, or disposed of at the time. The ontology captures this through a three-axis routing model: HazardType, ActionContext, and ContextualCondition.

4.3.1 Axis 2: Action Context

`ehs:ActionContext` represents the operational action being performed when a hazardous situation occurs. It is the verb of the incident. ActionContext is subclassed into nine categories:

TransportAction (loading, unloading, over-the-road transit) activates DOT/Pipeline and Hazardous Materials Safety Administration (PHMSA) under 49 CFR when hazardous materials cross or are at the property boundary. **StorageAction** (drum storage, tank farms, chemical warehousing) activates EPA RCRA for quantity thresholds and storage time limits. **ProcessingAction** (manufacturing, mixing, reacting) activates OSHA Process Safety Management (29 CFR 1910.119) for highly hazardous chemicals above threshold quantities. **MaintenanceAction** (repair, shutdown, servicing) activates OSHA Lockout/Tagout (29 CFR 1910.147) and confined space requirements (29 CFR 1910.146). **EmergencyResponseAction** activates OSHA Hazardous Waste Operations and Emergency Response (HAZWOPER, 29 CFR 1910.120) and EPA Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) for reportable quantities. **WasteHandlingAction** (collection, characterization, disposal) activates EPA RCRA (40 CFR 260-270) as the primary framework, with generator status determining specific requirements. **TransferAction** (decanting, dispensing between containers) activates OSHA Hazard Communication Standard (HCS) for secondary container labeling and NFPA 77 for grounding and bonding of flammables. **ConstructionAction** shifts jurisdiction from OSHA General Industry (29 CFR 1910) to OSHA Construction (29 CFR 1926) and may activate EPA National Emission Standards for Hazardous Air Pollutants (NESHAP) for asbestos. **SamplingAction** activates OSHA for worker protection during sample collection and NIOSH for sampling methodology.

Each ActionContext subclass carries its own `ehs:activatesFramework` assertions, adding frameworks to the baseline set activated by HazardType alone.

4.3.2 Axis 3: Contextual Conditions

`ehs:ContextualCondition` represents modifiers present at the time of the incident that further refine which specific regulations fire within the activated frameworks. Conditions answer three questions: Was it contained? Where exactly did it happen? How much was involved?

ContainmentStatus is the most consequential condition class. It has three subclasses. `ehs:FullyContained` indicates that secondary containment held and no material reached the environment. Per EPA/DOE guidance (U.S. Environmental Protection Agency 1992), a release into properly designed, constructed, and maintained secondary containment is not a release "to the environment" under CERCLA § 101(22). This determination is volume-independent: containment integrity, not quantity, is the primary regulatory gate. `ehs:PartialBreach` indicates some material escaped containment but was intercepted before reaching navigable waters, storm drains, or soil. `ehs:FullRelease` indicates material reached the environment, triggering the full regulatory cascade: CERCLA (if above the Reportable Quantity), Clean Water Act (CWA) Section 311 (if to waters), Clean Air Act (CAA) (if to air), and Emergency Planning and Community Right-to-Know Act (EPCRA) Sections 302/304 for community notification.

The ontology models an important edge case within FullyContained: volatile chemicals produce vapor emissions even when liquid containment is intact. CERCLA defines "environment" to include ambient air. If vapor mass from a contained spill exceeds the Reportable Quantity within 24 hours, CERCLA § 103 reporting may apply for the airborne component even when liquid containment is intact.

LocationContext determines jurisdictional boundaries. Its subclasses include LoadingDock (the jurisdictional interface between DOT and OSHA), ProductionFloor, StorageArea, ConfinedSpace (which automatically activates 29 CFR 1910.146 regardless of hazard type), PublicRoadway (DOT/PHMSA primary jurisdiction), WaterwayProximity (CWA Section 311 and potential USCG notification), and SecondaryContainmentArea.

QuantityThreshold determines whether reporting obligations are triggered. BelowReportableQuantity eliminates federal release reporting (though state thresholds may differ). AboveReportableQuantity triggers National Response Center (NRC) notification, EPCRA Section 304 emergency notification, and EPA CERCLA 103(a) reporting. AboveTPQ (Threshold Planning Quantity) triggers EPCRA Section 302 emergency planning. The ontology explicitly models the subordination relationship: QuantityThreshold is subordinate to ContainmentStatus. If containment is intact, CERCLA/EPCRA reporting does not apply regardless of quantity. The Reportable Quantity threshold only becomes the deciding factor when the release is to the environment.

Two object properties formalize how conditions modify the activation set. `ehs:narrowsActivation` removes a framework from the set (FullyContained narrows out CERCLA release reporting). `ehs:widensActivation` adds a framework (AboveReportableQuantity adds NRC notification). `ehs:shiftsJurisdiction` transfers primary authority from one framework to another (PublicRoadway shifts primary jurisdiction from OSHA to DOT).

4.3.3 The Routing Class

These three axes converge in `ehs:ContextualComplianceActivation`, a subclass of ComplianceActivation that represents the fully resolved regulatory routing for a specific scenario. Its properties are:

- `ehs:triggeredByType` (one or more HazardTypes)
- `ehs:hasActionContext` (the operational action)
- `ehs:hasContextCondition` (one or more contextual conditions)
- `ehs:activatesFramework` (the resulting set of activated frameworks)
- `ehs:specificCitation` (the precise CFR citations that apply)
- `ehs:scenarioDescription` (a plain-language narrative of the scenario)

The routing logic operates in sequence:

1. HazardType(s) fire the baseline framework set.
2. ActionContext adds or modifies frameworks based on the operational phase.
3. ContextualConditions further refine the output, narrowing, widening, or shifting the activation set.
4. The result is the precise regulatory obligation set for that specific incident in that specific context.

This architecture captures the operational reality described in the plain-language report (Section 1): the same drum of toluene triggers DOT 49 CFR in transit, OSHA 29 CFR in use, EPA 40 CFR in storage, and EPA RCRA in disposal. The chemical does not change. The regulatory routing changes entirely based on action context and conditions.

4.4 Compound Hazard Profiles

`ehs:CompoundHazardProfile`, a subclass of ComplianceActivation, models scenarios where multiple HazardTypes co-occur within a single HazardousExposureSituation. The design principle is explicit: compliance is additive, not selective. When Chemical and Ergonomic hazards co-occur, both OSHA HCS and NIOSH ergonomic guidelines fire simultaneously. The CompoundHazardProfile makes this union explicit.

```
ehs:ChemicalErgonomicProfile rdf:type ehs:CompoundHazardProfile ;
    ehs:combinesTypes ehs:ChemicalHazard , ehs:ErgonomicHazard ;
    ehs:activatesFramework ehs:OSHA_Framework , ehs:EPA_Framework ,
        ehs:GHS_Framework , ehs:ACGIH_Framework , ehs:NIOSH_Framework ,
        ehs:ANSI_Framework , ehs:ISO_Framework ;
    ehs:hasInteractionNote "Manual handling of chemical containers introduces
        MSD risk on top of chemical exposure. Awkward postures during chemical
        handling increase dermal absorption risk." .
```

The `ehs:combinesTypes` property links a CompoundHazardProfile to each constituent HazardType. The `ehs:hasInteractionNote` datatype property captures what is unique about the specific combination, including synergistic risks and special requirements that arise only from co-occurrence.

The ontology includes seven pre-defined profiles: Chemical + Ergonomic (manual drum handling), Chemical + Physical (heat increasing vapor pressure), Chemical + Electrical (flammable atmosphere near ignition sources), Biological + Chemical (disinfectant exposure alongside bloodborne pathogen risk), Mechanical + Ergonomic (repetitive machine operation), Physical + Psychosocial (noise exposure combined with shift work stress), and Chemical + Mechanical + Ergonomic (a triple overlap modeling manual drum handling near running machinery). Each profile carries the union of all constituent type activations plus any interaction-specific frameworks.

4.5 The Establishment Class

The ontology anchors facility-level regulatory obligations through `ehs:Establishment`, defined to match OSHA's definition in 29 CFR 1904.30(a): a single physical location where business is conducted or where services or industrial operations are performed. The Establishment class resolves a structural problem that arises when modeling EHS compliance: some obligations belong to individual employees (training, exposure monitoring, PPE fit testing) while others belong to the facility (chemical inventory reporting, OSHA 300 Log maintenance, Injury Tracking Application [ITA] electronic submission, Toxics Release Inventory [TRI] reporting).

`ehs:employs` (domain: Establishment, range: Employee) and its inverse `ehs:worksAt` connect the Establishment to the Employee pole. `ehs:hasNAICSCode` (North American Industry Classification System) determines TRI Section 313 coverage (manufacturing NAICS 31-33, metal mining 2122, electric utilities 2211, and others) and OSHA recordkeeping exemptions. `ehs:hasEstablishmentSize` (peak employee count) determines ITA electronic submission tiers: 20-249 employees in Appendix A industries submit 300A data only, 100+ in Appendix B industries submit 300, 301, and 300A data, and 250+ in all covered industries submit 300A data.

The Establishment class serves as the domain for three critical facility-level properties: `ehs:hasChemicalInventory` (linking to EPCRA Tier II in Module A), `ehs:hasEmissionUnit` (linking to Title V air permitting in Module B), and `ehs:hasInjuryIllnessRate` (linking to OSHA 300 recordkeeping in Module C).

