HHAI 2025
D. Pedreschi et al. (Eds.)

doi:10.3233/FAIA250646

Ontology-Based Risk Assessment in Smelting Plant Logistics

John GRANSTRÖM [1] and Andreas BRÄNNSTRÖM
Department of Computing Science
Umeå University, SE-901 87, Umeå, Sweden
ORCiD ID: John Granström https://orcid.org/0009-0007-1151-4065, Andreas Brännström https://orcid.org/0000-0001-9379-4281

Abstract. This paper presents an ontology-based approach to proactive risk assessment in high-risk industrial environments, with a focus on molten metal handling in smelting plants. Distributed cognition theory serves as the theoretical framework, guiding the grounded theory and thematic analysis by framing collaboration and communication as emerging from interactions among operators, equipment, and the environment. This lens emphasized shared knowledge and distributed responsibilities, helping identify domain-relevant concepts critical for risk assessment. Expert insights from professionals at a Swedish smelting plant were systematically elicited and structured into an ontology to support proactive risk management. A prototype system integrating this ontology with a dynamic interface showed strong alignment with expert decision-making and safety assessments. Moreover, the system identified overlooked risks—such as hazardous equipment containing molten metal—and received positive user feedback. While tailored for smelting operations, the methodology has broader applicability to improving information sharing, decision-making, and safety in other socio-technical systems.

Keywords. ontologies, distributed cognition theory, knowledge representation, risk assessment, industrial safety

1. Introduction

In high-stakes industrial environments, managing the logistics of hazardous materials is a complex, critical process that leaves little room for error. At a major smelting plant in Sweden (name withheld for confidentiality), operators and drivers handle molten metals at extreme temperatures, with even minor oversights potentially resulting in catastrophic incidents, such as explosions due to mishandled or improperly prepared ladles. This environment exemplifies the challenges that arise in high-risk settings, where effective, proactive risk assessment is essential to maintaining safety and operational efficiency—challenges that are also found in similar high-risk domains [1,2,3,4]. Achieving this level of safety requires not only a robust technical solution but also seamless collaboration across roles. Each stakeholder—operators, drivers, supervisors—contributes specific knowledge and perspective, from real-time positional awareness of ladles to the

[1] Corresponding Author: John Granström, johngranstrom@gmail.com

scheduling and procedural knowledge essential for managing operations. The expertise and responsibilities distributed among these roles are interdependent and gaps in communication or knowledge sharing can amplify risks. A successful solution must therefore capture, integrate and continuously update the various perspectives and inputs of each stakeholder, enabling the system to support coordinated, informed decisions across roles.

Traditionally, risk management in such dynamic settings has relied on statistical models [5,6,7,8] or rule-based systems [9,10,11,12] to predict and mitigate potential hazards. Recently, machine learning has emerged as a promising tool for risk assessment in various domains [13,14,15], offering advantages in identifying general patterns and predicting outcomes based on large datasets. However, while machine learning solutions provide benefits in some domains, they often fall short in high-stakes environments due to their dependence on extensive historical data [16,17,18]. In contexts like molten metal handling, where incidents must be minimized and data on rare events is limited, machine learning models struggle to provide proactive insights for situations not previously encountered. Furthermore, these models often lack the flexibility to adapt to fast-changing conditions or to capture the nuanced interdependencies within complex systems.

To enable proactive risk assessment, a solution must capture three essential aspects: (1) the entities involved, which include both physical entities, such as equipment and operator roles, and non-physical entities, such as information flows, tasks, and potential risks; (2) the relationships between these entities, which include procedural dependencies, communication flows, and causal links to risk factors that could compromise safety; and (3) the rules that govern interactions, specifying safety protocols, operational guidelines, and conditions under which certain risk causes may be triggered. Together, these aspects create a foundation for a system that can anticipate potential risks, interpret causal factors, and support informed, proactive decision-making in real-time. A computational method for modeling these requirements is the concept of an Ontology [19] and Semantic Web Rule Language (SWRL) [20].

While ontologies have proven effective for risk assessment in diverse fields, such as manufacturing systems [21], cyber-physical systems [22], construction safety [23], fault detection [24], and traffic management [25], their application to the specific challenges of molten metal handling, particularly regarding distributed knowledge and coordination needs, remains unexplored.

Our approach to addressing these requirements followed a systematic knowledge engineering process: 1) Acquiring domain understanding through on-site observations and interviews with experts; 2) Qualitative analysis through the lens of Distributed Cognition Theory [26]; 3) Ontology modeling and development; 4) Implementation and evaluation of a proof-of-concept prototype.

This research uniquely utilizes distributed cognition theory as a theoretical framework to develop an ontology tailored to smelting plant operations' risks and operational complexities. Through an in-depth knowledge elicitation process—including interviews, observations, and scenario analyses [27] we capture and encode essential knowledge into an ontology that reflects the critical entities, relationships, and rules within this environment.

An important aspect of our approach is the development of a digital twin [28], where our ontology-based knowledge representation captures not just data but also the relationships and rules governing the molten metal handling system. In this case, the digital twin represents entities like ladles, slag trucks, drivers, operators and other critical

components of the process, offering a dynamic, real-time digital model. This enables logical analysis of processes and risks, providing proactive support through an ontology-driven knowledge base, while laying the groundwork for future integration with various learning systems and sensors.

