Preprint CheMIn Technical Reports ISSN 3061-1075

Digitalising Laboratory Chemical Waste Management and Safety

Dragana Usovikj [1], Iman Topchi [1], Jan Leodolter [2], Alk Dransfeld [3], Aleksandar Kondinski [2], Dimitar Trajanov [1,4]

released: 16 July 2026

[1] Faculty of Computer Science and Engineering
Ss. Cyril and Methodius University, Skopje, Macedonia

[2] Institute of Physical and Theoretical Chemistry
Graz University of Technology
Stremayrgasse 9, 8010 Graz, Austria

[3] Institute of Inorganic Chemistry
Graz University of Technology
Stremayrgasse 9/IV, 8010 Graz, Austria

[4] Department of Computer Science,
Metropolitan College
Boston University
Boston, MA 02215, USA

Preprint No. 5



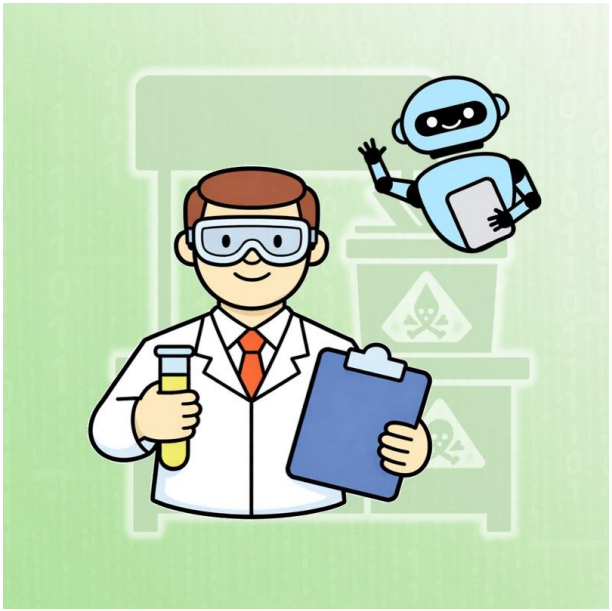
Keywords: digital twin, chemical waste management, laboratory safety, autonomous systems, sensor fusion, knowledge graphs

Edited by
Chemical Modelling and Informatics Group
Institute of Physical and Theoretical Chemistry
Graz University of Technology
Stremayrgasse 9, 8010 Graz
Austria (Europe)

E-Mail: kondinski@tugraz.at
World Wide Web: http://chemin.at/

Abstract

Managing small-volume but versatile chemical waste is a persistent safety and green chemistry challenge. Human error in manual logs often leads to hazardous mixtures, while existing tracking is unreliable. Although safety data is available, waste composition is often not captured at disposal. We propose an autonomous waste management system using a Digital Twin: a Physical Layer to monitor events, a Digital Backbone to fuse data, and a Decision Layer for real-time checks and automation. This system prevents incompatible mixing, creates an audit trail, and optimises resources.



Highlights

- A Digital Twin architecture for autonomous laboratory waste management.
- Real-time sensing and knowledge graphs make disposal traceable.
- Compatibility checks at disposal prevent hazardous mixing.
- Autonomous agents close the loop from observation to action.

Contents

1 Introduction **3**

2 Digital Twin Architecture **4**

3 The Physical Layer **6**

3.1 Computer vision for vessels and phases 6

3.2 Sensors 8

3.3 Humans in the loop 9

4 The Digital Backbone **10**

4.1 Physical-Digital Data Harmonisation 10

4.2 Ontology-based Knowledge Representation 11

4.3 Knowledge Services for Autonomous Agents 13

5 The Decision Layer **14**

5.1 Rule-Based and Model-Driven Intelligence 14

5.2 Real-Time Intervention and Control 16

5.3 Logistical Optimisation and Resource Management 17

6 Technology Gaps and Roadmap **18**

7 Conclusion and Outlook **19**

References **22**

1 Introduction

Chemical waste handling in research laboratories is usually organised as a small set of predefined collection streams rather than a single disposal point. In European academic environments this often includes separate containers for non halogenated organic solvents, halogenated organic solvents, aqueous waste, solid contaminated waste such as wipes, gloves, sorbents or silica, and dedicated special collections for substances that must be excluded from general containers, for example heavy metal containing waste, oxidisers, cyanides, mercury residues, peroxide formers, and pyrophoric or water reactive residues. These categories are well established, yet correct routing remains difficult in academic settings because waste streams are high mix and low volume, and they evolve continuously as experiments change. The practical bottleneck is therefore reliable capture of what is added to which container, when, and in what amount.

This challenge is becoming more consequential as laboratories face a broader shift toward circularity. Under the European Union's Green Deal and the Circular Economy Action Plan, waste is increasingly framed not as an unavoidable byproduct but as a preventable outcome of how processes are designed.[30] For academic laboratories, the implication is not simply stricter disposal, but stronger expectations on traceability and data quality across the chemical lifecycle. Figure 1 summarises the typical segregation workflow and highlights where a Digital Twin can enforce traceable, safety-checked disposal.

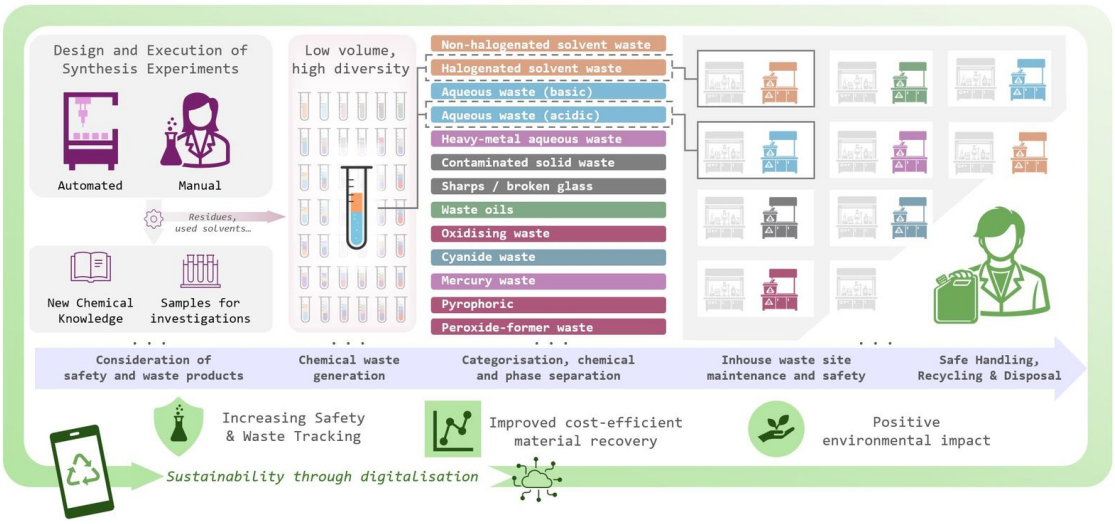


Figure 1: *Schematic of laboratory waste segregation and its digitalisation pathway. Low-volume, high-diversity residues are routed into dedicated waste streams; a Digital Twin can link container identity with transfer confirmation to enable traceable records, safer handling, and improved resource recovery.*

The impetus for change is no longer merely aspirational but is also legally codified. The emerging *Ecodesign for Sustainable Products Regulation*, which entered into force in 2024, and the associated Digital Product Passport (DPP) are shifting practice away from paper-based records and toward interoperable digital documentation.[12, 31] By 2030, the general "Chemicals" category will require machine-readable digital tracking throughout

the lifecycle.[53] Manual waste logs, which often rely on generic classifications such as "halogenated solvents", are fundamentally incompatible with these interoperable digital ecosystems. Furthermore, the "End of Waste" criteria under REACH (Registration, Evaluation, Authorisation and Restriction of Chemicals) impose a strict requirement. That is, laboratories attempting to recycle solvents must prove, through data, that the recovered substance is chemically identical to a registered commercial product.[29] Without a reliable provenance of what enters each container over time, compliance with the EU's "zero pollution" ambition becomes difficult to operationalise.

Beyond compliance, a persistent safety challenge distinguishes academic environments from their industrial counterparts. Data indicate that researchers in academic laboratories are often more likely to be injured than those in the industrial sector.[48] Structural factors such as high staff turnover, decentralised safety cultures, and the long tail of diverse chemical mixtures contribute to this disparity. Managing these unknowns is also a financial haemorrhage. The cost of identifying and disposing of a single unlabelled bottle can be 10 to 50 times that of a known substance.[60] Current manual handling, therefore, drives conservative disposal practices, for example, lab packing and broad classification, because the information required for safe segregation is often incomplete or unavailable at the time of disposal.