4.6 Module A: EPCRA Chemical Inventory and TRI Reporting

Module A formalizes the Emergency Planning and Community Right-to-Know Act (EPCRA) reporting structure. EPCRA, enacted as Title III of the Superfund Amendments and Reauthorization Act (SARA) of 1986 in response to the Bhopal disaster, establishes four reporting pillars: emergency planning (Sections 302-303), release notification (Section 304), chemical inventory reporting (Sections 311-312), and the Toxics Release Inventory (Section 313).

The module centers on `ehs:ChemicalInventory`, a facility-level inventory of hazardous chemicals linked to the Establishment via `ehs:hasChemicalInventory`. Individual chemicals are modeled as `ehs:InventoryChemical` instances, each tracking chemical identity, CAS number, maximum quantity present, average daily amount, number of days on site, storage type and conditions, and facility location. The `ehs:linkedToSDS` property creates the bridge between OSHA HCS and EPCRA: if OSHA requires a

Safety Data Sheet under 29 CFR 1910.1200, the chemical is potentially subject to EPCRA 311/312 reporting if present above threshold quantities.

`ehs:ExtremelyHazardousSubstance` (EHS) models the subset of chemicals listed in 40 CFR 355 Appendices A and B. EHS chemicals carry a Threshold Planning Quantity (TPQ) that triggers Section 302 emergency planning notification. For Tier II reporting under Section 312, the threshold for an EHS chemical is the TPQ or 500 pounds, whichever is less. For all other OSHA HCS hazardous chemicals, the Tier II reporting threshold is 10,000 pounds. The `ehs:ReportingThreshold` class formalizes this two-tier structure.

The reporting obligation chain flows through four classes: `ehs:Section302Notification` (initial EHS notification to the State Emergency Response Commission [SERC] and Local Emergency Planning Committee [LEPC] within 60 days of exceeding the TPQ), `ehs:Section311Notification` (one-time notification that reportable chemicals are present, submitted to SERC, LEPC, and local fire department), `ehs:Section312TierIIReport` (annual inventory report due March 1), and `ehs:Section304ReleaseNotification` (immediate notification when a release exceeds the Reportable Quantity).

4.6.1 TRI Reporting (Section 313)

The v3.1 expansion adds `ehs:TRIChemical` as a subclass of InventoryChemical. This subclass relationship encodes a critical structural fact: every TRI chemical is also an inventory chemical tracked for EPCRA 311/312, but not every inventory chemical is TRI-listed. As of 2024, approximately 800 individually listed chemicals and 33 chemical categories appear on the TRI list (40 CFR 372.65).

TRI applicability follows a three-prong test formalized through the Establishment class: covered NAICS code (`ehs:hasNAICSCode`), 10 or more full-time employees (`ehs:hasEstablishmentSize`), and at least one TRI chemical above activity thresholds. `ehs:TRIActivityThreshold` models the quantity triggers: 25,000 pounds per year for chemicals that are manufactured or processed, and 10,000 pounds per year for chemicals otherwise used (such as solvents used for cleaning).

`ehs:TRIDeMinimisExemption` formalizes the concentration exemption: chemicals present at less than 1% by weight in mixtures need not be counted toward activity thresholds, with a reduced 0.1% threshold for OSHA carcinogens and Persistent, Bioaccumulative, and Toxic (PBT) chemicals under 40 CFR 372.28. `ehs:TRIPersistentBioaccumulativeChemical` is a subclass of TRIChemical with lower reporting thresholds (100 pounds for lead compounds, 10 pounds for mercury compounds, 0.1 grams for dioxin/furan compounds) and ineligibility for Form A certification.

The ontology includes concrete chemical instances drawn from automotive manufacturing: Xylene (CAS 1330-20-7), Toluene (CAS 108-88-3), Chromium Compounds (category N090, with dual de minimis thresholds of 0.1% for hexavalent and 1.0% for trivalent), and Formaldehyde (CAS 50-00-0, which appears on three lists: TRI, EHS, and CAA Section 112(r)). A corrective instance, `ehs:MEK_Delisted`, documents that Methyl Ethyl Ketone was delisted from both TRI (June 2005) and the CAA HAP list (December 2005), demonstrating how regulatory status changes over time affect compliance routing.

4.7 Module B: Title V / Clean Air Act Air Permitting

Module B formalizes the Clean Air Act (CAA) Title V operating permit program (40 CFR Part 70). The core structure models how facilities determine whether they are major sources of air pollution and, if so, how their emissions are regulated.

`ehs:EmissionUnit` is the fundamental building block. Each emission unit represents a discrete piece of equipment, process, or activity that emits or has the potential to emit air pollutants. Examples include boilers, paint spray booths, solvent degreasers, and storage tanks. `ehs:FugitiveEmissionSource` is a subclass for sources that release pollutants through means other than a stack or vent, such as equipment leaks and open material handling. Each emission unit may use one or more `ehs:ControlDevice` instances (thermal oxidizers, scrubbers, baghouses, carbon adsorbers) linked via `ehs:hasControlDevice`. Control device efficiency determines the difference between pre-control and post-control Potential to Emit (PTE).

`ehs:PotentialToEmit` models the gateway calculation for all air permitting. PTE represents the maximum capacity of a source to emit a pollutant assuming maximum-rated capacity, 24 hours per day, 365 days per year, unless the source is subject to federally enforceable limitations. A facility's total PTE is the sum of all emission unit PTEs. PTE is calculated for each pollutant type against the applicable `ehs:MajorSourceThreshold`: 100 tons per year (tpy) for any criteria pollutant, 10 tpy for any single Hazardous Air Pollutant (HAP), 25 tpy for combined HAPs, and 100,000 tpy carbon dioxide equivalent (CO2e) for Greenhouse Gases (GHGs). `ehs:NonattainmentArea` models geographic areas that do not meet National Ambient Air Quality Standards (NAAQS), where lower major source thresholds apply.

The pollutant taxonomy includes four subclasses of `ehs:AirPollutant`. `ehs:CriteriaPollutant` covers the six NAAQS pollutants (SO2, NOx, CO, PM10/PM2.5, lead, and ozone regulated via VOC and NOx as precursors). `ehs:HazardousAirPollutant` covers the 187 toxic air pollutants listed in CAA Section 112(b). `ehs:VolatileOrganicCompound` covers carbon compounds that participate in atmospheric photochemical reactions. `ehs:GreenhouseGas` covers CO2, CH4, N2O, HFCs, PFCs, SF6, and NF3 as defined in the Mandatory Reporting Rule (40 CFR Part 98).

The permitting structure includes two classes. `ehs:TitleVPermit` is the comprehensive operating permit required for major sources, consolidating all applicable federal and state air requirements into a single federally enforceable document valid for five years. `ehs:FESOP` (Federally Enforceable State Operating Permit) allows a facility to accept restrictions that keep PTE below major source thresholds, creating a "synthetic minor" status that avoids Title V permitting.

Technology standards are modeled through two classes. `ehs:MACT` (Maximum Achievable Control Technology) represents emission standards for major sources of HAPs under CAA Section 112, codified in 40 CFR Part 63 as NESHAP. `ehs:NSPS` (New Source Performance Standards) represents standards for new and modified sources under CAA Section 111, codified in 40 CFR Part 60.

`ehs:ComplianceAssuranceMonitoring` (CAM) models the monitoring program under 40 CFR Part 64 for emission units at major sources that use control devices. CAM requires monitoring of operational parameters (temperature, pressure drop, flow rate) that indicate control device performance.

4.8 Module C: OSHA 300 Recordkeeping

Module C formalizes the OSHA injury and illness recording and reporting program (29 CFR 1904). This module models the decision logic that determines whether a workplace event is recordable, how it is classified, how rates are calculated, and when events must be reported directly to OSHA.

The recording decision operates as a two-gate model. The first gate is `ehs:WorkRelatedness`: an injury or illness is work-related if an event or exposure in the work environment caused, contributed to, or significantly aggravated the condition. Work-relatedness is presumed for events occurring in the work environment unless a specific exception in 29 CFR 1904.5(b)(2) applies. The ontology documents these exceptions, including presence as a member of the general public, symptoms from non-work events surfacing at work, voluntary wellness program participation, personal tasks outside assigned hours, motor vehicle accidents in parking lots during commute, and common cold or flu.

The second gate is `ehs:RecordingCriteria`. A work-related injury or illness is recordable if it results in any one of six outcomes, each modeled as a subclass: `ehs:DeathCriterion`, `ehs:DaysAwayCriterion`, `ehs:RestrictedWorkCriterion`, `ehs:MedicalTreatmentCriterion` (treatment beyond first aid), `ehs:LossOfConsciousnessCriterion`, and `ehs:SignificantDiagnosisCriterion` (cancer, chronic irreversible disease, fractured bone, punctured eardrum diagnosed by a physician or licensed healthcare professional [PLHCP]). Each criterion operates independently. Meeting any single criterion makes the case recordable.

The MedicalTreatmentCriterion carries particular formal significance because the boundary between first aid and medical treatment is defined by an exhaustive list in 29 CFR 1904.7(b)(5)(ii). The ontology documents this list in full. Any treatment not on the list is medical treatment by definition. The professional status of the provider is irrelevant: first aid administered by a physician remains first aid, and medical treatment administered by a non-physician remains medical treatment. Once medical treatment is provided, the case is recordable even if a second opinion later recommends first aid.