The rest of this paper is structured as follows. In Section 2, related work on ontology-based solutions in industrial safety systems are covered. In Section 3, the case study is presented. In Section 4, the knowledge elicitation and modeling process is presented. In Section 5, the resulting knowledge model—ontology—is presented. In Section 6, the evaluation of a proof-of-concept prototype is presented. Section 7, the limitations, strengths, and broader applicability of the approach is discussed. In Section 8, the paper concludes with future research directions.

2. Related Work

Ontology-based approaches have been widely used for safety and risk assessment. Ahmad et al. [29] and Kostov et al. [30] applied ontologies to aviation safety, improving accident analysis and safety information management. Similarly, Bhosale et al. [31] integrated safety-security risk assessment for industrial control systems, while Ding et al. [32] and Gao et al. [23] developed ontology-driven safety management frameworks in construction. In supply chain and manufacturing, Aboutorab et al. [33] highlighted risk identification techniques, and Abadi et al. [21] utilized digital twins and ontologies for maintenance optimization. Wu et al. [34] proposed a Bayesian network-based safety integrity evaluation for high-temperature molten metal operations, and Wu et al. [35] applied fuzzy CREAM and Bayesian networks to analyze human reliability in such operations. Salama et al. [36] utilized SWRL rules in a rule-based recommendation system, demonstrating their capability in automated reasoning. Reda et al. [37] leveraged SWRL rules for smart home scenarios, showcasing their utility in context-aware decision-making. The application of distributed cognition theory [26,38,39] in industrial risk assessment settings has been limited.

Distributed cognition theory, explored by Liu et al. [38] and Lindblom and Thorvald [40], remains underutilized in industrial risk assessment. Our approach introduces the first ontology-based proactive risk assessment model for molten metal handling. Our approach distinguishes itself by combining distributed cognition theory with expert knowledge and ontology-driven risk assessment, ensuring a practical and context-aware solution for molten metal handling. Unlike prior work, our methodology bridges theoretical ontology development with real-world application, capturing implicit knowledge and complex risk factors often overlooked in generalized frameworks.

3. Case Study

This research centers on a smelting plant in Sweden, where molten metals at temperatures approximately 1250°C are handled using large ladles, which are transported by specialized slag trucks (Figure 1a). The process involves three types of ladles: a drip ladle positioned under the furnace during melting to catch any spillage, a slag ladle for handling the residual slag byproduct, and a lead ladle for the final processed metal. Be-

fore lead extraction, the residual slag—which has a lower density than lead and thus accumulates at the top inside the furnace—must be carefully transferred to designated slag ladles. These handling procedures require precise adherence to safety protocols, as improper management presents severe explosion risks at two critical points: during emptying (endangering drivers) and pouring (endangering operators). When unsolidified slag contacts water from rain or snow, it can trigger dangerous steam explosions. To mitigate this risk, a mandatory 72-hour waiting period ensures complete slag solidification before ladles can be emptied and reused. This safety measure addresses explosion risks from unsolidified slag and residual moisture, with multiple serious incidents occurring.

The current solidification process relies on manual operations with critical information maintained on a physical whiteboard (Figure 1b), accessible only to operators. Drivers, while aware of ladle placement in real-time, lack access to this broader operational overview. This division of cognitive labor creates a system vulnerable to communication gaps and errors, particularly dangerous given the high-stakes nature of molten metal handling.

(a) Slag truck with a ladle [1]

(b) Current solution to track ladles[2]

Figure 1. Case environment.

4. Knowledge Elicitation

An anonymous interview study was conducted followed by a thematic analysis. Distributed Cognition Theory served as theoretical framework for the knowledge elicitation process[3]. This theory posits that cognitive processes are distributed across multiple elements in the environment, including tools, individuals, and practices. In the context of molten metal handling, this framework encompasses human agents (operators and drivers), artifacts (physical objects that play a role in cognitive processes), tools (ladles, trains, monitoring devices), and information flows between these components. The theory considers how knowledge and responsibilities are shared among stakeholders, making it particularly relevant for understanding complex industrial processes where information distribution impacts safety and efficiency.

[1] Image adapted and used with permission.

[2] Photographed and used by author with permission.

[3] Material related to the knowledge elicitation process, including an ethical assessment, interview transcripts, question formulations, evaluation and analysis frameworks, is available at https://github.com/RA-SPL/Ontology-RA-SPL; With this, we aim to support transparency and reproducibility.

The primary data was gathered through interviews with over 20 professionals at the plant, from whom their thoughts and feedback on the system were recorded. A detailed overview of the participants' demographics is presented in Table 1. Among these, four in-depth interviews were conducted, two with operators and two with slag truck drivers. These interviews, lasting between 1.5 and 3 hours each, were significantly more detailed and focused, serving as the foundation for the development of the ontology later created.