While self-driving laboratories have automated synthesis and analysis to accelerate chemical discovery, waste management remains largely unaddressed within the Lab 4.0 paradigm.[44, 65] The loop of autonomy breaks at the waste reservoir, where disposal remains a manual and error-prone process. Waste disposal currently indeed limits true 24/7 operation in self-driving laboratories. When a waste container reaches its capacity, experiments cannot continue unless a person is present to empty or replace it. This makes full unattended operation impractical with existing setups.[55]

This article presents a prospective solution to bridge this gap, an autonomous waste management system based on a Digital Twin concept. We propose a three-layer structure. The Physical reality layer uses sensors and computer vision to monitor pours, vessels, and operating conditions. The Digital Backbone integrates a knowledge graph with electronic laboratory notebooks and related data sources to ensure traceability for regulatory compliance. The Decision Layer performs real-time compatibility checks, applies rules and models, and automates logistics actions. By shifting from after-the-fact logging to verified, real-time recording and decision support at the point of disposal, this system addresses both the regulatory demands of the Green Deal and the practical safety constraints of academic laboratories.

2 Digital Twin Architecture

To organise the development of an autonomous waste management system, we propose a simple and testable three-layer architecture (see Figure 2) consisting of a Physical Reality layer that captures disposal-relevant events through monitoring devices and interactions with chemists/technicians and chemical-waste containers, a Digital Backbone that harmonises these observations and represents them through an ontology and knowledge graph while maintaining audit trails, and a Decision & Action layer that uses this digital

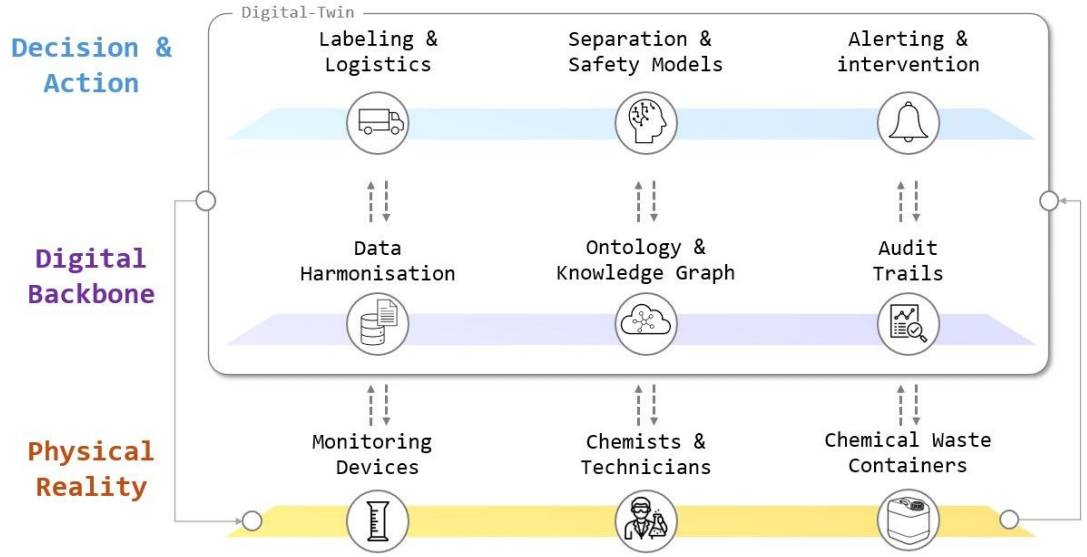


Figure 2: *Three-layer Digital Twin architecture for autonomous waste management: Physical Layer (sensing), Digital Backbone (harmonisation, knowledge graph, audit trail), and Decision & Action Layer (rules, models, interventions, logistics).*

representation to drive separation and safety models, labelling and logistics, and real-time alerting and intervention.

This layered design supports modular development. Each layer can be implemented and validated independently while preserving a clear information flow: observe, understand, log, decide, and act. A thin *digital thread* connects the layers. This ensures that each container and event has a persistent identifier, a timestamp, and a user ID, thereby linking sensor readings, notebook entries, and decision outcomes for compliance and auditing. Because the architecture is modular, a minimal deployment can start with a simple configuration, for example, a camera, a basic container ID with a database, and a small set of hard rules. It can then grow iteratively by adding sensors, richer knowledge representations, and machine-learning-based decision support as data and trust increase.

For example, a user selects an experiment in the electronic laboratory notebook, approaches a designated (smart-)waste container, and pours. The camera and scale detect the pour and forward an event with the container ID to the Digital Backbone. The Decision Layer checks compatibility and capacity, then either permits and logs the pour or issues an immediate intervention, such as a warning, a temporary block, label printing, or pickup scheduling. This illustrates how the digital thread keeps the system auditable and how small, verifiable actions close the loop.

3 The Physical Layer

The Physical Layer is the foundation for reliable, real-time observation in our Digital Twin, tasked with detecting events at the waste station, including vessel identity, phase interfaces, liquid levels, and pour events. Given its critical role, this layer needs to integrate three key elements: (1) the sensing and perception hardware, such as sensors and cameras, (2) the physical infrastructure, such as the canisters, waste containers, and the station itself, and (3) the human users whose actions such as experimentation, pouring, labelling, are being monitored. The combination of sensors and actors enables the production of verifiable, auditable events ("who, when, what, how much") from these interactions with minimal user burden.

3.1 Computer vision for vessels and phases

Computer vision is becoming a practical, cost-effective baseline for observing laboratory vessels, monitoring material transfer, and distinguishing material phases.[26] In laboratory automation, an effective paradigm is to first detect the geometry of the vessel ("container first") and subsequently analyse its contents. The HeinSight system introduced by Jason E. Hein et al. demonstrates that vessel-geometry detection, combined with image-based phase analysis, can provide robust feedback for autonomous chemical operations.[27] Figure 3**a,b** show the HeinSight3.0 hardware setup and the corresponding vision workflow. Sagi Eppel and co-workers introduced the Vector-LabPics dataset, which contains over two thousand images of mostly transparent vessels with annotated filling levels, material phases, labels, and corks, and serves as one such foundation. Models trained on this dataset achieve high accuracy for vessel detection and separating materials such as liquids, solids, and foams, though performance drops in complex multiphase scenes.[28] Even standard webcams with sufficient resolution can provide cost-effective monitoring, aiding decision-making and supporting phase classification in multiphase scenarios (e.g., distinguishing liquid from solid phases).

In real laboratory environments, vision systems must operate under imperfect conditions. Condensation can form on container walls or camera lenses, reducing image clarity, while glare from fume hood sashes and overhead lighting can obscure labels and liquid surfaces. In addition, hands and tools may partially occlude the camera view during pouring, leading to a temporary loss of visual identification. [23, 27, 28]

Furthermore, vision pipelines developed for liquid-liquid extraction can output phase-splitting and interfacial cues and detect pour events in real time. This capability is directly applicable to residue handling.[23] When multiple phases coexist, the system can automatically classify them into distinct categories, such as liquid versus solid or immiscible liquid layers, each corresponding to different waste streams.[23, 28] This distinction is critical for automated separation or routing logic, as a single vessel may contain both recyclable solvent and hazardous residue.[27] High-resolution web cameras provide a practical, low-cost option for continuous monitoring and can be integrated into waste stations or fume hoods without major infrastructure changes.

Vision also provides simple kinetic and optical proxies (e.g., colour change, turbidity)

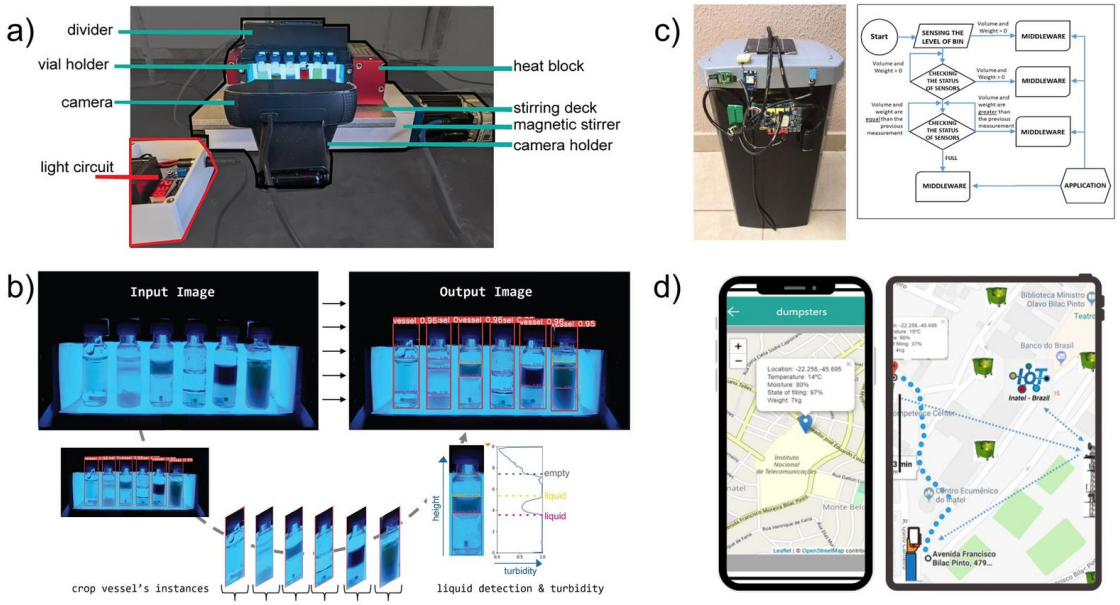


Figure 3: *Overview of vision-enabled monitoring approaches relevant to autonomous waste management.* ***(a,b)*** *HeinSight3.0 high-throughput liquid–liquid separation platform:* ***(a)*** *experimental hardware setup and* ***(b)*** *computer-vision workflow for vessel detection and phase/turbidity analysis (adapted from Elkhawaldeh et al. [27], CC BY).* ***(c,d)*** *Smart waste-bin monitoring system:* ***(c)*** *prototype and data-processing workflow;* ***(d)*** *mobile interfaces for bin status and routing (adapted from Pardini et al. [50], CC BY 4.0).*