Three special recording criteria are subclassed separately: `ehs:NeedlestickCriteria` (29 CFR 1904.8, with privacy case designation), `ehs:HearingLossCriteria` (29 CFR 1904.10, requiring both a Standard Threshold Shift and a 25 dB total hearing level), and `ehs:MedicalRemovalCriteria` (29 CFR 1904.9, for substance-specific medical surveillance programs including lead, cadmium, benzene, and formaldehyde).

`ehs:CaseClassification` models the mutually exclusive categories entered on the 300 Log: injury, skin disorder, respiratory condition, poisoning, hearing loss, and all other illnesses.

The module formalizes four rate calculations as subclasses of `ehs:InjuryIllnessRate`, all calculated per 200,000 hours worked (equivalent to 100 full-time employees working 40 hours per week for 50 weeks). `ehs:TRIR` (Total Recordable Incident Rate) includes all recordable cases. `ehs:DART` (Days Away, Restricted, or Transferred rate) excludes other-recordable cases and is considered a better indicator of severity. `ehs:LTIR` (Lost Time Incident Rate) includes only cases with days away from work. `ehs:SeverityRate` measures total days away and restricted per 200,000 hours. All rate calculations are anchored to the Establishment class via `ehs:hasInjuryIllnessRate`, because rates are calculated per establishment, not per employee.

The three OSHA forms are modeled as classes: `ehs:OSHA300Log` (the running log, maintained per establishment), `ehs:OSHA300ASummary` (the annual summary, posted February 1 through April 30), and `ehs:OSHA301Report` (the detailed incident report for each recordable case). `ehs:OSHAReportingObligation` captures the events that must be reported directly to OSHA: fatality within 8 hours, and in-patient hospitalization, amputation, or loss of an eye within 24 hours. `ehs:ElectronicSubmission` models the ITA annual submission requirements, with tiers determined by establishment size and NAICS code.

4.9 Module D: Incident Management

Module D extends the existing Outcome hierarchy (Accident, Incident, OccupationalDisease) with investigation workflow, root cause analysis, corrective action tracking, and feedback loops to the ARECC framework and Hierarchy of Controls.

`ehs:IncidentSeverity` classifies events for investigation prioritization. Its subclasses span the full severity spectrum: Fatality (triggering OSHA 8-hour reporting and the most comprehensive investigation), LostTimeIncident, RestrictedDutyIncident, MedicalTreatmentIncident, FirstAidIncident (not OSHA-recordable but tracked as a leading indicator), NearMiss (the most valuable leading indicator, per Heinrich's Triangle), PropertyDamageIncident, and EnvironmentalIncident (unplanned releases that may trigger EPA, EPCRA, or CERCLA reporting independent of worker injury).

The `ehs:alignsWithRecordingCriteria` property (domain: IncidentSeverity, range: RecordingCriteria) maps Module D severity classifications to Module C recording criteria: Fatality aligns with DeathCriterion, LostTimeIncident with DaysAwayCriterion, RestrictedDutyIncident with RestrictedWorkCriterion, and MedicalTreatmentIncident with MedicalTreatmentCriterion. FirstAidIncident, NearMiss, and PropertyDamageIncident have no corresponding recording criteria because they are not OSHA-recordable.

`ehs:IncidentInvestigation` models the systematic examination of an incident. It connects to the investigation team (`ehs:InvestigationTeam`, scaled by severity), the root cause analysis method, the corrective actions generated, and the ARECC framework via `ehs:connectsToARECC` (range: ARECCProcess). An investigation is fundamentally a Confirm-phase activity: it confirms whether prior ARECC phases produced adequate controls.

Root cause analysis methods are formalized as subclasses of `ehs:RootCauseAnalysis`: FiveWhys (best suited for single causal chains), FishboneDiagram (Ishikawa, for brainstorming multiple contributing factors), FaultTreeAnalysis (Boolean logic for complex, high-consequence incidents), and TapRooT (proprietary structured methodology for standardized causal factor trending).

`ehs:CorrectiveAction` is linked to the Hierarchy of Controls via `ehs:addressedByControl` (range: ControlMeasure), forcing explicit documentation of where each corrective action falls in the hierarchy and justification if it falls toward the less effective end. `ehs:CorrectiveActionStatus` tracks the lifecycle: Open, In Progress, Completed, Verified, and Overdue. Verification is the critical step. A corrective action is not complete until its effectiveness is confirmed through follow-up observation or audit.

`ehs:EnvironmentalIncident` carries a dedicated property, `ehs:triggersComplianceActivation` (range: ContextualComplianceActivation), that routes environmental incidents directly into the three-axis compliance activation model. A chemical spill that reaches a floor drain triggers an EnvironmentalIncident. The `triggersComplianceActivation` property routes it to the ContextualComplianceActivation instance that evaluates HazardType, ActionContext, and ContextualConditions to determine which EPA, EPCRA, CERCLA, CWA, or state reporting obligations fire.

4.10 Cross-Module Integration

The four modules are not independent silos. The ontology wires them together through shared classes and cross-linking properties that reflect how regulatory obligations interconnect in practice.

The primary cross-link runs between Module A (chemical inventory) and Module B (air permitting) through the `ehs:usesChemical` property (domain: EmissionUnit, range: InventoryChemical) and its inverse `ehs:emittedFromUnit`. A single chemical like toluene appears in the facility's chemical inventory (tracked for Tier II reporting thresholds under Module A) and as a pollutant emitted from an emission unit (tracked for PTE and permit limits under Module B). When TRI reporting applies, the same chemical also appears as a TRIChemical with release quantities calculated from the emission unit's data, creating a three-way linkage: inventory tracking, emission unit PTE, and TRI release quantity.

The `ehs:worksAtEmissionUnit` property (domain: Employee, range: EmissionUnit) connects Module B back to the Employee pole. A paint spray booth operator's personal exposure (assessed via the IH concepts in the core model) is driven by the same chemicals whose emissions are calculated for PTE and reported on TRI Form R. This property enables queries such as: "Which employees work at emission units that emit HAPs above the occupational exposure limit?"

Module C and Module D are aligned through the `ehs:alignsWithRecordingCriteria` property, ensuring that a severity classification in the incident management system reliably predicts the OSHA recording outcome. Module D feeds back into the core model through `ehs:connectsToARECC`, closing the loop between incident investigation findings and the proactive hazard management cycle.

The Establishment class serves as the integration anchor across all four modules. It is the domain for chemical inventory (Module A), emission units (Module B), injury and illness rates (Module C), and TRI reporting obligations. A company may have multiple establishments, each maintaining its own OSHA 300 Log, its own Tier II report, its own Title V permit, and its own TRI submissions.

This cross-module architecture reflects the operational reality that EHS practitioners manage: a single chemical arrives at the loading dock (triggering DOT and OSHA in the core routing model), enters the chemical inventory (Module A), is consumed by an emission unit that emits it to the air (Module B), exposes workers who may be injured (Modules C and D), and generates waste that must be disposed of (back to core routing under WasteHandlingAction). The ontology traces this lifecycle through formal, traversable relationships rather than leaving the connections implicit in practitioner knowledge.

5 Scenarios: Routing Validation

This section validates the ontology's routing logic through nine worked scenarios. Each scenario presents a concrete workplace event, walks it through the three-axis routing model (HazardType × ActionContext × ContextualCondition), and produces the complete regulatory output. The first six scenarios involve chemical hazards in various operational contexts. The final three demonstrate that the routing model applies equally to non-chemical hazard types, including pure mechanical/physical events, chronic noise exposure combined with psychosocial stressors, and a real-world recordable incident drawn from the author's professional practice.

For each scenario, the walkthrough follows the same structure: identify the HazardType(s), identify the ActionContext, identify the ContextualConditions, resolve the routing, and state the resulting regulatory obligation set. This repetitive structure is intentional. It demonstrates that the routing model is systematic rather than ad hoc, producing deterministic outputs from classified inputs.

5.1 Scenario 1: Chemical Spill at Loading Dock During Offloading (Uncontained)

A chemical drum ruptures while being offloaded from a delivery truck at the facility loading dock. The spill reaches the ground outside secondary containment. The carrier's driver is still present at the dock.

Axis 1 — HazardType: Chemical. Baseline activation: {OSHA, EPA, GHS, ACGIH, DOT}.

Axis 2 — ActionContext: TransportAction (loading/unloading at property boundary). Adds DOT/PHMSA (49 CFR) jurisdiction. Per 49 CFR 171.1(c)(2)-(3), unloading in the presence of the carrier is "incidental to movement" and within DOT jurisdiction.

Axis 3 — ContextualConditions:

- LocationContext: LoadingDock (jurisdictional interface between DOT and OSHA).
- ContainmentStatus: FullRelease (material reached environment, no containment barrier).
- QuantityThreshold: AboveReportableQuantity.

Routing resolution:

- DOT fires: 49 CFR 171.15 (immediate telephonic notice) and 49 CFR 171.16 (written hazmat incident report).
- EPA CERCLA fires: Section 103(a), NRC notification required because the release is to the environment and above the Reportable Quantity (RQ).
- EPA EPCRA fires: Section 304, emergency notification to LEPC and SERC.
- EPA RCRA fires: 40 CFR 262, contaminated soil and absorbent pads from cleanup are hazardous waste requiring characterization, manifesting, and disposal via licensed Treatment, Storage, and Disposal Facility (TSDF).
- OSHA fires: 29 CFR 1910.120 (HAZWOPER) if the release is uncontrolled. 29 CFR 1910.1200 (HCS) for employee exposure during cleanup, requiring SDS access, labeling, and PPE.