Role	Count	Age Range	Experience (years)	Gender
Operators	12	25-55	3-30	12M
Slag Truck Drivers	7	30-50	5-20	5M, 2F
Supervisors	3	40-63	15-30	2M, 1F
Maintenance Staff	2	35-57	10-15	2M

Table 1. Participant Demographics

The operators, with 20 and 7 years of experience respectively, provided in-depth insights into control room operations and safety protocols. The slag truck drivers, with 8 and 3 years of experience in molten metal transportation, contributed highly relevant logistical knowledge. The interviews consisted of 20 questions[3], covering background and experience (1-2), operational challenges before, during, and after pouring (3–6), critical information needs for ladle management (7–11), collaboration and safety concerns (12–13), system requirements and usability preferences (14–20), with a focus on safety, efficiency, and information flow. The answers were transcribed and analyzed from multiple perspectives using distributed cognition theory as the theoretical framework. In addition to the interviews, direct observations of the processes were conducted, and the broader feedback gathered from the other participants was considered. Together, these elements played a crucial role in shaping the system's ontology and deepening the understanding of the operational context.

4.1. Knowledge Capture and Formalization

To illustrate the knowledge elicitation methodology, consider the response to Question 5: "What information is most critical to have available about a slag ladle?" Driver #1's answer provides key insights into the critical information needed for safe slag ladle handling. The response, highlighting key domain concepts (marked in `monospace`), will serve as a running example throughout the paper to demonstrate how our methodology transforms raw expert knowledge into formal ontological structures and ultimately into executable system rules. Statement by Driver #1:

> "Ensuring that the `[ladle]` ends up in the right place (at the `[solidification site]`). If its placed on the wrong `[area]` the next driver could take that (`[unsolidified] [ladle]`) and empty it, if that happens the consequence could be catastrophic. If the `[slag]` has not been `[solidified]` for at least `[72 hours]`. At the moment we have no way of checking this, the `[information]` is only available to the `[operators]` in their control room. We have to call and ask if the next `[ladle]` with a specific `[number]` is actually `[solidified]` for the actual time. And I know that not every `[driver]` calls, they trust that next one is safe."

This response illustrates the type of domain knowledge gathered, highlighting key concepts that informed our ontological model, from physical entities (`ladle`, `area`) to process requirements (`72 hours`) and safety constraints (`solidified`). As we proceed

through our methodology, we will show how these concepts are formalized into classes, properties, and rules in our final system.

Through the lens of distributed cognition theory, we analyzed each transcribed response through six key categories that reveal how information is distributed and transformed throughout the system: `Roles` (identifying key actors), `Tasks` (describing their actions), `Tools` (physical artifacts used), `Information` (critical data points), `Rules` (governing protocols), and `Relationships` (interactions between components).

The analysis identified two key roles: `Operator`, coordinating from the control room, and `Driver`, managing slag ladle transport to the solidification site. The tools category highlighted the ladles themselves, while information elements encompassed critical data like solidification times, ladle numbers, and area assignments. Rules emerged around safety protocols, such as the 72-hour solidification requirement and proper area placement, with relationships capturing the vital yet problematic communication patterns between operators and drivers. The analysis also revealed critical risks, specifically one driver's concern that "[...] if that happens the consequence could be catastrophic" regarding unsolidified slag.

Table 2. Subset of concepts derived through the knowledge elicitation process.

Category	Concepts
Roles	Operator, Slag truck driver
Tasks	Lubrication, Visual Inspection, Emptying of ladle, Emptying of Lead Ladle, Emptying Slag Ladle, Lead Pouring at refinery, Pouring of lead, Slag pouring, Clear gas duct, Cooling of lead, Cooling OK, Lead Pouring Done, Leave Slag Ladle, Pickup Slag ladle, Slag Solidification, Maintenance, Reparations, Inspection, Milling, Sandblasting, Furberg Garage [...]
Tools	Ladle, Lead Ladle, Slag Ladle, Drip Ladle, Train [...]
Information	Time for solidification, Temperature of Ladle, Temperature of Lead, Temperature of Slag, Slag Ladle not to be emptied, Fleetskill Order Pouring, Solidification Site, Leave on area, Get on area, Area, Level of molten metal, LadleNr, Date, Time, Shift [...]
Risks	Slag ladle not empty, Liquid in Slag ladle, Ladle not OK to lift, Hole in ladle, Fire Hazard, Explosion Risk, Smoke Hazard, Loss of Production, Slag Ladle not warm [...]

To transform this distributed cognition analysis into a formal ontological structure, we performed further aggregation of the experts' insights. This process involved distilling the categorized information into three key ontological components: classes (and subclasses), object properties (relationships between classes),and data properties (attaches a class to a datatype value).

Analyzing participants' descriptions of, e.g., ladle placement and types, we identified classes and their hierarchical relationships, with `Ladle` as a parent class and `SlagLadle` as a specialized subclass. The described communication needs between operators and drivers informed our object properties, such as `hasLadleInArea` connecting ladles to their locations. The specific requirements mentioned— "[...] solidified for at least 72 hours" —led to data properties like `hoursSolidified`, while the area and ladle identification systems mentioned became `AreaNr` and `LadleNr` properties. This transformation from the participants' practical knowledge into formal ontological structures provided the foundation for developing a comprehensive knowledge model capable of supporting automated reasoning about the safety and operational decisions that the participants' highlighted as critical. Refer to Table 2 for a selection of identified concepts categorized under each main category.