that have proven useful for reaction monitoring and can inform routing or recovery decisions for residuals.[70] For example, a sudden colour shift or opacity increase could warn that an incompatible chemical is being poured into the wrong container, prompting an automatic alert or a visual cue to stop the operation. Combined with phase-detection algorithms, these optical and kinetic cues turn basic camera feeds into decision-support mechanisms that reduce human error and generate an auditable record of each pouring and filling event for integration within a Digital Twin framework, as demonstrated in the work of Jonathan D. Hirst et al.[17, 23]

In practice, the robustness of vessel detection improves with controlled camera placement and background, and by using vessel geometry as an anchor for downstream inference. Vision alone has limits (e.g., opaque fluids, look-alike liquids). Therefore, lab systems typically pair cameras with simple orthogonal cues to raise event confidence.[54]

Studies on waste and medical-waste sorting demonstrate that object detection and segmentation models can flag overflowing containers, foreign objects, and hazardous items in noisy, real-world scenes.[18, 38] Edge-ready detection models (e.g., You Only Look Once (YOLO) variants used in waste sorting) show that real-time monitoring can run on modest hardware in distributed lab deployments.[9] Spectral imaging (Short-wave infrared (SWIR) / multispectral) supplements RGB vision where visual appearance is ambiguous, enabling the separation of look-alike liquids and material classes.[45] Taken together, vi-

sion forms the foundational sensing modality in the physical layer, as it identifies vessels, measures bulk phases, and produces pour/fill events that feed the Digital Twin when fused with other sensors and the ELN.

3.2 Sensors

Reliable sensing starts with robust physical infrastructure. This includes labelled, closable canisters, dedicated waste stations, and clear pickup logistics (scheduled collection, contracted carriers).[22, 37] Canisters should support persistent identifiers (barcodes/RFID) and seals for linking sensor events to disposal records.[22, 52] Pressure or differential sensors on closed canisters can provide early warning of gas build-up or pressure drift that would not be visible to a camera, and are therefore a low-cost safety addition to the infrastructure.[72]

A pragmatic sensing stack begins with RGB webcams as the baseline perception channel. This is supplemented by inexpensive orthogonal sensors for event validation and includes chemical-specific probes only where they materially change a routing decision.[26, 54] By "RGB camera/webcam" we mean a colour visible-light camera (red/green/blue channels) at sufficient resolution (typical working minimum of 2–5 MP for bench/mid-range setups). These units operate with controlled lighting and fixed mounts to stabilise the background and viewpoint. They are inexpensive, widely available, and easy to integrate.[28] Webcams are cost-effective because they provide continuous spatial context (vessel geometry, pour motion, splashes) that simple point sensors cannot. Good placement and simple diffuse lighting greatly improve accuracy.[26]

Complementary low-cost sensors can make vision events verifiable. These include load cells (mass), ultrasonic / time-of-flight level probes (height), and simple proximity/beam-break switches that provide independent signals to corroborate a detected pour. A disposal event is only confirmed when both the vision system and the load cell indicate a change consistent with liquid transfer. When image, weight, and level change together, a pour can be auto-logged with high confidence. Conflicts between sensors are expected in real operation. For example, the camera may detect a pouring motion while the load cell shows no measurable change in weight. This can occur due to glare, very small volumes, or liquid falling outside the measurement area. In such cases, the event is treated as uncertain and ignored until confirmation is obtained. In case of disagreement, priority can be given to the load cell for confirming volume change, while vision primarily supports the detection of the pouring action and the identification of the vessel. Unresolved conflicts are flagged for human verification rather than being automatically accepted.[9, 54] For sealed or semi-closed containers, adding a pressure sensor and a simple gas sensor (e.g., NO_x/H_2S for specific lab contexts) catches slowly developing hazards or leaks that cameras and scales would miss.[72]

Smart-bin prototypes integrate cameras, ultrasonic sensors, load cells, and IoT gateways to provide fill status, classification, and alerts.[54] Figure 3c,d show an example prototype and its mobile interface from an IoT waste-management deployment.[50] Typical prototypes pair a Convolutional Neural Network (CNN) (or lightweight detector) for coarse visual classification with ultrasonic/weight signals for fill detection and a Bluetooth/Wi-

Fi link for remote dashboards and pickup scheduling. YOLO/ResNet class models and ultrasonic sensors are common building blocks in these systems.[9] In labs, they flag overfill, detect contamination, and log to ELN/LIMS.[22] Limitations include dataset quality for chemical classification, privacy concerns from video, and the need to tune models for transparent vessels and small volumes typical in labs. These are all reasons to treat smart-bins as a pragmatic complement (not a replacement) to other sensors and human verification.[33]

Sensor combination rules should be simple and auditable. This means they should require the concurrence of two orthogonal signals for automated acceptance (e.g., camera and weight) and fall back to a human verification flow when confidence is low.[54] Optimised, edge-ready models (lightweight CNNs or YOLO variants) enable real-time inference at the waste station without streaming raw video to the cloud, reducing latency and privacy concerns.[9] These models can be highly effective. For instance, the classification accuracy of the proposed CNN-based architecture is 95.3%, and the System Usability Scale (SUS) score is 86%.[54]

3.3 Humans in the loop

Current self-driving laboratories (SDLs) still depend on human specialists to translate abstract research goals into concrete experimental actions, to configure equipment, and to verify safety before runs commence. Rihm et al. emphasise that the vision of fully autonomous laboratories is one where humans would only set high-level goals based on which the AI scientist derives appropriate experiments that are designed, executed, and analysed fully autonomously.[55] However, the current reality necessitates substantial human oversight throughout the experimental process. The authors argue that, while actions within a narrowly defined experimental system are automated, peripheral data and activities are decoupled from the design and execution of experiments and reside in data silos.[55] This creates a critical gap where human intervention becomes essential for bridging disconnected information streams and ensuring experimental integrity.

The paper advocates for a human-in-the-loop approach that recognises the interim technology gap between current capabilities and the envisioned fully autonomous systems.[55] As the authors explain, in comparison with state-of-the-art SDLs, far less human input would be required and, more importantly, less error-prone and potentially arbitrary human input would be involved.[55] This approach acknowledges that humans must continue to play critical roles in breaking down high-level objectives into actionable plans and intervening when sensor data are ambiguous or unexpected.

There are three critical human roles that contribute essential tacit expertise not captured by raw sensor streams alone. Chemists provide indispensable domain expertise and laboratory experience that bridge theoretical knowledge and practical execution, helping define and constrain the experimental search space in current SDL development.[24] Principal investigators (PIs) set experimental goals, risk tolerances, and the overall scientific direction and they provide strategic oversight that aligns experimental campaigns with broader objectives.[55] Lab managers coordinate operational and logistical support, including reagent availability and equipment maintenance, which directly shapes feasible

scheduling and safe day-to-day operation.[55]

4 The Digital Backbone

The Digital Backbone serves as the central information-processing and knowledge-management layer. It bridges the gap between the raw sensor data from the physical layer and the actionable intelligence required by the decision layer. Its primary function is to ingest, contextualise, and structure information to create a verifiable, real-time model of the laboratory waste streams. This is accomplished through data harmonisation, ontological representation, and queryable knowledge services.

4.1 Physical-Digital Data Harmonisation

Physical-digital harmonisation is the pipeline that transforms raw sensor data from the physical layer into verified, auditable waste events. This process involves real-time event-stream processing, in which data from multiple sensors (e.g., cameras and load cells) are fused to identify a single, high-confidence event, such as a pour. This event is then contextualised through external data integration. Information from Electronic Laboratory Notebooks (ELN) or Laboratory Information Management Systems (LIMS) is used to link the event to a specific user and experiment.[8] Data from Building Information Management Systems (BIMS) can provide further spatial and environmental context.[46] The final output of this harmonisation is an automated, time-stamped audit trail that forms the persistent digital thread.