Result: Three federal agencies (DOT, EPA, OSHA), four or more regulatory programs, multiple notification calls, and written reports within 30 days.

Edge case (DOT carrier-presence): If the carrier's driver has already departed before the spill, the unloading is no longer "incidental to movement." DOT jurisdiction does not apply. The routing output drops to {EPA, OSHA}, demonstrating that a single contextual condition (carrier presence) gates an entire agency's jurisdiction.

5.2 Scenario 2: Chemical Spill in Secondary Containment (Contained)

The same chemical, the same drum, and the same quantity as Scenario 1. This time, the drum is sitting on a spill containment pallet inside the facility's chemical storage area. An employee organizing the storage room knocks the drum over. The containment pallet catches every drop.

Axis 1 — HazardType: Chemical. Baseline activation: {OSHA, EPA, GHS, ACGIH, DOT}.

Axis 2 — ActionContext: StorageAction (organizing chemical storage). No transport context.

Axis 3 — ContextualConditions:

- LocationContext: SecondaryContainmentArea.
- ContainmentStatus: FullyContained. Per EPA/DOE guidance (U.S. Environmental Protection Agency 1992), a release into properly designed, constructed, and maintained secondary containment is not a release "to the environment" under CERCLA § 101(22). This exemption is volume-independent.
- QuantityThreshold: Subordinate to ContainmentStatus. Because containment is intact, CERCLA/EPCRA reporting does not apply regardless of the quantity spilled.

Routing resolution:

- DOT does not fire: No transportation context.
- EPA CERCLA does not fire: No release to the environment. Containment held.
- EPA EPCRA does not fire: Same basis as CERCLA.
- EPA RCRA fires: 40 CFR 262 and 40 CFR 264.175. The contained chemical must be properly characterized and disposed of as hazardous waste.
- OSHA HCS fires: 29 CFR 1910.1200. The employee performing cleanup needs SDS access, proper labeling, and appropriate PPE.

Result: One agency for disposal (EPA RCRA) and one OSHA standard for worker protection during cleanup. Compare this to Scenario 1's cascade of three agencies and four or more programs.

The routing difference: The only variable that changed between Scenarios 1 and 2 is ContainmentStatus. The chemical did not change. The quantity did not change. Containment status, a single ContextualCondition, is the primary gate that determines whether one phone call or a dozen is required. The ontology captures this gate formally through the `narrowsActivation` property on FullyContained, which removes CERCLA release reporting from the activation set.

Edge case (volatile vapor emissions): If the contained chemical is volatile (benzene, toluene, paint thinner), the liquid stays in the containment tray but vapors enter ambient air. CERCLA defines "environment" to include ambient air. If vapor mass exceeds the RQ within 24 hours, CERCLA § 103 reporting may apply for the airborne component even when liquid containment is intact. For benzene (vapor pressure 95 mmHg, RQ 10 pounds), partial evaporation of a 55-gallon drum could exceed this threshold. The ontology models this edge case within the FullyContained class definition, noting that containment catches liquid but does not catch vapor.

5.3 Scenario 3: Chemical + Ergonomic During Manual Drum Handling

A worker manually lifts and tilts a 55-gallon drum of solvent to decant the contents into a smaller container on the production floor. The drum weighs approximately 400 pounds when full.

Axis 1 — HazardType: CompoundHazardProfile (Chemical + Ergonomic). Chemical hazard: inhalation of solvent vapors and dermal contact during pouring. Ergonomic hazard: heavy lifting, awkward posture, and repetitive drum tilting. Compound activation is the union of both constituent types: {OSHA, EPA, GHS, ACGIH, DOT}   {OSHA, ANSI, ISO} = {OSHA, EPA, GHS, ACGIH, DOT, ANSI, ISO, NIOSH}.

Axis 2 — ActionContext: TransferAction (decanting between containers). Adds OSHA HCS for secondary container labeling and NFPA 77 for grounding/bonding of flammables.

Axis 3 — ContextualConditions:

- LocationContext: ProductionFloor.
- ContainmentStatus: Not applicable (no spill has occurred; this is routine operation).

Routing resolution:

- OSHA HCS fires: 29 CFR 1910.1200 for chemical labeling, SDS access, and training.
- OSHA General Duty Clause fires: Section 5(a)(1) for the ergonomic hazard. OSHA does not have a specific ergonomic standard but cites ergonomic hazards under the General Duty Clause.
- ACGIH TLV fires: Exposure assessment for the solvent vapor. If the solvent has a skin notation, the awkward posture during handling increases dermal absorption risk.
- NIOSH Lifting Equation fires: Manual handling assessment for the 400-pound drum.
- ANSI Z365 fires: Ergonomic program requirements.
- GHS fires: Labeling and SDS format for the solvent.

Result: Both compliance pathways fire simultaneously. Neither cancels the other. The Job Hazard Analysis (JHA) for this task must address both the chemical and ergonomic dimensions. A JHA that addresses only the chemical exposure or only the lifting hazard is incomplete.

Interaction note: The compound profile carries synergistic risk. Awkward postures during chemical handling increase respiration rate (higher inhalation dose) and expose more skin surface area (higher dermal absorption). The two hazard types do not merely coexist. They amplify each other.

5.4 Scenario 4: Hazmat Release on Public Roadway

A tanker truck overturns on a public highway, releasing a hazardous chemical. The spill reaches the roadway and shoulder.

Axis 1 — HazardType: Chemical. Baseline activation: {OSHA, EPA, GHS, ACGIH, DOT}.

Axis 2 — ActionContext: TransportAction (over-the-road transit). DOT is primary jurisdiction.

Axis 3 — ContextualConditions:

- LocationContext: PublicRoadway. `shiftsJurisdiction` transfers primary authority from OSHA to DOT. OSHA does not regulate the driving activity.
- ContainmentStatus: FullRelease (material reached environment).
- QuantityThreshold: AboveReportableQuantity.

Routing resolution:

- DOT fires as primary agency: 49 CFR 171.15/171.16 (hazmat incident reporting), 49 CFR 172.600 (Emergency Response Guidebook information).
- EPA CERCLA fires: Section 103(a), NRC notification.
- EPA EPCRA fires: Section 304, LEPC/SERC notification.
- EPA CWA fires: Section 311 if the spill is near navigable waters.
- USCG fires: If the spill occurs in a coastal or port area.
- OSHA fires only for response workers: 29 CFR 1910.120 (HAZWOPER) applies to the emergency responders at the scene, not to the public or the truck driver (who is under DOT jurisdiction).
- State DOT and state environmental agency emergency notifications are also triggered.

Result: DOT is primary. OSHA activates only for response workers, not for the transportation activity itself. The `shiftsJurisdiction` property on PublicRoadway formally captures this jurisdictional transfer, distinguishing this scenario from loading dock scenarios where OSHA and DOT share jurisdiction simultaneously.

5.5 Scenario 5: Biological + Chemical in Healthcare Maintenance (Confined Space)

A maintenance worker enters an autoclave room (permit-required confined space) to repair sterilization equipment at a healthcare facility. The space contains residual biological contamination and glutaraldehyde vapor.

Axis 1 — HazardType: CompoundHazardProfile (Biological + Chemical). Biological hazard: bloodborne pathogen exposure from residual contamination. Chemical hazard: glutaraldehyde vapor inhalation and skin contact. Compound activation: {OSHA, EPA}   {OSHA, EPA, GHS, ACGIH, DOT} = {OSHA, EPA, GHS, ACGIH, DOT, NIOSH}.

Axis 2 — ActionContext: MaintenanceAction (repair of equipment). Adds OSHA LOTO (29 CFR 1910.147) for the sterilization equipment.

Axis 3 — ContextualConditions:

- LocationContext: ConfinedSpace. Automatically activates OSHA 29 CFR 1910.146 (Permit-Required Confined Spaces) regardless of hazard type.

Routing resolution:

Three OSHA standards fire simultaneously:

- OSHA Bloodborne Pathogens (BBP): 29 CFR 1910.1030 for the biological hazard.
- OSHA HCS: 29 CFR 1910.1200 for the glutaraldehyde chemical exposure.
- OSHA Permit-Required Confined Spaces (PRCS): 29 CFR 1910.146 for the confined space entry.
- OSHA LOTO: 29 CFR 1910.147 for the maintenance work on the sterilization equipment.
- NIOSH REL for glutaraldehyde fires as the more protective exposure limit (the NIOSH REL is lower than the OSHA PEL for glutaraldehyde, and best practice is to use the more protective limit).
- EPA RCRA fires for the mixed biological/chemical waste generated during cleanup and maintenance.

Result: Four OSHA standards fire simultaneously for a single task. The ConfinedSpace location condition adds PRCS regardless of whether the hazard is chemical, biological, or both. The MaintenanceAction context adds LOTO regardless of hazard type. This scenario demonstrates that ActionContext and LocationContext contribute independently to the regulatory output, and that their contributions are additive.