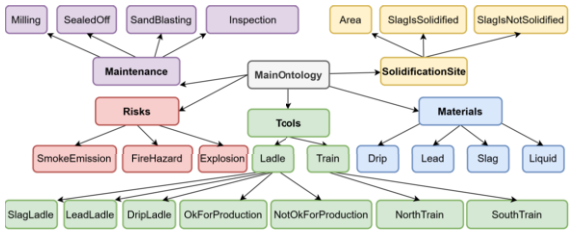

Figure 2. Classes and subclasses with relationships in the ontology.

5. Ontology and System Architecture

The following section presents the ontological framework that emerged from our knowledge elicitation process and demonstrates how it enables automated reasoning about safety-critical situations in ladle management.

5.1. Ontological Framework

Our ontology captures the distributed knowledge and safety protocols identified during the knowledge elicitation. The framework developed in the ontology editor Protégé centers around six main conceptual areas that emerged as critical for ladle management: `Tools` (including trains where the ladles are delivered to the driver and various ladle types), `Materials` (encompassing lead, slag, and potential contaminants), `SolidificationSite` (representing areas where slag ladles solidify), `Maintenance` procedures, and `Risks`. Figure 2 presents this ontology with its concepts, relationships, and hierarchies.

5.2. Description Logic and Reasoning

The ontology's power lies in its ability to encode and reason about complex safety rules through its structural components and logical relationships. Building upon our running example of Driver #1's statement about solidification verification - "We have to call and ask if the next ladle with a specific number is actually solidified" - we can examine how the ontology represents and automates this knowledge. The hierarchical structure uses classes (like `Ladle`) and subclasses (like `SlagLadle`), connected through object properties (such as `hasMaterial`) that establish relationships between entities. Data properties like `hoursSolidified` capture critical values for safety assessment, such as the 72-hour solidification requirement. These components work together through Description Logic (DL) queries and Semantic Web Rule Language (SWRL) rules. SWRL rules consist of a body (antecedent) and head (consequent), connected by an implication arrow ($\rightarrow$). The body contains class assertions and property relationships combined using logical and ($\wedge$). For example, to automate the driver's manual verification process, a rule checks if an entity is a Ladle, has solidification hours, and if those hours are less than or equal ($\leq$) to 72, classifying such ladles as not solidified. The relationship between improper solidification and explosion risks is captured through similar rules and risk classifications. Table 3 presents a subset of SWRL rules derived from the knowledge elicitation process.

Table 3. Subset of SWRL rules: Defining context-based decisions

Rule and Description
(#1) $Ladle(?x) \wedge hasMaterial(?x,?y) \wedge Drip(?y) \rightarrow DripLadle(?x)$ A ladle with drip material is classified as DripLadle.
(#2) $Ladle(?x) \wedge hasMaterial(?x,?y) \wedge Slag(?y) \rightarrow SlagLadle(?x)$ A ladle with slag material is classified as SlagLadle.
(#3) $Ladle(?x) \wedge hasMaterial(?x,?y) \wedge Lead(?y) \rightarrow LeadLadle(?x)$ A ladle with lead material is classified as LeadLadle.
(#4) $SlagLadle(?x) \wedge hasMaterial(?x,?y) \wedge Liquid(?y) \rightarrow Explosion(?x)$ Liquid in a slag ladle poses an explosion risk.
(#5) $Ladle(?x) \wedge hoursSolidified(?x,?t) \wedge swrlb:lessThanOrEqual(?t,72) \rightarrow SlagIsNotSolidified(?x)$ If less than 72 hours have passed, the slag is not solid.
(#6) $Ladle(?x) \wedge hoursSolidified(?x,?t) \wedge swrlb:greaterThan(?t,72) \rightarrow SlagIsSolidified(?x)$ A ladle is okay for production if solidified over 72 hours.
(#7) $Ladle(?x) \wedge SlagIsNotSolidified(?x) \rightarrow NotOkForProductionSolidification(?x)$ A ladle is NOT okay for production if solidified under 72 hours.
(#8) $Ladle(?x) \wedge hasMaintenanceStatus(?x,true) \rightarrow NotOkForProductionMaintenance(?x)$ A ladle is NOT okay for production if it has some type of maintenance.
(#9) $swrlb:equal(?currentArea,\text{"}southTrain\text{"}) \wedge findNextLadle(?nextLadle) \wedge hasArea(?nextLadle,?nextArea) \wedge OkForProduction(?nextLadle) \rightarrow NextLadleForSouthTrain(?nextLadle,?nextArea)$ Identify the next ladle for the south train.
(#10) $Ladle(?x) \wedge hasMaterial(?x,?m) \wedge swrlb:notEqual(?m,\text{empty}) \wedge hoursSolidified(?x,?h) \wedge swrlb:add(?newH,?h,1) \rightarrow hoursSolidified(?x,?newH)$ Increment solidification time for non-empty ladles.
(#11) $Ladle(?x) \wedge hasArea(?x,?oldA) \wedge hasTargetArea(?x,?newA) \wedge swrlb:notEqual(?oldA,?newA) \rightarrow hasArea(?x,?newA)$ Handle movement of ladles between areas.
(#12) $Ladle(?x) \wedge hasArea(?x,\text{southTrain}) \rightarrow hasMaterial(?x,\text{empty}) \wedge hasPouringCompleted(?x,\text{false}) \wedge hoursSolidified(?x,0)$ Reset ladle properties upon moving to the south train.