To support interoperability and machine-readable harmonisation, standard laboratory data models, such as the Allotrope Data Framework and the Analytical Information Markup Language (AnIML), can be adopted. These standards provide structured schemas for experimental and analytical data, enabling consistent metadata capture across instruments, ELNs, and LIMS, and reducing ambiguity in how waste-related events are represented and exchanged.[6, 10]

ELNs and LIMS form the backbone of the contextualised event. Rapid access to and analysis of laboratory notebook data enable researchers to obtain insights from results and improve reproducibility across experiments.[36] For instance, some systems employ colour-coding schemes that classify solvents into sustainability tiers, ranging from recommended to highly hazardous, while providing immediate feedback to researchers on the sustainability of their choices.[17] While ELNs capture experimental intent, Laboratory Information Management Systems function as the central data hub for laboratory operations. A LIMS is a computerised system designed to gather, process, and store information generated by laboratories. Originally conceived to automate experimental data management, these systems have evolved into digital hubs that coordinate many aspects of laboratory operations.[64] Together, LIMS and ELNs create a comprehensive digital ecosystem that bridges material management with experimental execution. The openBIS ELN-LIMS exemplifies this integration through its two interconnected components. The LIMS maintains up-to-date information about materials and methods, while the ELN

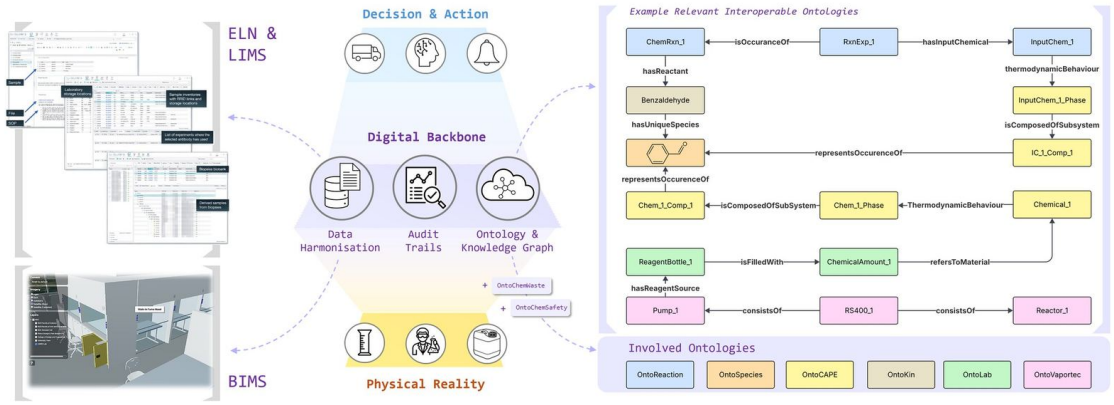


Figure 4: *Digital Backbone concept and example interoperable knowledge representation. The ELN/LIMS example (left) is adapted from Argento [8]. The facilities/BIMS example (bottom left) is adapted from Rihm et al. [56]. The ontology-based knowledge-graph example (right) is redrawn from Bai et al. [11]. All sources are CC BY.*

stores and describes experimental data in a scientifically meaningful manner.[14]

Extending this integrated ecosystem, BIMS further enrich the data harmonisation process by representing the facility's layout and assets. Recent work on BIM-IoT integration demonstrates how sensor data and building models can be fused to enable real-time operational monitoring and decision-making.[20] In our envisioned autonomous waste management system, the building layer (BIMs) indicates the location of each canister, and when integrated with sensor event data (pour, movement) and laboratory data (from ELNs/LIMs) with a time-stamped audit trail, it creates a comprehensive digital thread. Figure 4 summarises the Digital Backbone and illustrates how ELN/LIMS context, facility models (BIMS/BIM), and ontology-based knowledge representation can be integrated into an auditable digital thread.

4.2 Ontology-based Knowledge Representation

This knowledge representation forms the foundation of the Digital Twin, where physical and digital domains are unified within a knowledge graph. An ontology, or a formal specification of concepts and classes, is required to define the core entities of the system. As the paper explains, an ontology is a holistic approach to semantically describing data, information, and knowledge within a domain.[63] It provides the terms, relations, and logic necessary to semantically annotate and link data, which is the foundational work for building a knowledge graph.[63] This is crucial because it ensures that data is not only stored but also FAIR (Findable, Accessible, Interoperable, and Reusable).[63]

Ontologies play a vital role in representing modern chemical knowledge. At the core are upper-level frameworks such as the Basic Formal Ontology (BFO), which supports interoperability across the Open Biological and Biomedical Ontology (OBO) Foundry on-

tologies. .[63] The OBO Foundry provides a curated ecosystem of life science ontologies aligned with shared best practices. Building on these foundations, domain-specific ontologies capture chemical knowledge in detail.[63] Chemical Entities of Biological Interest (ChEBI) classifies chemical entities and their roles and is widely integrated into resources like PubChem and HMDB. The Chemical Information Ontology (CHEMINF) complements this by standardising software terms and data formats for chemical informatics.[63] In other domains, such as waste biorefineries, specialised ontologies formalise flows and processes. For instance, the Waste Biorefinery Ontology enables SPARQL-based discovery of valorisation paths that connect waste resources to conversion technologies and final products.[15]

These ontological frameworks enable practical application. Ontology-driven knowledge graphs provide a unifying semantic layer for various municipal waste streams, enabling automated integration, interpretation, and circular economy decision support.[47] The proposed semantic model captures entities such as producers, authorities, and consumers and links them through Resource Description Framework (RDF) and Web Ontology Language (OWL) representations that can be queried via neural search, reducing the need for manual SPARQL Protocol and RDF Query Language (SPARQL) construction.[47] A similar ontology framework has been applied to hazardous chemical risk management, where a domain-specific schema (including classes for chemicals, incidents, equipment, etc.) supports high precision named entity recognition and enriches the graph with safety relationships.[73]

The most important entities include `WasteContainer`, `ChemicalSubstance`, `User`, `Experiment`, and `PourEvent`, along with their properties and inter-relationships. In this schema, a `Chemical Substance` follows the relational trajectory described by Ghibaudi and Cerruti, that is, chemical substance - material - product - goods - waste, so that each transition can be captured as a state change of the entity.[34] The bibliometric mapping of food-waste research shows that these state changes give rise to distinct research clusters such as anaerobic digestion, composting and other value-added conversion pathways.[49] Accordingly, a `PourEvent` records the moment a material is transferred into a `WasteContainer`, triggering a shift along the `ChemicalSubstance`-waste trajectory and linking the event to the relevant conversion cluster.[49] This ontology provides the formal schema for the knowledge graph, ensuring that data is stored in a consistent, unambiguous, and machine-readable format. The knowledge graph itself is the dynamic instance of this ontology, populated with real-time data to reflect the current state of all waste containers and activities.

The knowledge graph must also capture time-dependent state changes of chemical waste. Certain solvents, such as ethers and dioxane, can form explosive peroxides upon standing for extended periods, making the age of a waste container a critical safety attribute. Therefore, each `WasteContainer` and `ChemicalSubstance` instance should be associated with timestamps and temporal rules that allow the system to infer increasing hazard over time and trigger alerts or disposal actions when safe storage durations are exceeded.[21, 42]

By aligning with established standards, the system is positioned for integration with initiatives such as the Chemistry Consortium within the National Research Data Infrastructure (NFDI4Chem) and with emerging research data infrastructures in Macedonia, such as the

new *Vezilka* ("Embroidery") project led by Ss. Cyril and Methodius University in Skopje (UKIM), ensuring interoperability with broader European open science networks.[19, 62] NFDI4Chem establishes and maintains a national research data infrastructure for the field of chemistry in Germany, facilitating innovative services and scientific endeavours grounded in research data.[62] Through this ontology-based approach, chemical waste management evolves into a FAIR-compliant data system capable of contributing to cross-institutional sustainability analytics and European research initiatives.

4.3 Knowledge Services for Autonomous Agents

Knowledge services provide the interpretive and reasoning layer that enables autonomous agents to query, synthesise, and act on information from heterogeneous laboratory systems. They bridge sensing from the physical layer and actuation from the decision layer by providing machine-readable semantic context that links raw data, rules, and ontological knowledge into actionable form. In digital-lab architectures, this layer is increasingly implemented through knowledge graphs (KGs) and ontologies that structure scientific knowledge for autonomous reasoning and auditability.[17, 27]

Within autonomous laboratories, knowledge services operate as machine-readable intermediaries between the Electronic Laboratory Notebook (ELN), Laboratory Information Management System (LIMS), and the Digital Twin, ensuring that data produced by instruments and experiments are contextualised, interoperable, and queryable. Ontologies and knowledge graphs (KGs) provide the semantic backbone for this layer: they encode chemical entities, equipment, processes, and safety relationships in a standardised, queryable format (e.g., RDF/OWL), allowing agents to retrieve information through SPARQL queries or API endpoints.[27] Recent research highlights dynamic and distributed knowledge graph infrastructures as the foundation for such systems.[11] demonstrates that each laboratory or instrument can be represented as a node in a global, synchronised KG, where autonomous "knowledge agents" continuously update states, metadata, and experiment outcomes, forming a live, self-consistent representation of laboratory operations.