5.6 Scenario 6: Electrical + Chemical in Confined Space

An electrician enters a tank interior (permit-required confined space) to repair a level sensor. The tank previously held flammable solvent. The atmosphere contains solvent vapor below the Lower Explosive Limit (LEL) but at a measurable concentration.

Axis 1 — HazardType: CompoundHazardProfile (Electrical + Chemical). Electrical hazard: arc flash potential from energized sensor circuitry. Chemical hazard: flammable atmosphere from residual solvent vapor. Compound activation: {OSHA, NFPA, ANSI}   {OSHA, EPA, GHS, ACGIH, DOT} = {OSHA, NFPA, ANSI, EPA, GHS, ACGIH, DOT}.

Axis 2 — ActionContext: MaintenanceAction. Adds OSHA LOTO (29 CFR 1910.147).

Axis 3 — ContextualConditions:

- LocationContext: ConfinedSpace. Adds OSHA PRCS (29 CFR 1910.146).

Routing resolution:

- OSHA Subpart S fires: Electrical safety standards for the energized sensor work.
- NFPA 70E fires: Arc flash safety, electrical safety program requirements.
- NFPA 70 fires: Articles 500-506, hazardous location classification for the flammable atmosphere. This determines whether intrinsically safe electrical tools are required.
- OSHA HCS fires: 29 CFR 1910.1200 for the solvent vapor exposure.
- OSHA PRCS fires: 29 CFR 1910.146 for the confined space entry.
- OSHA LOTO fires: 29 CFR 1910.147 for the equipment repair.
- ACGIH TLV fires: Exposure assessment for the flammable solvent.

Result: The compound electrical + chemical profile produces a regulatory set that neither single-type activation would fully cover alone. The flammable atmosphere classification under NFPA 70 arises specifically from the co-occurrence of electrical work and chemical vapor. An ontology that modeled only the electrical hazard or only the chemical hazard would miss this interaction-specific requirement.

Interaction note: The `hasInteractionNote` on the ChemicalElectricalProfile captures this: electrical ignition sources in a flammable atmosphere trigger NFPA hazardous location classification, which governs the type of electrical equipment permitted in the space. This requirement does not exist for electrical work alone or for chemical exposure alone.

5.7 Scenario 7: Electrical + Mechanical During Conveyor Maintenance

An electrician performs maintenance on a conveyor system in a packaging area. The task requires accessing an energized electrical panel to troubleshoot a motor controller while the conveyor belt and drive mechanism are locked out but adjacent conveyors remain operational.

Axis 1 — HazardType: CompoundHazardProfile (Electrical + Mechanical). Electrical hazard: arc flash and shock from the energized panel. Mechanical hazard: caught-in/caught-between from adjacent operational conveyors and from the conveyor under repair if LOTO is breached. Compound activation: {OSHA, NFPA, ANSI}   {OSHA, ANSI} = {OSHA, NFPA, ANSI}.

Axis 2 — ActionContext: MaintenanceAction. Adds OSHA LOTO (29 CFR 1910.147).

Axis 3 — ContextualConditions:

- LocationContext: ProductionFloor.

Routing resolution:

- OSHA Subpart S fires: 29 CFR 1910 Subpart S for electrical safety during troubleshooting of the energized panel.
- OSHA Subpart O fires: 29 CFR 1910 Subpart O for machine guarding on adjacent operational conveyors.
- OSHA LOTO fires: 29 CFR 1910.147 for the lockout/tagout of the conveyor under repair. The scope of LOTO must address whether adjacent conveyors create a group lockout requirement.
- NFPA 70E fires: Arc flash risk assessment, approach boundaries, and PPE selection for the energized panel work.
- ANSI fires: Machine safety standards for the conveyor system.

Result: {OSHA, NFPA, ANSI}. No EPA. No DOT. No chemical involvement. The routing model produces a regulatory output governed entirely by workplace safety standards. This scenario demonstrates that the three-axis model is not dependent on chemical hazards or environmental release conditions. The same routing architecture (HazardType activates baseline, ActionContext adds LOTO, LocationContext is neutral) produces the correct output for a pure mechanical/electrical event.

5.8 Scenario 8: Physical + Psychosocial on Night Shift (Chronic Exposure)

Workers in an automotive stamping plant operate heavy presses on a rotating night shift schedule. Noise levels on the production floor exceed the OSHA PEL of 90 dBA as an 8-hour TWA. Workers report chronic fatigue, difficulty sleeping during daytime hours, and increasing stress from sustained production pressure.

Axis 1 — HazardType: CompoundHazardProfile (Physical + Psychosocial). Physical hazard: noise exposure above the OSHA PEL. Psychosocial hazard: rotating shift work, chronic fatigue, workplace stress from production demands. Compound activation: {OSHA, NRC, ACGIH}   {OSHA, ISO} = {OSHA, ACGIH, ISO, NIOSH}.

Axis 2 — ActionContext: ProcessingAction (manufacturing operations). No additional frameworks added beyond baseline for this context.

Axis 3 — ContextualConditions:

- LocationContext: ProductionFloor.
- No containment, quantity, or release conditions apply. This is a chronic exposure scenario, not an acute incident.

Routing resolution:

- OSHA Noise Standard fires: 29 CFR 1910.95. At 90 dBA TWA, the employer must implement a hearing conservation program including annual audiometric testing, hearing protection, employee training, and noise monitoring. The Action Level (85 dBA) triggers program initiation; the PEL (90 dBA) triggers mandatory hearing protection.
- ACGIH TLV for noise fires: ACGIH recommends a more protective 85 dBA TWA ceiling, which would classify the 90 dBA exposure as above the recommended limit and strengthen the case for engineering noise controls rather than relying on hearing protection alone.
- ISO 45003 fires: Psychological health and safety at work. Addresses shift work design, workload management, and organizational factors contributing to worker stress.

- NIOSH Total Worker Health fires: Integrated approach addressing the intersection of occupational noise exposure and shift work as compounding health risk factors.
- OSHA General Duty Clause: Applicable to the psychosocial stressors if specific conditions (workplace violence risk, documented stress-related health outcomes) support a citation.

Result: {OSHA, ACGIH, ISO, NIOSH}. No EPA. No DOT. No chemical involvement. No acute incident. This scenario demonstrates that the routing model handles chronic, non-incident exposures where the "event" is sustained workplace conditions rather than a discrete spill, release, or equipment failure. The Physical + Psychosocial compound profile activates standards that neither type would fully activate alone: noise standards do not address shift work stress, and psychosocial guidelines do not address hearing conservation. Both apply simultaneously.

Interaction note: Chronic noise exposure and shift work stress share a compounding health pathway. Noise-induced stress amplifies psychosocial strain. Sleep disruption from shift work reduces tolerance to noise exposure. The interaction is bidirectional, and the compound profile captures this through the `hasInteractionNote` property.

5.9 Scenario 9: Physical + Mechanical at Loading Dock (Dock Plate Failure)

An employee is standing on a dock plate while unloading a delivery truck. The employee was trained to engage the dock lock, a mechanical device that secures the trailer to the dock and prevents the truck from pulling away prematurely. The dock is equipped with go/no-go indicator lights on both the interior and exterior, designed to communicate safety status to two different audiences. The interior light faces the employee: green indicates the dock lock is engaged and it is safe to enter the trailer; red indicates the lock is not engaged and it is unsafe to enter. The exterior light faces the truck driver: green indicates the dock lock is disengaged and it is safe to pull away; red indicates the lock is engaged and the driver must remain in position. When functioning correctly, the interior and exterior lights display opposite colors (interior green/exterior red when locked, interior red/exterior green when unlocked). The employee does not engage the dock lock. Both lights malfunction: the interior light shows green (safe to enter) and the exterior light shows green (safe to pull away) simultaneously, a state that should be mechanically impossible. The truck driver, seeing a green exterior light, pulls away from the dock. The dock plate drops. The employee falls, dislocating their arm.

Axis 1 — HazardType: CompoundHazardProfile (Physical + Mechanical). Physical hazard: fall from the dock plate to the ground below. Mechanical hazard: the dock plate as a moving mechanical element that shifted when the truck departed, creating a struck-by/caught-between condition. Compound activation: {OSHA, NRC, ACGIH} {OSHA, ANSI} = {OSHA, ANSI}.

Axis 2 — ActionContext: TransportAction (unloading at the dock). No additional frameworks beyond baseline for this non-hazmat context. DOT does not activate because no hazardous material is involved.

Axis 3 — ContextualConditions:

- LocationContext: LoadingDock. In this scenario, the loading dock is relevant as a physical location (elevated platform with drop-off), not as a DOT/OSHA jurisdictional boundary, because no hazardous material is present.

Routing resolution:

- OSHA Walking-Working Surfaces fires: 29 CFR 1910.28/1910.29 for fall protection at the dock edge.
- OSHA General Duty Clause fires: Section 5(a)(1) for the dock lock procedure and indicator light failure.
- ANSI fires: Dock safety equipment standards.

Result: {OSHA, ANSI}. No EPA. No DOT. No chemical involvement.

5.9.1 Module C Walkthrough: Recordkeeping Decision

This scenario provides the opportunity to walk the event through the OSHA 300 recordkeeping decision model formalized in Module C.

Gate 1 — Work-Relatedness: The injury occurred in the work environment during a work task (unloading a delivery). No exceptions under 29 CFR 1904.5(b)(2) apply. Work-relatedness is established.