5.3. Ontological Reasoning in Practice

The ontology demonstrates its practical value through automated reasoning about various safety-critical scenarios. Consider a situation where an operator or driver discovers liquid residue in a ladle. By marking this in the system, the ontology applies the subclass of `Material`, which in this case is `Liquid` which in turn classes that ladle as the `ExplosionRiskLiquidInSlagLadle` rule (see Table 3), which immediately classifies the ladle as an explosion risk, making it visible to all users through red warning indicators. The system's reasoning extends beyond basic solidification time tracking to encompass multiple risk factors that could compromise safety. Maintenance scenarios similarly benefit from this shared knowledge framework. When either profession marks a ladle for maintenance—whether due to observed cracks, needed sandblasting, or scheduled inspections—the ontology evaluates the `NotOkForProductionMaintenance` rule. This automatically removes the ladle from consideration for production tasks, visually indicated through yellow warning markers. The system maintains this status until maintenance is completed and verified, ensuring that no ladle with pending repairs or inspections enters the production cycle.

The ontology also supports "look ahead" through the application's time simulation capability. Users can fast forward through 1, 5, or 12-hour intervals to visualize future states of the solidification areas. This feature runs all relevant rules against projected times, showing which ladles will become available for production and identifying potential bottlenecks. For example, when advancing time, the system continuously evaluates the

`SolidificationOver72Hours` rule to determine which ladles will be safe for production. This allows operators and drivers to plan their work more effectively, particularly during shift changes or when coordinating maintenance schedules.

5.4. Implementation Architecture

The implementation consists of a web-based front end built with React 18.2 and a Python 3.11 Flask backend, integrated with an OWL 2.0 ontology using the Owlready2 reasoner (version 0.41). The architecture enables real-time queries and updates to the knowledge base through a RESTful API interface. The system processes SWRL rules via the HermiT reasoner integrated within Owlready2, providing continuous evaluation of safety conditions and operational constraints. The interface, shown in Figure 3, displays solidification sites with interactive ladle icons and status indicators. Color-coding reflects operational states: red (not ready for production), yellow (under maintenance), and green (completed 72-hour solidification without risks). The system includes dedicated areas for north and south train operations, where both train systems transporting their respective ladles from the furnace to designated pickup points for the slag truck drivers.

The system also features practical utilities such as a ladle library, search functionality (enabled by translating searches to DL Queries), and time simulation controls that allow users to project ladle states forward. Notably, the SWRL terminal window visible in the screenshot provides transparency into the system's decision-making process. For each operation, it displays the corresponding Description Logic query, SWRL rule being applied, and the resulting determination. This feature allows users to understand how the system arrives at its safety and operational decisions, providing clarity and accountability in the risk assessment process.

Figure 3. Proof-of-concept prototype; Visible ontology response provides transparency.

6. Evaluation

The evaluation process was conducted in two phases to assess 1) the decision-making capabilities of the ontology, and 2) the usability of the dynamic interface. The first phase focused on scenario-based testing, presenting participants with seven realistic operational situations to compare their expert judgment against the system's automated decisions. The second phase comprised an usability assessment of the collaborative digital tool, including hands-on interaction with the interface and structured feedback collection.

Only participants with the roles of driver or operator were included in the evaluation, as they are the primary end-users of the developed technology. This group is best positioned to assess how accurately the system's suggestions align with real-world conditions. By incorporating both drivers and operators, the study captured valuable insights into how information gaps and communication challenges emerge from different operational perspectives.

In Table 4 a summary of the experts' answers and the ontology's responses is presented, with Central Rules referencing Table 3. In this context, correct placement refers to positioning newly poured slag ladles to ensure the next ladle to be emptied has been solidified for the required minimum of 72 hours. This placement decision must account for multiple factors including maintenance status, potential liquid content, solidification time of adjacent ladles, and scenarios with multiple empty areas where the safest placement option may not be immediately apparent.

6.1. Scenario-Based Evaluation Results

Table 4. Scenario-Based Evaluation Overview. Further details available online[3].

#	Description	Query	Participant Responses	Ontology Response & Central Rules
1	Empty area between newly poured and solidified ladle	`FindNextArea AND LadleOKForProduction`	All participants correctly identified optimal area placement	Correct area identification and ladle selection (Rule #9, #11)
2	Multiple empty areas with varying solidification times	`FindNextArea AND SolidificationOver72Hours`	One participant initially suggested unsafe area; others chose correctly	Selected safest area considering solidification times (Rules #6, #7, #11)
3	Risk assessment - ladle contains liquid/slag shell	`ExplosionRiskLiquidInSlagLadle AND LadleOKForProduction`	Noted no way to know liquid content without inspection	Identified and avoided risky ladle (Rules #2, #4)
4	Maintenance status tracking	`hasMaintenanceStatus AND LadleOKForProduction`	Noted difficulty tracking maintenance status	Identified and avoided maintenance ladle (Rule #8)
5	Available ladles in storage	`AreaNr AND LadleOKForProduction`	Relied on phone communication to verify availability	Real-time visibility of available ladles (Rules #9, #11, #12)
6	Ladle emptying time constraints	`FindNextLadleForProduction AND Time` (not allowed between 23:00 and 07:00)	Manual communication about emptied ladles	Automatic tracking of ladle status (Rules #10, #12)
7	Optimal ladle selection with multiple constraints	`FindNextLadleForProduction AND LadleOKForProduction AND SolidificationOver72Hours`	All correctly identified safe alternatives	Selected optimal ladle meeting all constraints (Rules #6, #7, #8)