A key component is *Standardised API for Interoperability*, which provides a machine-readable query interface. This allows high-level AI agents to access complex, relational information, such as retrieving the full history and calculated risk profile of a waste mixture. Semantic-web-based interfaces (SPARQL endpoints, JSON-LD/GraphQL wrappers) are common implementations because they support efficient, expressive queries at scale.[11]

The *Decision Support Engine* implements logic to address complex operational questions, including disposal routing and real-time recommendations. For disposal routing, it determines the closest and most appropriate disposal location for a given waste mixture by combining geospatial data from BIMS with hazard profiles from LIMS and applying graph-based pathfinding algorithms. For recommendation support, it provides real-time advisories on whether a proposed chemical addition is safe and indicates the likely resulting state, e.g., single-phase or multiple phases. This service draws on predictive models integrated into the knowledge graph, similar to those used in agent-based digital twins for

fault handling in chemical processes.[35]

Recent multi-agent laboratory frameworks emphasise modular, knowledge-driven coordination. For example,[41] introduces a multi-agent framework where specialised agents query a shared knowledge graph to orchestrate experimental workflows, directly applicable to waste logistics. Similarly,[1] demonstrates how knowledge graphs combined with LLM-powered agents automate scientific discovery, offering a scalable pattern for query-driven decision support.

The evolution of large language models (LLMs) has expanded the role of knowledge services from passive data repositories to active reasoning partners.[16] demonstrates that LLM-driven agents can autonomously plan and execute experiments by combining unstructured text knowledge with structured laboratory interfaces. [58] further shows that integrating LLMs with hazardous-chemical KGs enables retrieval-augmented reasoning for safe experiment and waste-handling decisions. These systems illustrate how LLMs can serve as natural-language front-ends to structured knowledge, translating human or system queries into executable, auditable agent actions. Recent advances show that large language models can be extended to embodied, multimodal agents that directly consume visual and continuous-sensor data while producing sequential action plans for robots (PaLM-E). This demonstrates a practical pathway for coupling semantic knowledge graphs with LLM-based planners so that waste-management agents can both understand sensor context and generate auditable, executable actions.[25]

Finally, the *Actionable Intelligence Outputs* format queries results into clear, machine-actionable instructions (for example, standardised JSON-LD or RDF messages such as "Redirect Canister X to Collection Point Y") so that downstream schedulers, robots, or human operators can execute and audit the decision.[57] This aligns with accountable AI agent designs using knowledge graphs for governance and auditability.[32]

5 The Decision Layer

The Decision Layer is the active, executive component of the Digital Twin. It translates the contextualised knowledge from the digital backbone layer into verifiable, autonomous actions. Its primary purpose is to close the loop, ensuring that all laboratory activities, particularly waste disposal, are safe, compliant, and efficient. It achieves this by querying the knowledge services of the digital backbone and issuing commands to the user-facing interfaces or physical hardware of the physical layer.

5.1 Rule-Based and Model-Driven Intelligence

The Rule-Based and Model-Driven Intelligence implements the core logic to model and predict the dynamic state of chemical waste mixtures, enabling proactive decision-making within the Digital Backbone. Rule systems enforce known chemical compatibility constraints, whereas model-driven methods estimate emergent behaviour (e.g., phase formation, cumulative hazards) when explicit rules are unavailable. This intelligence layer bridges structured data from harmonisation and ontological modelling, applying formal

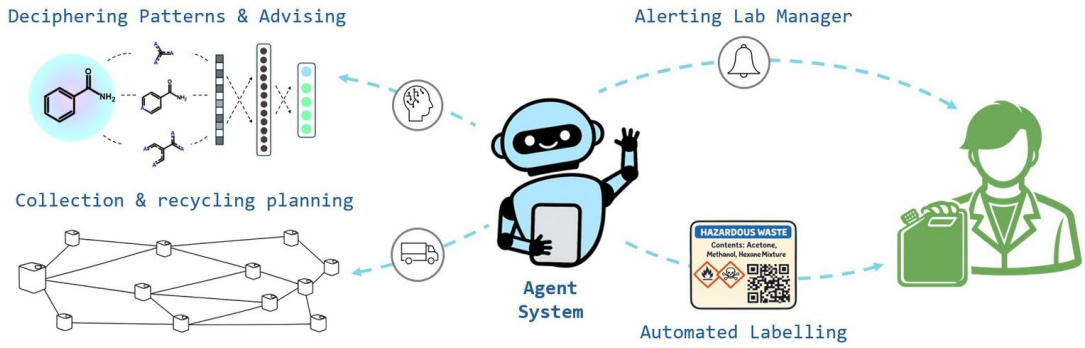


Figure 5: *Illustrative decision-layer schematic for autonomous laboratory waste management. An agent system combines (i) pattern deciphering and advisory functions informed by ML-based monitoring approaches [69], (ii) collection and recycling planning informed by route-optimisation frameworks [5], and (iii) automated labelling and alerting actions within a closed-loop workflow.*

rules and predictive models to ensure safety, compliance, and efficiency. By continuously updating the waste state based on sensor inputs and external knowledge, it prevents hazardous conditions and optimises disposal or recovery paths.

Formal class definitions establish the foundational entities for logical reasoning: Waste Container (e.g. Canister with attributes like capacity and location), Chemical Substance (properties such as density, solubility, hazards from LIMS), Waste Mixture (dynamic state tracking cumulative composition), Waste Streams (disposal rules per regulatory standards) and User (Humans/Agents for accountability).[51] These classes, grounded in semantic ontologies like RDF/OWL, allow extensible modelling of complex interactions, as seen in knowledge engineering for chemistry.[43] Figure 5 sketches the Decision Layer as an agent system that converts verified waste events into guidance (pattern deciphering), logistics decisions (collection/recycling planning), and operational actions (automated labelling and alerting).

Dynamic Chemical State Computation continuously calculates the overall hazard profile (cumulative risk) and predicts phase status (single phase or multi-phase) using integrated physical properties from LIMS data. For cumulative risk, rule-based aggregation applies explicit decision rules to the known constituents and their classifications (e.g., GHS/CLP hazard classes and, where applicable, concentration-based mixture rules), flagging thresholds that trigger alerts.[22] Phase prediction employs model-driven approaches, such as machine learning models trained on solubility and density data to forecast immiscibility or emulsion formation.[23] Recent advancements in ML for mixture phase prediction, like physics-embedded neural networks, enhance accuracy by incorporating thermodynamic principles (e.g. UNIversal Functional Activity Coefficient (UNIFAC) for activity coefficients).[71] This prediction informs disposal recommendations: single-phase mixtures might route to neutralisation, while multi-phase ones suggest separation, reducing waste volume and costs as demonstrated in digital twins for wastewater.[68]

The Compatibility and Rule Engine implements proactive safety checks utilising LIMS hazard data to flag incompatible additions before or immediately after a pour event. Rule-

based systems, such as if-then logic from chemical compatibility charts, prevent reactions like (potentially violent) exothermic reactions and boiling/splashing by cross-referencing SDS data. In lab contexts, intelligent waste systems use CNNs for classification but integrate rules for real-time validation.[54] Model-driven extensions, e.g. graph neural networks for reactivity prediction, forecast outcomes for novel mixtures, aligning with SDLs where agents simulate states to avoid errors.[26] This engine maintains an audit trail, ensuring regulatory compliance by logging all checks and overrides.

Hybrid rule-ML approaches further enhance the layer, combining physical knowledge (e.g. group contributions) with machine learning to predict thermodynamic properties under data-scarce conditions.[39, 40] For instance, matrix completion methods fill gaps in interaction parameters, enabling reliable state computations for diverse waste streams.[40] In digital twins, these hybrids optimise processes like treatment scheduling, as seen in wastewater applications where mechanistic models integrate with ML for phase and property estimation.[68] Overall, this intelligence layer transforms passive data into actionable insights, supporting autonomous agents in the decision layer while addressing gaps like data sparsity through hybrid rule-ML approaches.

5.2 Real-Time Intervention and Control

In the Decision Layer of the waste-management digital twin, software agents continuously monitor live sensor data and apply knowledge-based rules to enact safety interventions on the fly. This Real-Time Intervention and Control component responds to flagged risks from compatibility checks by issuing alerts and initiating automated safeguards, blending human oversight with autonomous actions to maintain laboratory safety. By processing real-time data from the Physical Layer and applying reasoning from the knowledge graph, it prevents incompatible mixtures or sensor anomalies from escalating into incidents such as fires, leaks, or overfills.