Gate 2 — Recording Criteria: The employee dislocated their arm. Reduction of a dislocation is medical treatment beyond first aid (it does not appear on the exhaustive first aid list in 29 CFR 1904.7(b)(5)(ii)). The MedicalTreatmentCriterion is met. If the employee missed any subsequent workdays, the DaysAwayCriterion is also met. If the employee returned to work with restrictions (limited use of the affected arm), the RestrictedWorkCriterion is met.

Recording determination: Recordable. Entered on the OSHA 300 Log within seven calendar days. Case classification: injury. Contributes to the establishment's TRIR. If days away or restriction occurred, contributes to the DART rate.

OSHA 301 Report: Completed for this case. Captures the event description (employee fell from dock plate when truck pulled away), the object that directly harmed the employee (dock plate/ground), and the nature of the injury (dislocation, upper extremity).

OSHA Reporting Obligation: A dislocation requiring reduction in a hospital does not trigger the 24-hour reporting requirement unless it involves in-patient hospitalization (admission, not just emergency department treatment). If the employee was treated and released from the emergency department, no direct OSHA reporting is required beyond the 300 Log entry.

5.9.2 Module D Walkthrough: Incident Investigation

Severity classification: MedicalTreatmentIncident at minimum. LostTimeIncident if days away occurred. `alignsWithRecordingCriteria` maps the severity to the corresponding recording criterion, confirming consistency between Module D and Module C.

Investigation: Initiated within 24 hours. The investigation team includes the employee's supervisor, an EHS representative, and the affected employee.

Root cause analysis: A 5 Whys or Fishbone analysis reveals two concurrent causal chains:

The first chain addresses the employee's failure to engage the dock lock. Why? The investigation must determine whether the employee relied on the indicator light showing green (safe) and concluded the lock was engaged, or whether the employee habitually skipped the dock lock step. If the employee relied on the false-safe indicator, the root cause shifts from procedural noncompliance to engineering control failure: the indicator system provided false confirmation that removed the behavioral cue to engage the manual backup.

The second chain addresses the indicator light failure. Why did both lights show green simultaneously? In correct operation, the interior and exterior lights should always display opposite states. Both showing green is a condition that should be mechanically impossible, indicating a fundamental failure in the indicator circuit or sensor. The system failed in the most dangerous possible mode: it told the employee it was safe to enter and told the driver it was safe to leave at the same time. The investigation must determine whether the failure was electrical, mechanical (sensor misalignment), or systemic (a design flaw that permitted a dual-green state rather than defaulting to a safe failure mode where any malfunction produces red on both sides).

Corrective actions mapped to Hierarchy of Controls:

The highest-priority corrective action is engineering: redesign the indicator system to eliminate the possibility of a dual-green state (fail-safe design where any malfunction defaults to red/unsafe on both the interior and exterior lights). This maps to `ehs:EngineeringControl` (hierarchyRank 3) via the `addressedByControl` property.

A second engineering control: install an interlock that prevents the dock plate from deploying unless the dock lock is physically engaged, removing the human decision point entirely. This approaches elimination (hierarchyRank 1) by making it mechanically impossible to stand on an unsecured dock plate.

An administrative corrective action: revise the dock loading SOP to require manual verification of dock lock engagement regardless of indicator light status (belt-and-suspenders approach). This maps to `ehs:AdministrativeControl` (hierarchyRank 4). The investigation should document why administrative controls alone are insufficient, since the dual-green indicator failure demonstrated that the existing system could provide false confirmation to both the employee and the truck driver simultaneously, overriding trained behavior on both sides of the dock.

ARECC feedback: The investigation connects to the Confirm phase via `connectsToARECC`. The Confirm phase failed: the existing engineering control (indicator lights) was not confirmed to be functioning correctly through routine inspection or preventive maintenance. Investigation findings feed back into Anticipate (proactively identifying similar false-safe failure modes on other docks), Recognize (adding dock lock indicator testing to the inspection program), and Control (implementing the interlock and fail-safe redesign).

Corrective action tracking: Each action receives an `ehs:CorrectiveActionStatus`. The interlock installation is tracked through Open, In Progress, Completed, and Verified. Verification requires confirming that the interlock physically prevents dock plate deployment without lock engagement under real operating conditions, not just bench testing.

5.10 Summary of Routing Outputs

The table below summarizes the nine scenarios and their regulatory outputs, demonstrating how the routing model produces different obligation sets from different input combinations.

Scenario	HazardType(s)	ActionContext	Key Condition	Agencies Activated
1. Dock spill (uncontained)	Chemical	Transport	FullRelease, AboveRQ	DOT + EPA + OSHA
2. Dock spill (contained)	Chemical	Storage	FullyContained	EPA (RCRA) + OSHA
3. Drum handling	Chemical + Ergonomic	Transfer	—	OSHA + EPA + GHS + ACGIH + NIOSH + ANSI
4. Roadway release	Chemical	Transport	PublicRoadway, FullRelease	DOT + EPA + USCG + OSHA (responders only)
5. Healthcare maintenance	Biological + Chemical	Maintenance	ConfinedSpace	OSHA (4 standards) + EPA + NIOSH
6. Tank sensor repair	Electrical + Chemical	Maintenance	ConfinedSpace	OSHA + NFPA + ANSI + GHS + ACGIH
7. Conveyor maintenance	Electrical + Mechanical	Maintenance	—	OSHA + NFPA + ANSI
8. Night shift stamping	Physical + Psychosocial	Processing	Chronic exposure	OSHA + ACGIH + ISO + NIOSH
9. Dock plate failure	Physical + Mechanical	Transport	LoadingDock	OSHA + ANSI

Scenarios 1 and 2 demonstrate that a single ContextualCondition (ContainmentStatus) can reduce a three-agency cascade to a single-agency disposal event. Scenarios 3, 5, 6, and 9 demonstrate that CompoundHazardProfiles produce additive regulatory obligations. Scenarios 7 and 8 demonstrate that the routing model operates independently of chemical hazards. Scenario 9 demonstrates the integration of the core routing model with Module C (recordkeeping) and Module D (incident investigation), tracing the event from the initial fall through recording determination, root cause analysis, corrective action, and ARECC feedback.

The ontology produces the correct regulatory output in all nine scenarios. The routing is deterministic: given the same HazardType, ActionContext, and ContextualConditions, the same regulatory obligation set results. The practitioner's judgment is encoded in the classification of inputs, not in the routing logic itself.

6 Discussion

The ontology presented in this paper formalizes a body of professional knowledge that has, until now, existed only in the judgment of experienced EHS practitioners. This section discusses what that formalization enables, presents a designed extension that addresses the ontology's most significant scope limitation, identifies remaining limitations, and outlines future work.

6.1 What the Ontology Enables

6.1.1 Automated Compliance Routing

The most immediate application is automated determination of regulatory obligations from classified inputs. The three-axis routing model (HazardType × ActionContext × ContextualCondition) produces a deterministic output: given a classified hazard, a defined action, and a set of conditions, the ontology returns the complete set of applicable regulatory frameworks and specific CFR citations. A software system built on this ontology could accept structured input from an EHS practitioner (a chemical spill occurred during truck offloading at the loading dock, containment failed, the quantity exceeds the Reportable Quantity) and return the full obligation set: DOT 49 CFR 171.15, EPA CERCLA 103(a), EPCRA Section 304, EPA RCRA 40 CFR 262, OSHA HAZWOPER 29 CFR 1910.120, and OSHA HCS 29 CFR 1910.1200.

This does not replace practitioner judgment. The practitioner classifies the inputs. The ontology routes them. The value is that the routing logic is explicit, auditable, and consistent rather than dependent on the individual practitioner's experience and memory. A junior EHS professional with two years of experience and a senior practitioner with twenty years of experience receive the same routing output for the same classified inputs.

6.1.2 Training and Professional Development

The ontology's explicit structure creates opportunities for training that current EHS education struggles to deliver. EHS degree programs teach the 5 E's, the seven IH concepts, ARECC, and the Hierarchy of Controls as separate frameworks. They teach OSHA, EPA, and DOT regulations in separate courses. The integration of these frameworks, the understanding of how a hazard type classification activates regulatory obligations across independent agencies simultaneously, is left to on-the-job experience accumulated over years.

The ontology makes this integration visible and teachable. A training system built on the ontology could present a scenario, ask the trainee to classify the inputs (What is the hazard type? What was the action? What were the conditions?), and then reveal the complete routing output. The trainee could modify a single condition (change ContainmentStatus from FullRelease to FullyContained) and observe the regulatory cascade collapse from three agencies to one. This kind of interactive, consequence-driven learning is possible only when the underlying compliance logic is formalized.

6.1.3 Knowledge Preservation

The most significant long-term application may be the least visible. When a senior EHS practitioner with decades of experience retires, the organization loses more than a person. It loses the integrated understanding of how regulations interact, which edge cases matter (the volatile vapor exception to the containment exemption, the carrier-presence requirement for DOT jurisdiction at the loading dock), and how facility-specific conditions modify standard regulatory obligations.

This knowledge currently transfers through mentorship, apprenticeship, and institutional memory. These transfer mechanisms are slow, incomplete, and vulnerable to disruption. The ontology provides a formal vessel for this knowledge. The routing logic, the edge cases, the interaction notes on compound hazard profiles, and the contextual conditions that gate agency jurisdiction can be captured in a structure that persists independent of any individual practitioner. The ontology does not replace experience. It prevents the loss of the structural knowledge that experience produces.