The evaluation process demonstrated the system's reasoning and decision-making. On all scenarios, the four experts evaluated the system's proposed actions as correct, suggesting that the new system's logic matches experienced operators' decision-making. Scenario 2 was particularly significant. When faced with a situation where there were multiple empty areas on the solidification site, one expert suggested using a ladle placement that could have resulted in an explosion. The system correctly prioritized the safest area, a

decision later supported by the expert. Operator #2 noted, "An accident I was involved in when driving the slag truck similar to this would never have happened with this system."

The evaluation revealed several additional critical findings regarding safety and operational efficiency. The system demonstrated significant value in addressing maintenance tracking deficiencies that pose current safety risks. The limitations of the current whiteboard system, which lacks clear maintenance status indicators, create a dangerous scenario where defective ladles can inadvertently enter production. As one operator acknowledged when discussing real-world implications, a ladle requiring maintenance would be "guaranteed to enter production" under the existing manual system, as stated by Operator #1. The system also standardized placement procedures for newly poured slag ladles, resolving inconsistencies in operators' and drivers' approaches and establishing clear safety-first protocols. Furthermore, the digital system's comprehensive documentation capabilities addressed a crucial challenge in shift transition management, preventing the information loss that frequently occurs during handovers between teams. These improvements collectively enhance both operational safety and process efficiency.

6.2. System Evaluation

The second part involved system demonstration of the interface followed by structured evaluation using 23 questions (14 quantitative, 9 qualitative) on a 5-point Likert scale (1 = Do not agree, 5 = Strongly agree). Figure 4 presents the quantitative scores across the five key categories (Usability, Functionality, Visualization, Efficiency, and Safety), inspired by [41].

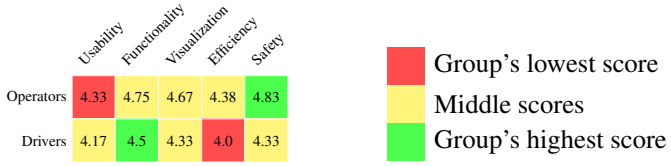

Figure 4. Heatmap of Mean Scores by Role

The system evaluation results indicate that operators rated safety the highest (4.83), while drivers rated functionality the highest (4.5). This aligns with the scenario-based evaluation, where safety-critical decision-making—such as risk assessment for slag ladles—was a key benefit of the system. The lower usability (Operators 4.33, Drivers 4.17) and efficiency (Operators 4.38, Drivers 4.0) scores suggest that while the system provides valuable decision support, further refinements may be needed to streamline interactions and reduce the effort required to access and act on information. The higher safety rating from operators reflects their stronger reliance on risk assessment, while drivers higher rating on functionality suggests they value the system's ability to support operational tasks. These findings highlight a trade-off between detailed safety assessment presentation and ease of use, which should be considered in future development efforts.

7. Discussion

A significant limitation of the model is its reliance on precise sensor readings from the environment. Effectively integrating ontology-based reasoning with real-world data

presents a challenge, particularly in dynamic industrial settings where sensor inconsistencies can impact decision-making. While this challenge can be mitigated through extended knowledge bases and integration with machine learning methods designed to infer and validate facts from the environment, ensuring continuous reliability remains an ongoing concern. Another limitation is that the current model infers risks based on static states but does not capture dynamic transitions between states. Hence, the model cannot reason multiple steps ahead to predict risks. This is a focus for future work.

The evaluation results highlight the model's effectiveness in improving safety, decision-making support, and operational transparency. The system's reasoning aligns closely with expert judgment, successfully identifying risks that were overlooked in traditional methods. Furthermore, transparency is a key strength of the approach—not only in decision-making but also in the knowledge elicitation process. All concepts within the ontology can be traced back to their original expert interviews, ensuring clarity and validation of the knowledge model. Additionally, the ontology-powered intelligent interface (the evaluated prototype) supports visibility into risk factors and maintenance status, minimizing human error. The structured representation of risk factors enhances collaboration among operators and drivers, fostering a shared understanding of critical safety protocols. As a result, this approach offers a scalable foundation for future industrial applications, particularly when integrated with real-time monitoring and adaptive learning techniques.

An ethical assessment of the study is available online[4]; the interview study and expert evaluation were judged appropriate for implementation by authorized decision-maker at Umeå University, Sweden. A key ethical concern in experiments involving automated reasoning is the risk of unpredictable system behavior affecting participants. To address this, all evaluation took place in a synthetic, controlled setting. Expert input was used to construct an ontology, and participants interacted only with predefined, paper-based scenarios rather than the system. The same scenarios were submitted to the system, allowing comparison of human and system responses against expert benchmarks. Further considerations detailed in the ethical assessment.