Alerts are graduated based on severity: minor issues trigger visual cues (e.g., LED warnings on waste stations), while critical ones send notifications via apps or emails to lab personnel or issue an audible or visual alarm to nearby users.[54] For example, if a pour event risks overfill or incompatibility, the system might display a "stop" signal and log the event for audit trails. In self-driving labs, similar mechanisms use multi-level feedback to halt processes, as seen in platforms that use computer vision to detect anomalies and pause operations.[26]

Automated interventions include physical controls like solenoid locks on canister lids or diversion valves for waste streams, activated when rules detect imminent dangers. To prevent alert fatigue, modern systems gate notifications on a Confidence Score derived from sensor data and historical user reactions.[67] While static thresholds often generate excessive irrelevant alarms, adaptive strategies that personalise limits to situational context significantly reduce alarm burden without sacrificing sensitivity.[66] Machine learning models, specifically random forests, continuously refine these thresholds by learning from prior interactions to suppress false positives. This approach ensures that as a user's experience increases and their reaction patterns stabilise, the system becomes more accurate at presenting only the most relevant alerts.[67]

Automated labelling and compliance management are also part of this layer. Connected IoT printers can generate hazard labels whenever a new chemical deposit or pour event occurs, thereby keeping labels up to date without delay. Each verified waste transfer, confirmed by sensors or RFID, is logged in a secure audit trail to support regulatory review.[22] When a fill level sensor reports that a container is full, a rule can trigger a pickup scheduling workflow and notify the waste team. In one implementation, each virtual object polls its IoT sensors in a loop and allows operations to proceed only when all disposition conditions are satisfied. If a condition is not satisfied, the system immediately notifies the user through the interface.[61] These interventions include alerts, label printing, lid locking, event logging, and pickup scheduling. They illustrate how the Decision Layer uses knowledge graph reasoning and sensor feedback within an event-driven, agent-based architecture. Ontology-based rules can continuously run over the integrated knowledge graph, triggering immediate responses when thresholds are violated. For example, Semantic Web Rule Language (SWRL) logic embedded in the twin can interpret sensor inputs and initiate emergency workflows without human intervention.[61]

Human-in-the-loop protocols allow overrides for verified safe actions, with all interventions recorded immutably for regulatory reviews.[22] This balance supports moderate autonomy, where agents confirm decisions before execution, aligning with SDL frameworks that prioritise safety in uncertain environments.[65] Ultimately, this layered, knowledge-driven design transforms raw sensor data into structured, actionable intelligence. The Decision Layer autonomously interprets live data, infers hazards, and dispatches control actions or alerts within milliseconds, reducing error rates and cutting operational costs through proactive hazard management.

5.3 Logistical Optimisation and Resource Management

Beyond immediate safety, the Decision Layer provides long-term logistical optimisation. By aggregating data from the Digital Backbone over time, this layer can identify patterns in waste generation. It can then propose resource-efficient actions, such as identifying and segregating high-purity solvent streams for simple on-site treatment or reuse. It can also optimise pickup schedules based on predicted fill rates rather than a fixed calendar, reducing both costs and the risks associated with storing full containers. This function directly supports the Green Chemistry goal of proactive and resource-efficient management.

Traditional waste collection systems operate on predetermined schedules that do not account for the actual fill levels of waste containers, resulting in suboptimal resource utilisation and elevated operational costs. In these processes, the frequency of collection routes is established in advance based on fixed plans, regardless of the waste levels in bins. This approach prevents the efficient utilisation of collection trucks while attempting to meet all waste collection demands.[13] Modern IoT-enabled systems address this limitation through the application of predictive analytics. For example, smart bin systems equipped with ultrasonic fill sensors and GPS transmitters can digitally incorporate sensor characteristics to align with real-world smart bin telemetry.[7] Machine learning models, particularly Extreme Gradient Boosting (XGBoost) classifiers, have demonstrated high predictive accuracy for overflow events. XGBoost classification models can predict the likelihood of bin overflow based on historical fill level patterns with accuracies reaching

approximately ninety-four per cent and recall rates exceeding ninety-five per cent.[7]

A critical component of a systems-level approach is the integration of intelligent decision logic to identify and prioritise "High-Value Recovery" streams in real-time.[3] By utilising Machine Learning algorithms to analyse waste stream characteristics (solvent identity, composition, and purity), the system can detect when a container exceeds a viable threshold for reuse.[3] For example, since acetone represents a significant portion of industrial solvent use, the logic can identify high-purity streams (e.g. >98%) and suggest solvent recovery instead of standard disposal via incineration.[59]

The Decision Layer synthesises predictive scheduling, route optimisation, and resource recovery into a unified decision support framework.[4] Machine learning algorithms can be designed to fine-tune the parameters of optimisation algorithms by analysing the inputs and outputs of real-world trips.[2] This integration supports both operational efficiency and sustainability objectives.[7] By shifting laboratory waste management from a reactive, schedule-based approach to a proactive, data-driven system, the Decision Layer enhances safety results, lowers operational expenses, and promotes environmental sustainability in alignment with Green Chemistry principles.

6 Technology Gaps and Roadmap

Although the Physical Layer, Digital Backbone, and Decision Layer each have mature components in isolation, a deployable autonomous waste-management system requires these components to function as a single, verifiable loop under conditions typical of academic laboratories (high-mix, low-volume, frequent edge cases). The main obstacles are therefore not only performance-related, but integration-related: (i) perception must remain robust under realistic visual and material variability, (ii) sensor events must be bound to chemical and experimental context without manual effort, and (iii) safety decisions must generalise beyond fixed rule sets and database lookups to novel compounds and mixtures.

To make the transition from *gap identification* to *implementation logic* explicit, each layer-specific gap is mapped to the capability that must be strengthened and to the enabling AI development that most directly addresses it. This gap-to-roadmap mapping is summarised in Table 1 and serves as the organising bridge from the diagnostic analysis to the proposed development sequence.

The mapping in Table 1 implies a practical roadmap that closes the autonomy loop in three complementary steps. First, the Physical Layer bottleneck is dominated by data sparsity and edge cases (rare hazardous conditions, multiphase scenes, opaque fluids, and variable lighting). Here, Generative AI offers a pragmatic mechanism to expand training coverage without recreating dangerous scenarios: diffusion models can be used to synthesise photorealistic hazard-relevant images (e.g., crystallisation, phase separation, fuming) and difficult operational conditions (occlusions, glare), thereby augmenting datasets such as *Vector-LabPics* and improving robustness of vessel-, phase-, and pour-event inference.

Second, even reliable sensing is insufficient if events remain disconnected from identity and context. The dominant gap in the Digital Backbone is the manual, error-prone binding

Table 1: *Technology gaps and corresponding AI roadmap components for achieving autonomous laboratory waste management.*

Layer	Key Gap	Required Capability	Roadmap Component
Physical	Robust perception under opaque, multiphase, and variable lighting conditions.	Improved edge-case detection and visual generalisation.	Generative AI (synthetic hazard data)
Digital	Manual linking of pour events to ELN, user, and chemical identity.	Automatic semantic eventcontext binding.	Autonomous AI agents
Decision	Inability to predict hazards for novel or complex mixtures beyond rule sets.	Structure-to-property inference and predictive compatibility.	Foundational chemistry models
General	High-mix, low-volume laboratory environment requiring adaptability.	Integrated perception, context, and predictive reasoning.	Combined deployment of all three components

between a detected pour and the corresponding experiment, user, and chemical identity in ELN/LIMS workflows. Autonomous AI agents address this by shifting from passive logging to goal-oriented coordination: an agent can query ELN/LIMS context at the time of the event, reconcile container identity and contents, and trigger operational actions (e.g., label updates, audit logging, pickup requests) with minimal human burden. In multi-station settings, agents can additionally coordinate load balancing by directing users to a nearby compatible container when capacity or compatibility constraints are detected.

Third, the Decision Layer currently remains limited by non-generalising logic: compatibility charts and fixed lookups do not cover the long tail of experimental compounds and complex mixtures. Foundational models of chemistry (e.g., ChemBERTa, MolFormer) provide a route to structure-to-property inference that can extend decision-making beyond explicit database entries. When a new reagent is not in the compatibility knowledge base, the model can generate representations supporting predictive estimates of reactivity- and hazard-relevant behaviour, enabling compatibility scoring and intervention logic even for novel species. In combination, these three roadmap elements operationalise a continuous path from observation (robust perception), to meaning (automatic event–context binding), to action (predictive compatibility and intervention), which is the minimum requirement for a fully autonomous and proactive chemical-waste management system.

7 Conclusion and Outlook

Significant challenges remain in laboratory waste handling, particularly the need for reliable, real-time monitoring. The Digital Twin approach addresses this by integrating multiple data sources. This will require new sensors and smart waste containers. Seamless interoperability with electronic lab notebooks, laboratory information management systems, and building information models shifts safety checks—such as preventing misrouting, overfill, or incompatible mixtures—from manual logs to automated, point-of-disposal

control. The audit trail ensures regulatory compliance and provides traceable records of waste composition and timing.

The next step is real-world deployment and validation. Detection performance must be tested under realistic lab conditions, including glare, occlusions, and complex residues. Event-to-experiment linking should be streamlined while preserving uncertainty flags and an auditable override path. Risk assessment should be enhanced by combining rule-based checks with predictive models that provide calibrated confidence. From a sustainability perspective, the initial focus should be on recovering high-purity solvent streams using accurate composition histories. Standardising event records will facilitate sharing waste data across institutions as regulations develop.