6.2 The Geo-Compliance Extension: From Federal Baseline to Facility-Specific Routing

The most significant scope limitation of the v3.1 ontology is that it models only federal regulatory frameworks. In practice, EHS compliance operates across a minimum three-tier regulatory stack: federal, state, and county/local. A manufacturing facility in Oakland, California simultaneously faces federal OSHA (29 CFR 1910), Cal/OSHA (Title 8 CCR), and City of Oakland industrial hygiene ordinances. The same facility's air emissions are governed by EPA (40 CFR Part 70), the California Air Resources Board (CARB), and the Bay Area Air Quality Management District (BAAQMD), which serves as the primary permitting authority for stationary sources in the nine-county Bay Area. Its chemical inventory is reported under federal EPCRA Tier II at a 10,000-pound threshold, California's Hazardous Materials Business Plan (HMBP) at a 55-gallon or 500-pound threshold, and the Alameda County Certified Unified Program Agency (CUPA). Three tiers of government, three different thresholds, three different submission systems, for the same chemical at the same facility.

To address this gap, I have designed a Geo-Compliance Extension (ehs-geo:) that adds a fourth routing axis, FacilityJurisdiction, to the existing three-axis model. The extension is architecturally complete as a design document and is summarized here.

6.2.1 Four-Axis Routing

The extended routing model processes compliance activation in four ordered steps:

1. HazardType(s) fire the federal framework baseline.
2. ActionContext adds operational modifiers.
3. ContextualConditions refine scope.
4. FacilityJurisdiction applies the geographic overlay.

Step 4 does not replace Steps 1 through 3. It accepts the output of `ehs:ContextualComplianceActivation` and augments it with sub-federal regulatory obligations. The extension lives in a separate OWL namespace (ehs-geo:) and imports the v3.1 ontology without modifying any existing class or property. This preserves backward compatibility: the federal routing model functions identically with or without the extension loaded.

6.2.2 Jurisdictional Layer Architecture

The extension models sub-federal authority through `ehs-geo:JurisdictionalLayer`, an abstract superclass with three concrete tiers:

StateJurisdiction captures state-level regulatory authority through three subclasses. `ehs-geo:OSHAStatePlan` models the 22 states and territories that operate OSHA-approved State Plans with standards that must be "at least as effective" as federal OSHA but may be more stringent (California's Heat Illness Prevention standard, Oregon's ergonomic standard, and Virginia's Heat Illness standard have no federal equivalent). `ehs-geo:StateEnvironmentalProgram` models state EPA equivalents that administer delegated federal programs (CAA, CWA, RCRA) with state-specific additional requirements. `ehs-geo:StateFireMarshalProgram` models state-level fire code adoption and enforcement.

CountyJurisdiction captures county-level authority through three subclasses. `ehs-geo:CountyAirDistrict` models the 130+ Air Quality Management Districts and Air Pollution Control Districts that serve as primary air permitting authorities in many regions, often with emission thresholds lower than federal major source levels. `ehs-geo:CertifiedUnifiedProgramAgency` models the California CUPA system, where a single county-level agency administers six environmental and safety programs (HMBP, California Accidental Release Prevention [CalARP], underground storage tanks, hazardous waste generator, on-site treatment, aboveground petroleum storage) under a unified permit. While California-specific, the CUPA class models a general pattern of unified local program agencies that exists in other states under different names. `ehs-geo:CountyHealthDepartment` models county public health authority over environmental health programs.

CityJurisdiction captures municipal authority through three subclasses. `ehs-geo:CityFireDepartment` models the Authority Having Jurisdiction (AHJ) for local fire code enforcement, including hazardous materials permits at thresholds that often differ from EPCRA. `ehs-geo:CityBuildingDepartment` models occupancy classification and fire suppression requirements. `ehs-geo:CityEnvironmentalOffice` models city-level chemical use ordinances that exceed state and federal minimums, as adopted by cities such as San Francisco, Chicago, and Seattle.

6.2.3 Conflict Resolution

When multiple jurisdictional tiers impose requirements on the same regulated subject, the extension models three relationship types as first-class OWL classes rather than free-text annotations. `ehs-geo:StatePreempts` indicates that state law expressly prohibits local regulation in a subject area. `ehs-geo:LocalSupplements` indicates that local ordinances add requirements on top of the federal and state floor (both apply; neither cancels the other). `ehs-geo:MostStringentApplies` indicates that when multiple tiers set different numerical standards for the same subject, the most protective requirement governs. Modeling these relationships as typed classes enables SPARQL queries to automatically filter and resolve conflicting requirements based on jurisdiction.

6.2.4 Jurisdiction Resolution and Integration

A five-step jurisdiction resolution algorithm determines the complete set of applicable layers from an establishment's physical address: geocoding to coordinates, state layer lookup (OSHA State Plan status, state environmental agency, state right-to-know programs), county layer lookup (air district, CUPA, health department), city layer lookup (fire department permits, building department, environmental office), and composition into a single `ehs-geo:FacilityJurisdiction` instance linked to the establishment via `ehs-geo:hasJurisdiction`.

The extension integrates with the v3.1 ontology through two connection points. The `ehs-geo:hasJurisdiction` property links `ehs:Establishment` to its resolved FacilityJurisdiction. The `ehs-geo:hasGeoActivation` property links `ehs:ContextualComplianceActivation` to the sub-federal `ehs-geo:GeoComplianceRequirement` instances that also apply. These two properties are the only additions to the existing ontology's class structure.

6.2.5 Worked Example

The extension's design document includes worked scenarios that demonstrate the geographic overlay in practice. An automotive paint shop in Wayne County, Michigan receives the federal baseline (OSHA HCS, EPA Title V, ACGIH TLVs) plus state-level overlays from Michigan Occupational Safety and Health Administration (MIOSHA) Part 301 (Michigan-specific PELs) and the Michigan Department of Environment, Great Lakes, and Energy (EGLE) Air Quality Division (state Reasonably Available Control Technology [RACT] requirements for VOC coating operations). The same facility relocated to Oakland, California receives a substantially different overlay: Cal/OSHA Title 8 CCR § 5155 (stricter PELs), Cal/OSHA § 3395 (Heat Illness Prevention, no federal equivalent), BAAQMD Regulation 8 Rule 16 (spray coating VOC limits stricter than federal NESHAP), Alameda County CUPA HMBP (55-gallon threshold versus 10,000-pound EPCRA threshold), and City of Oakland fire prevention hazmat permits. The federal routing output is identical for both locations. The geographic overlay produces materially different compliance obligations.

6.2.6 Phased Implementation

The extension is designed for incremental population across six phases: State OSHA Plan status (low effort, immediate value), state environmental programs (medium effort), major county air districts (medium effort, covering approximately 80% of industrial facilities in nonattainment areas), California CUPA program (medium effort, critical for California facilities), city fire department hazmat permit thresholds (high effort, wide variation), and city-level chemical use ordinances (high effort, limited to progressive municipalities). Data population draws from three source types: static lookup tables for state-level data, API-driven resolution

(Census Geocoder, EPA Facility Registry System, EPA Enforcement and Compliance History Online [ECHO]) for county and city layers, and expert-curated requirement instances for high-impact local regulations.

The Geo-Compliance Extension is architecturally complete as a design document. Full implementation and population represent a substantial effort that merits its own dedicated publication.

6.3 Limitations

This ontology has several limitations that must be acknowledged.

6.3.1 Validation

The nine scenarios presented in Section 5 demonstrate that the routing logic produces correct outputs for those specific input combinations. They do not constitute exhaustive validation. The ontology has not been tested against a comprehensive corpus of real-world EHS incidents to verify that its HazardType, ActionContext, and ContextualCondition taxonomies capture the full range of situations practitioners encounter. Categories that the current taxonomy may underrepresent include agricultural operations, maritime environments, mining (beyond the MSHA framework class), and laboratory research settings.

6.3.2 Ontology Versus Implementation

The ontology is a formal knowledge model, not a software application. It defines classes, properties, and relationships in OWL. It does not include inference rules, SPARQL queries, SWRL rules, or SHACL constraints that would automate the routing logic computationally. A practitioner cannot load the TTL file into a reasoner and receive compliance determinations. Building the inference layer that transforms the ontology from a knowledge representation into a functioning compliance routing engine is a substantial engineering effort that lies outside the scope of this paper. Relational schema prototyping is underway for Modules A (EPCRA/TRI) and C (OSHA 300 Recordkeeping), translating the ontology's class structure into auditable database models that separate events from decisions from records. This work is ongoing and its results are not yet reported here.

6.3.3 Static Regulatory Modeling

Regulations change. OSHA issues new standards, EPA modifies reporting thresholds, chemicals are added to or removed from the TRI list (as the MEK delisting in the ontology demonstrates), and court decisions reinterpret enforcement boundaries (as the Tampa Electric HAZWOPER decision illustrates). The current ontology represents a point-in-time snapshot of the regulatory landscape. It does not include a versioning mechanism for tracking regulatory changes, a change management process for updating the ontology when regulations are amended, or a provenance model for documenting why specific regulatory mappings were made. A production implementation would require all three.