8. Conclusion and Future Work

This work demonstrates how an ontology-based knowledge model, integrated with an adaptive user interface, enables a digital twin of smelting plant logistics, providing a shared knowledge base for proactive risk assessment. The approach facilitates improved safety, decision-making, and operational transparency by structuring distributed knowledge and critical safety constraints. The evaluation suggests that the system aligns with expert decision-making while addressing existing communication gaps.

In future work, we aim to extend the theoretical framework to incorporate reasoning about dynamic changes in the domain. A promising direction is the development of a transition system that models operational states, enabling forward reasoning for risk prediction and backward reasoning for causal analysis. This would enhance the model's capability to support human-aware planning and improve proactive risk mitigation in industrial settings.

[4]Ethical assessment available at: https://github.com/RA-SPL/Ontology-RA-SPL

References

[1] Rao LL, Xu Y, LI S, Li Y, Zheng R. Effect of perceived risk on nuclear power plant operators' safety behavior and errors. Journal of Risk Research. 2015 05;20:1-9. 10.1080/13669877.2015.1031267.

[2] Tyagi A, Tripathi R, Bouarfa S. Safety management system and hazards in the aircraft maintenance industry: a systematic literature review. Aviation. 2023 11;27:221-3. 10.3846/aviation.2023.19851.

[3] Aven T, Vinnem JE, Wiencke HS. A decision framework for risk management, with application to the offshore oil and gas industry. Reliability Engineering & System Safety. 2007;92(4):433-48. 10.1016/j.ress.2005.12.009.

[4] Komljenovic D, Groves W, Kecojevic V. Injuries in U.S. mining operations – A preliminary risk analysis. Safety Science. 2008 06;46:792-801. 10.1016/j.ssci.2007.01.012.

[5] Smith C, Kastenberg W. On risk assessment of high level radioactive waste disposal. Nuclear Engineering and Design. 1976 12;39:293-333. 10.1016/0029-5493(76)90078-9.

[6] Bobbio A, Portinale L, Minichino M, Ciancamerla E. Improving the analysis of dependable systems by mapping Fault Trees into Bayesian Networks. Reliability Engineering & System Safety. 2001 03;71:249-60. 10.1016/S0951-8320(00)00077-6.

[7] Eschenbach T, Harper W. Modeling Rare Events: Using a Fixed/Variable Rate Model for an Arctic Oil Spill Rate Example. Engineering Management Journal; EMJ. 2013 06;25:46-58. 10.1080/10429247.2013.11431974.

[8] Sangnawakij P, Boehning D, Holling H, Jansen K. Nonparametric estimation of the random effects distribution for the risk or rate ratio in rare events meta-analysis with the arm-based and contrast-based approaches. Statistics in Medicine. 2023 12;43:723-6. 10.1002/sim.9981.

[9] Leão A, Costa E. Fuzzy techniques for assessing workers' risk exposure – A case study in dynamic work environments. Journal of Safety Research. 2019;68:187-96. 10.1016/j.jsr.2018.11.006.

[10] Harvey J, Erdos G, BOLAM H, GREGORY D. An examination of different safety cultures in a nuclear processing plant. Risk, Decision and Policy. 2002 04;7:69-80. 10.1017/S1357530902000510.

[11] Yang X, Zhu C. Industrial Expert Systems Review: A Comprehensive Analysis of Typical Applications. IEEE Access. 2024 01;PP:1-24. 10.1109/ACCESS.2024.3419047.

[12] Kumar S, Agarwal S. Rule based complex event processing for IoT applications: Review, classification and challenges. Expert Systems. 2024 03;41:1-18. 10.1111/exsy.13597.

[13] Varshney KR. Engineering safety in machine learning. In: 2016 Information Theory and Applications Workshop (ITA); 2016. p. 1-8. 10.1109/ITA.2016.7888195.

[14] Agrell C, Eldevik S, Hafver A, Pedersen F, Stensrud E. In: Pitfalls of machine learning for tail events in high risk environments. Taylor & Francis; 2018. p. 3043-51. 10.1201/9781351174664-381.

[15] Lwakatare LE, Raj A, Crnkovic I, Bosch J, Olsson H. Large-Scale Machine Learning Systems in Real-World Industrial Settings A Review of Challenges and Solutions. Information and Software Technology. 2020 07;127:106-10. 10.1016/j.infsof.2020.106368.

[16] Zhang G, Thai VV. Expert elicitation and Bayesian Network modeling for shipping accidents: A literature review. Safety Science. 2016;87:53-62. 10.1016/j.ssci.2016.03.019. Available from: https://www.sciencedirect.com/science/article/pii/S092575351630025X.

[17] Lane P. Meta-analysis of incidence of rare events. Statistical Methods in Medical Research. 2012 01;22:117-32. 10.1177/0962280211432218.

[18] DeBrusk C. The Risk of Machine-Learning Bias (and How to Prevent It); 2018. 10.1007/978-3-319-55961-2_25.

[19] Gruber TR. A translation approach to portable ontology specifications. Knowledge Acquisition. 1993;5(2):199-220. 10.1006/knac.1993.1008.

[20] Horrocks I, Patel-Schneider PF, Boley H, Tabet S, Grosof B, Dean M, et al. SWRL: A semantic web rule language combining OWL and RuleML. W3C Member submission. 2004;21(79):1-31.