List of abbreviations

AI: Artificial Intelligence
AnIML: Analytical Information Markup Language
API: Application Programming Interface
BFO: Basic Formal Ontology
BIMS: Building Information Management Systems
CEAP: Circular Economy Action Plan
ChEBI: Chemical Entities of Biological Interest
CHEMINF: Chemical Information Ontology
CLP: Classification, Labelling and Packaging
CNN: Convolutional Neural Network
DPP: Digital Product Passport
ELN: Electronic Laboratory Notebook
ESPR: Ecodesign for Sustainable Products Regulation
FAIR: Findable, Accessible, Interoperable, and Reusable
GHS: Globally Harmonised System of Classification and Labelling of Chemicals
IoT: Internet of Things
JSON-LD: JavaScript Object Notation for Linked Data
KG: Knowledge Graph
LIMS: Laboratory Information Management System
LLM: Large Language Model
ML: Machine Learning
NFDI4Chem: Chemistry Consortium in the National Research Data Infrastructure
OBO: Open Biological and Biomedical Ontology
OWL: Web Ontology Language
RDF: Resource Description Framework
REACH: Registration, Evaluation, Authorisation and Restriction of Chemicals
RFID: Radio-Frequency Identification
RGB: Red, Green, Blue
SDL: Self-Driving Laboratory
SDS: Safety Data Sheet
SPARQL: SPARQL Protocol and RDF Query Language
SUS: System Usability Scale
SWIR: Short-Wave Infrared

SWRL: Semantic Web Rule Language
UNIFAC: UNIversal Functional Activity Coefficient
XGBoost: Extreme Gradient Boosting
YOLO: You Only Look Once

Acknowledgements

A.K. thanks the Field of Expertise - Advanced Materials Science of Graz University of Technology for Competitive Start-up Funding support (AF22-635-1). DT Thanks the Ministry of Education and Science of the Republic of North Macedonia through the project "Utilising AI and National Large Language Models to Advance Macedonian Language Capabilities".

Conflicts of interest

No conflict of interest to declare.

Author contributions

Dragana Usovikj and Iman Topchi contributed equally to this work. All authors contributed to the conceptualisation and writing of the manuscript.

Correspondence

Corresponding authors: Aleksandar Kondinski (kondinski@tugraz.at) and Dimitar Trajanov (dimitar.trajanov@finki.ukim.mk).

Data availability

No new data are associated with this article.

References

[1] N. Aamer, M. N. Asim, S. Munir, and A. Dengel. Automating ai discovery for biomedicine through knowledge graphs and llm agents. bioRxiv preprint, 2025. Available at: https://www.biorxiv.org/content/10.1101/2025.05.08.652829v1. Accessed: 2026-03-02.

[2] M. Abdallah, M. Adghim, M. Maraqa, and E. Aldahab. Simulation and optimization of dynamic waste collection routes. *Waste Manag. Res.*, 37(8):793–802, 2019. ISSN 1096-3669. doi:10.1177/0734242X19833152. URL https://doi.org/10.1177/0734242X19833152.

[3] E. A. Aboagye, J. D. Chea, and K. M. Yenkie. Systems level roadmap for solvent recovery and reuse in industries. *iScience*, 24(10):103114, 2021. ISSN 2589-0042. doi:10.1016/j.isci.2021.103114. URL https://doi.org/10.1016/j.isci.2021.103114.

[4] A. M. Alabdali. Blockchain based solid waste classification with ai powered tracking and iot integration. *Sci. Rep.*, 15(1):15197, 2025. ISSN 2045-2322. doi:10.1038/s41598-025-97030-2. URL https://www.nature.com/articles/s41598-025-97030-2.

[5] B. Ali, M. A. Javed, A. A. K. Alharbi, S. Alotaibi, and M. Alkhathami. Internet of things-assisted vehicle route optimization for municipal solid waste collection. *Appl. Sci.*, 14(1):287, 2024. ISSN 2076-3417. doi:10.3390/app14010287. URL https://doi.org/10.3390/app14010287.

[6] Allotrope Foundation. Allotrope data format specification. Technical specification, Allotrope Foundation, 2024. URL https://docs.allotrope.org/. Version 1.5.3.

[7] R. Anitha and A. Parthiban. Ai-iot-graph synergy for smart waste management: a scalable framework for predictive, resilient, and sustainable urban systems. *Front. Sustain.*, 6:1675021, 2025. ISSN 2673-4524. doi:10.3389/frsus.2025.1675021. URL https://www.frontiersin.org/journals/sustainability/articles/10.3389/frsus.2025.1675021/full.

[8] N. Argento. Institutional eln/lims deployment: highly customizable eln/lims platform as a cornerstone of digital transformation for life sciences research institutes. *EMBO Rep.*, 21(3):e49862, 2020. doi:10.15252/embr.201949862. URL https://doi.org/10.15252/embr.201949862.

[9] A. Arishi. Real-time household waste detection and classification for sustainable recycling: A deep learning approach. *Sustainability*, 17(5):1902, 2025. ISSN 2071-1050. doi:10.3390/su17051902. URL https://doi.org/10.3390/su17051902.

[10] ASTM International. Standard specification for analytical data interchange protocol for chromatographic data. Standard ASTM E1947-98(2004), West Conshohocken, PA, 2004. URL https://www.astm.org/.

[11] J. Bai, S. Mosbach, C. J. Taylor, D. Karan, K. F. Lee, S. D. Rihm, J. Akroyd, A. A. Lapkin, and M. Kraft. A dynamic knowledge graph approach to distributed self-driving laboratories. *Nat. Commun.*, 15(1):462, 2024. doi:10.1038/s41467-023-44599-9. URL https://doi.org/10.1038/s41467-023-44599-9.

[12] BanQu. The digital product passport guide: Eu regulatory compliance and beyond. Web page, 2025. Available at: https://www.banqu.co/blog/the-digital-product-passport-guide. Accessed: 2025-11-18.

[13] T. Bányai, P. Tamás, B. Illés, Ž. Stankevičiūtė, and Á. Bányai. Optimization of municipal waste collection routing: Impact of industry 4.0 technologies on environmental awareness and sustainability. *Int. J. Environ. Res. Public Health*, 16(4):634, 2019. ISSN 1660-4601. doi:10.3390/ijerph16040634. URL https://doi.org/10.3390/ijerph16040634.

[14] C. Barillari, D. S. M. Ottoz, J. M. Fuentes-Serna, C. Ramakrishnan, B. Rinn, and F. Rudolf. openbis eln-lims: an open-source database for academic laboratories. *Bioinformatics*, 32(4):638–640, 2016. ISSN 1367-4803. doi:10.1093/bioinformatics/btv606. URL https://doi.org/10.1093/bioinformatics/btv606.

[15] F. Barla, F. Lykokanellos, and A. C. Kokossis. Discovering valorisation paths in waste biorefineries using an ontology engineering approach. In Z. Kravanja and M. Bogataj, editors, *26th European Symposium on Computer Aided Process Engineering*, volume 38 of *Comput. Aided Chem. Eng.*, pages 2079–2084. Elsevier, 2016. doi:10.1016/B978-0-444-63428-3.50351-9. URL https://doi.org/10.1016/B978-0-444-63428-3.50351-9.

[16] D. A. Boiko, R. MacKnight, B. Kline, and G. Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023. doi:10.1038/s41586-023-06792-0. URL https://doi.org/10.1038/s41586-023-06792-0.

[17] S. Boobier, J. C. Davies, I. N. Derbenev, C. M. Handley, and J. D. Hirst. Ai4green: An open-source eln for green and sustainable chemistry. *J. Chem. Inf. Model.*, 63 (10):2895–2901, 2023. doi:10.1021/acs.jcim.3c00306. URL https://doi.org/10.1021/acs.jcim.3c00306.

[18] A. Bruno, C. Caudai, G. R. Leone, M. Martinelli, D. Moroni, and F. Crotti. Medical waste sorting: A computer vision approach for assisted primary sorting. In *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2023 - Workshops, Rhodes Island, Greece, June 4–10, 2023*, pages 1–5. IEEE, 2023. doi:10.1109/ICASSPW59220.2023.10193520. URL https://doi.org/10.1109/ICASSPW59220.2023.10193520.

[19] ChatMED Consortium. Vezilka: A new AI beacon for the region. Web page, 2025. Published: 2025-12-17. Available at: https://chatmed-project.eu/vezilka-a-new-ai-beacon-for-the-region/. Accessed: 2026-03-02.

[20] Y. Chen, X. Wang, Z. Liu, J. Cui, M. Osmani, and P. Demian. Exploring building information modeling (bim) and internet of things (iot) integration for sustainable building. *Buildings*, 13(2):288, 2023. ISSN 2075-5309. doi:10.3390/buildings13020288. URL https://doi.org/10.3390/buildings13020288.