6.3.4 Classification Dependence

The routing model produces deterministic outputs from classified inputs, but the classification itself requires practitioner judgment. Is a given situation a Chemical hazard, a Physical hazard, or both? Is the action context TransportAction or StorageAction when drums are being moved from a truck to a storage area? Is the containment FullyContained or PartialBreach when a small amount of vapor escapes but no liquid does? These classification decisions carry the same professional judgment that the ontology seeks to formalize at the routing level. The ontology shifts judgment from "which regulations apply?" to "how do I classify this situation?" This is a genuine improvement in consistency, but it is not the elimination of subjective decision-making.

6.4 Future Work

Several directions extend naturally from this foundation.

The most immediate next step is the development of inference rules (SWRL or SHACL) that automate the routing logic. The class structure and property relationships defined in the ontology provide the schema. The inference rules would provide the engine. This work would transform the ontology from a knowledge representation into a queryable compliance routing system.

Implementation and population of the Geo-Compliance Extension represents the path to practical adoption. A pilot targeting a single state with both a State OSHA Plan and delegated EPA programs (Michigan and California are the strongest candidates given the worked scenarios in the extension design document) would test whether the v3.1 architecture accommodates state-level overlays without structural modification.

Cross-linking the ontology with existing chemical knowledge bases would increase its practical value. The GHS hazard classifications formalized in (Pascazio et al. 2023) and (Medici et al. 2022) provide a natural integration point: if a chemical's GHS classification is known, the HazardType classification (and therefore the baseline regulatory activation) can be partially automated.

Validation against a larger corpus of real-world EHS incidents, drawn from OSHA inspection records, EPA enforcement actions, and National Transportation Safety Board (NTSB) transportation incident reports, would test the routing model's coverage and identify gaps in the current taxonomies.

Finally, the ontology's training applications could be evaluated through a controlled study comparing the performance of EHS trainees who learn compliance routing through ontology-based interactive scenarios against trainees who learn through traditional curriculum. The hypothesis is that making the routing logic explicit and interactive accelerates the development of the integrated multi-agency understanding that currently requires years of practice to acquire.

7 Conclusion

EHS practitioners perform a complex intellectual task that no existing ontology has formalized: routing a single workplace event through the overlapping jurisdictions of multiple independent federal agencies, where hazard type, operational context, and contextual conditions determine which regulatory frameworks activate and which obligations apply. This paper presented an ontology that captures that routing logic as a formal system.

The ontology is structured around the Employee-Hazard bipolar model, mediated by the 5 E's of Safety and assessed through the ARECC decision-making framework. Its central contribution is the treatment of hazard type classification as a compliance activation mechanism rather than a descriptive label. The three-axis routing model (HazardType × ActionContext × ContextualCondition) produces deterministic regulatory outputs from classified inputs. CompoundHazardProfiles formalize the additive nature of compliance: when multiple hazard types co-occur, the resulting obligation is the union of all constituent activations. Four regulatory modules (EPCRA chemical inventory and TRI reporting, Clean Air Act Title V permitting, OSHA 300 recordkeeping, and incident management) extend the core model into specific compliance domains, wired together through shared classes and cross-linking properties that reflect how regulatory obligations interconnect in practice.

Nine worked scenarios validated the routing logic across chemical, biological, electrical, mechanical, physical, ergonomic, and psychosocial hazard types, in contexts ranging from acute spills to chronic noise exposure to a real-world dock plate failure drawn from the author's professional practice. The scenarios demonstrated that the same chemical produces different regulatory outputs based solely on contextual conditions, that compound hazard profiles produce additive obligations, and that the routing model operates independently of any single hazard domain.

The gap between this ontology and a production compliance tool is real. The ontology defines the knowledge structure. It does not yet include the inference rules, the state and local regulatory overlays, or the change management processes that a functioning system requires. What it does provide is a formal, transferable model of the integrated multi-agency compliance knowledge that currently exists only in the professional judgment of experienced practitioners. That knowledge has never been captured in a structure that can be audited, taught, shared, and preserved independent of the individuals who hold it.

This paper is a first step. The architecture is mapped. The formalization is built. The routing logic is validated. What remains is the engineering work to make it operational and the empirical validation to prove it comprehensive.

Dimyadi, Johannes, Pieter Pauwels, and Robert Amor. 2016. "Modelling and Accessing Regulatory Knowledge for Computer-Assisted Compliance Audit." *Journal of Information Technology in Construction* 21: 317–36. http://www.itcon.org/2016/21.

Gallina, Barbara, Gergő László Steierhoffer, Thomas Young Olesen, Eszter Parajdi, and Mike Aarup. 2025. "Towards an Ontology for Process Compliance with the (Machinery) Legislations." *Journal of Software: Evolution and Process* 37 (1): e2728. https://doi.org/10.1002/smr.2728.

Gill, Navdeep Singh. 2026. *ElixirClaw-ElixirData Manufacturing Use Cases: AI Agents.* ElixirData Blog. https://www.elixirdata.co/blog/elixirclaw-elixirdata-manufacturing-use-cases-ai-agents.

Ławrynowicz, Agnieszka, and Ilona Ławniczak. 2016. "Towards a Core Ontology of Occupational Safety and Health." In *Ontology Engineering*, edited by Valentina Tamma, Mauro Dragoni, Rafael Gonçalves, and Agnieszka Ławrynowicz. Springer International Publishing.

Medici, Marco, Damion Dooley, and Maurizio Canavari. 2022. "PestOn: An Ontology to Make Pesticides Information Easily Accessible and Interoperable." *Sustainability* 14 (11): 6673. https://doi.org/10.3390/su14116673.

Mortensen, Holly M., Jaleesia D. Amos, Thomas E. Exner, et al. 2024. "NNI Nanoinformatics Conference 2023: Movement Toward a Common Infrastructure for Federal nanoEHS Data Computational Toxicology: Short Communication." *Computational Toxicology* 30 (June): 100316. https://doi.org/10.1016/j.comtox.2024.100316.

Onut Badea, Daniel, Victoria-Rodica Cioca, Doru Costin Darabont, Timur Vasile Chis, and Raluca Maria Iordache. 2024. "Ontology-Based Occupational Safety and Health Management for Workers with Disabilities." *Polish Journal of Management Studies* 30 (1): 24–41. https://doi.org/10.17512/pjms.2024.30.1.02.

Oyasiji, Odunayo, Adeola Okesiji, Adeyinka Lawal, Bisayo Oluwatosin Otokiti, and Sibongile Gobile. 2024. "Developing Conceptual AI Models for Legal Text Interpretation and Regulatory Compliance Automation." *International Journal of Multidisciplinary Research and Growth Evaluation* 5 (2): 1098–104. https://doi.org/10.54660/.IJMRGE.2024.5.2.1098-1104.

Park, Hyunsoung, and Sangyun Shin. 2025. "Development of Safety Domain Ontology Knowledge Base for Fall Accidents." *Buildings* 15 (13): 2299. https://doi.org/10.3390/buildings15132299.

Pascazio, Laura, Simon Rihm, Ali Naseri, Sebastian Mosbach, Jethro Akroyd, and Markus Kraft. 2023. "Chemical Species Ontology for Data Integration and Knowledge Discovery." *Journal of Chemical Information and Modeling* 63 (21): 6569–86. https://doi.org/10.1021/acs.jcim.3c00820.

Sanchez-Pi, Nayat, Luis Martí, Jose Manuel Molina, and Ana Cristina Bicharra Garcia. 2015. "Ontology Definition and Cognitive Analysis in Ocupational Health and Security (OHS) Environments." *Proceedings of the 30th Annual ACM Symposium on Applied Computing* (Salamanca, Spain), April, 201–6. https:

//doi.org/10.1145/2695664.2695891.

Single, Johannes I., Jürgen Schmidt, and Jens Denecke. 2020. "Computer-Aided HAZOP: Ontologies and AI for Hazard Identification and Propagation." In *Computer Aided Chemical Engineering*, vol. 48. Elsevier. https://doi.org/10.1016/B978-0-12-823377-1.50298-6.

Tan, X. C., K. H. Yew, and T. J. Low. 2012. "Ontology Design for Process Safety Management." *2012 International Conference on Computer & Information Science (ICCIS)* 1 (June): 114–19. https://doi.org/10.1109/ICCISci.2012.6297223.

United Nations. 2023. *Globally Harmonized System of Classification and Labelling of Chemicals (GHS), Tenth Revised Edition.* United Nations.

U.S. Environmental Protection Agency. 1992. *SARA Orientation Manual for DOE Facilities.* EPA/542/R-92/005. Office of Solid Waste; Emergency Response.

Vigneron, A, F Guarnieri, and B Rallo. 2013. "The Contribution of Ontologies to the Creation of Knowledge Bases for the Management of Legal Compliance in Occupational Health and Safety." In *Safety, Reliability and Risk Analysis*, edited by R Steenbergen, P Van Gelder, S Miraglia, and A Vrouwenvelder. CRC Press. https://doi.org/10.1201/b15938-230.

Zhang, Zhuoqun, Chunhui Zhao, Danli Li, et al. 2026. "Human-in-the-Loop Semantic Rule Base Generation and Dynamic Updating for Automated BIM Compliance Checking: A Knowledge Graph Approach." *Buildings* 16 (4): 719. https://doi.org/10.3390/buildings16040719.