[21] Abadi M, Chaimae A, Asmae A, Ben-azza H. A Smart Decision Making System for the Optimization of Manufacturing Systems Maintenance using Digital Twins and Ontologies. International Journal of Advanced Computer Science and Applications. 2022 01;13:77-88. 10.14569/IJACSA.2022.0130811.

[22] Bengler K, Damm W, Lüdtke A, Rieger J, Austel B, Biebl B, et al. A References Architecture for Human Cyber Physical Systems, Part II: Fundamental Design Principles for Human-CPS Interaction. ACM Transactions on Cyber-Physical Systems. 2023 09;8:2-28. 10.1145/3622880.

[23] Gao S, Ren G, Li H. Knowledge Management in Construction Health and Safety Based on Ontology Modeling. Applied Sciences. 2022 08;12:1-18. 10.3390/app12178574.

[24] Nagy L, Ruppert T, Abonyi J. Towards an Ontology-Based Fault Detection and Diagnosis Framework - A Semantic Approach; 2023. p. 1267-72. 10.1109/CoDIT58514.2023.10284094.
[25] Buechel M, Hinz G, Rühl F, Schroth H, Györi C, Knoll A. Ontology-Based Traffic Scene Modeling, Traffic Regulations Dependent Situational Awareness and Decision-Making for Automated Vehicles; 2017. p. 2-7. 10.1109/IVS.2017.7995917.
[26] Pea RD. Practices of Distributed Intelligence and Designs for Education. In: Salomon G, editor. Distributed Cognitions: Psychological and Educational Considerations. New York: Cambridge University Press; 1993. p. 47-83. 10.1017/CBO9780511526335.003. Available from: https://doi.org/10.1017/CBO9780511526335.003.
[27] Cooke N. Varieties of Knowledge Elicitation Techniques. Int J Hum-Comput Stud. 1994 12;41:801-49. 10.1006/ijhc.1994.1083.
[28] Meierhofer J, Schweiger L, Jinzhi L, Züst S, West S, Stoll O, et al. Digital Twin-Enabled Decision Support Services in Industrial Ecosystems. Applied Sciences. 2021 12;11:1-18. 10.3390/app112311418.
[29] Ahmad J, Kostov B, Lališ A. Ontological foundation of hazards and risks in stamp. Semantic Web Journal. 2018:1-14. Available from: https://www.semantic-web-journal.net/system/files/swj2414.pdf.
[30] Kostov B, Ahmad J, Kremen P. Towards Ontology-Based Safety Information Management in the Aviation Industry; 2017. p. 242-51.
[31] Bhosale P, Kastner W, Sauter T. Integrated Safety-Security Risk Assessment for Industrial Control System: An Ontology-based Approach; 2023. p. 1-8. 10.1109/ETFA54631.2023.10275530.
[32] Ding LY, Zhong BT, Wu S, Luo HB. Construction risk knowledge management in BIM using ontology and semantic web technology. Safety Science. 2016;87:202-13. 10.1016/j.ssci.2016.04.008.
[33] Aboutorab H, Hussain O, Saberi M, Hussain F, Chang E. A survey on the suitability of risk identification techniques in the current networked environment. Journal of Network and Computer Applications. 2021 01;178:102-6. 10.1016/j.jnca.2021.102984.
[34] Wu Yj, Xu Kl, Wang Rj, Xu Xh. Safety Integrity Evaluation of High Temperature Molten Metal Operation. Journal of Northeastern University (Natural Science). 2021;42(4):576-7. Available from: https://xuebao.neu.edu.cn/natural/EN/Y2021/V42/I4/576.
[35] Wu Y, Xu K, Wang R, Xu X. Human reliability analysis of high-temperature molten metal operation based on fuzzy CREAM and Bayesian network. PLOS ONE. 2021 08;16:1-16. 10.1371/journal.pone.0254861.
[36] Salama A, Shaheen M, Alfeel H. Rule-based Recommendation System based on Semantic Web of Things. International Journal of Engineering Research and Technology. 2020 06;13:1455-64. 10.37624/IJERT/13.6.2020.1455-1465.
[37] Reda R, Carbonaro A, Boer V, Siebes R, van der Weerdt R, Nouwt B, et al. Supporting Smart Home Scenarios Using OWL and SWRL Rules. Sensors. 2022 05;22:1-18. 10.3390/s22114131.
[38] Liu Z, Nersessian N, Stasko J. Distributed Cognition as a Theoretical Framework for Information Visualization. IEEE transactions on visualization and computer graphics. 2008 11;14:1-9. 10.1109/TVCG.2008.121.
[39] Hollan J, Hutchins E, Kirsh D. Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. vol. 7. Association for Computing Machinery; 2000. 10.1145/353485.353487.
[40] Lindblom J, Thorvald P. Manufacturing in the wild – viewing human-based assembly through the lens of distributed cognition. Production & Manufacturing Research. 2017 06;5:57-80. 10.1080/21693277.2017.1322540.
[41] Kumar AA, Burgard M, Stacey S, Sandaradura I, Lai T, Coorey C, et al. An evaluation of the user-friendliness of Bayesian Forecasting programs in a clinical setting. British Journal of Clinical Pharmacology. 2019:2436-41. 10.1111/bcp.14066.