[21] D. E. Clark. Peroxides and peroxide-forming compounds. *Chem. Health Saf.*, 8(5): 12–22, 2001. doi:10.1016/S1074-9098(01)00247-7. URL https://doi.org/10.1016/S1074-9098(01)00247-7.

[22] M. E. Cournoyer, M. M. Maestas, D. R. Porterfield, and P. Spink. Chemical inventory management: The key to controlling hazardous materials. *Chem. Health Saf.*, 12(5):15–20, 2005. doi:10.1016/j.chs.2005.01.018. URL https://doi.org/10.1016/j.chs.2005.01.018.

[23] J. Daglish, A. J. Blacker, G. de Boer, A. Crampton, D. R. J. Hose, A. R. Parsons, and N. Kapur. Determining phase separation dynamics with an automated image processing algorithm. *Org. Process Res. Dev.*, 27(4):627–639, 2023. doi:10.1021/acs.oprd.2c00357. URL https://doi.org/10.1021/acs.oprd.2c00357.

[24] K. Darvish. Self-driving laboratories for chemistry and materials science. ChemRxiv preprint, 2024. Available at: https://doi.org/10.26434/CHEMRXIV-2024-RJ946-V2. Accessed: 2026-03-02.

[25] D. Driess, F. Xia, M. S. M. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, W. Huang, Y. Chebotar, P. Sermanet, D. Duckworth, S. Levine, V. Vanhoucke, K. Hausman, M. Toussaint, K. Greff, A. Zeng, I. Mordatch, and P. Florence. PaLM-e: An embodied multimodal language model. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 8469–8488. PMLR, 2023. URL https://proceedings.mlr.press/v202/driess23a.html.

[26] R. El-khawaldeh, M. Guy, F. Bork, N. Taherimakhsousi, K. N. Jones, J. M. Hawkins, L. Han, R. P. Pritchard, B. A. Cole, S. Monfette, and J. E. Hein. Keeping an "eye" on the experiment: computer vision for real-time monitoring and control. *Chem. Sci.*, 15:1271–1282, 2024. doi:10.1039/D3SC05491H. URL https://doi.org/10.1039/D3SC05491H.

[27] R. El-khawaldeh, A. Mandal, N. Yoshikawa, W. Zhang, R. Corkery, P. Prieto, A. Aspuru-Guzik, K. Darvish, and J. E. Hein. From eyes to cameras: Computer vision for high-throughput liquid-liquid separation. *Device*, 2(7):100404, 2024. ISSN 2666-9986. doi:10.1016/j.device.2024.100404. URL https://doi.org/10.1016/j.device.2024.100404.

[28] S. Eppel, H. Xu, M. Bismuth, and A. Aspuru-Guzik. Computer vision for recognition of materials and vessels in chemistry lab settings and the vector-labpics data set. *ACS Central Sci.*, 6(10):1743–1752, 2020. doi:10.1021/acscentsci.0c00460. URL https://doi.org/10.1021/acscentsci.0c00460.

[29] European Chemicals Agency (ECHA). Guidance on waste and recovered substances. Guidance document (REACH), 2010. Available at: https://echa.europa.eu/documents/10162/23036412/waste_recovered_en.pdf. Accessed: 2026-02-23.

[30] European Commission. Circular economy action plan: For a cleaner and more competitive europe. Web page, 2020. Available at: https://environment.ec.europa.eu/strategy/circular-economy_en. Accessed: 2025-11-18.

[31] European Parliament and Council. Regulation (eu) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products. Official Journal of the European Union, 2024. Available at: https://data.europa.eu/eli/reg/2024/1781/oj. Accessed: 2025-11-18.

[32] R. Fieblinger, M. T. Alam, and N. Rastogi. Actionable cyber threat intelligence using knowledge graphs and large language models. In *2024 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 100–111, 2024. doi:10.1109/EuroSPW61312.2024.00018. URL https://doi.org/10.1109/EuroSPW61312.2024.00018.

[33] F. Fotovvatikhah, I. Ahmedy, R. M. Noor, and M. U. Munir. A systematic review of ai-based techniques for automated waste classification. *Sensors*, 25(10):3181, 2025. ISSN 1424-8220. doi:10.3390/s25103181. URL https://doi.org/10.3390/s25103181.

[34] E. Ghibaudi and L. Cerruti. Chemical substance, material, product, goods, waste: a changing ontology. *Found. Chem.*, 19(2):97–123, 2017. ISSN 1572-8463. doi:10.1007/s10698-017-9281-8. URL https://doi.org/10.1007/s10698-017-9281-8.

[35] M. S. Gill, J. Vyas, A. Markaj, F. Gehlhoff, and M. Mercangöz. Leveraging llm agents and digital twins for fault handling in process plants. arXiv preprint, 2025. Available at: https://arxiv.org/abs/2505.02076. Accessed: 2026-03-02.

[36] S. G. Higgins, A. A. Nogiwa-Valdez, and M. M. Stevens. Considerations for implementing electronic laboratory notebooks in an academic research environment. *Nat. Protoc.*, 17(2):179–189, 2022. doi:10.1038/s41596-021-00645-8. URL https://doi.org/10.1038/s41596-021-00645-8.

[37] X. Jiang and X. Li. Research on hazardous chemicals information management system. In *2011 IEEE 3rd International Conference on Communication Software and Networks (ICCSN)*, pages 630–633, 2011. doi:10.1109/ICCSN.2011.6014649. URL https://ieeexplore.ieee.org/document/6014649.

[38] S. Jin, Z. Yang, G. Królczyk, X. Liu, P. Gardoni, and Z. Li. Garbage detection and classification using a new deep learning-based machine vision system as a tool for sustainable waste recycling. *Waste Manag.*, 162:123–130, 2023. ISSN 0956-053X. doi:10.1016/j.wasman.2023.02.014. URL https://doi.org/10.1016/j.wasman.2023.02.014.

[39] F. Jirasek and H. Hasse. Combining machine learning with physical knowledge in thermodynamic modeling of fluid mixtures. *Annu. Rev. Chem. Biomol. Eng.*, 14

(1):31–51, 2023. doi:10.1146/annurev-chembioeng-092220-025342. URL https://doi.org/10.1146/annurev-chembioeng-092220-025342.

[40] F. Jirasek, R. Bamler, S. Fellenz, M. Bortz, M. Kloft, S. Mandt, and H. Hasse. Making thermodynamic models of mixtures predictive by machine learning: matrix completion of pair interactions. *Chem. Sci.*, 13(17):4854–4862, 2022. doi:10.1039/D1SC07210B. URL https://doi.org/10.1039/D1SC07210B.

[41] N. S. Johnson. An autonomous, large language model driven x-ray microscope with contextual understanding of domain specific experimental workflows. *Microsc. Microanal.*, 31(Suppl_1):ozaf048.1092, 2025. doi:10.1093/mam/ozaf048.1092. URL https://doi.org/10.1093/mam/ozaf048.1092.

[42] R. J. Kelly. Review of safety guidelines for peroxidizable organic chemicals. *Chem. Health Saf.*, 3(5):28–36, 1996. doi:10.1021/acs.chas.8b03515. URL https://pubs.acs.org/doi/abs/10.1021/acs.chas.8b03515.

[43] A. Kondinski, J. Bai, S. Mosbach, J. Akroyd, and M. Kraft. Knowledge engineering in chemistry: From expert systems to agents of creation. *Acc. Chem. Res.*, 56(2):128–139, 2023. doi:10.1021/acs.accounts.2c00617. URL https://doi.org/10.1021/acs.accounts.2c00617.

[44] J. Leodolter, S. D. Rihm, M. Kraft, and A. Kondinski. Digitale chemie: Ein neues zeitalter. *Nachr. Chem.*, 73(11):40–45, Nov. 2025. doi:10.1002/nadc.20254154356.

[45] A. Manakkakudy, A. D. Iacovo, E. Maiorana, F. Mitri, and L. Colace. Waste material classification: A short-wave infrared discrete-light-source approach based on light-emitting diodes. *Sensors*, 24(3):809, 2024. ISSN 1424-8220. doi:10.3390/s24030809. URL https://doi.org/10.3390/s24030809.

[46] J. Merino, X. Xie, N. Moretti, J. Y. Chang, and A. Parlikad. Data integration for digital twins in the built environment based on federated data models. *Proc. Inst. Civ. Eng.: Smart Infrastruct. Constr.*, 176(4):194–211, 2023. doi:10.1680/jsmic.23.00002. URL https://doi.org/10.1680/jsmic.23.00002.

[47] M. Motevallian, A. M. Esfar-E-Alam, A. Taherkordi, and G. Abbasi. Semantic modeling of waste dataflow for automating circular economy systems. In *2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*, pages 677–684, 2024. doi:10.1109/DCOSS-IoT61029.2024.00105. URL https://doi.org/10.1109/DCOSS-IoT61029.2024.00105.

[48] I. M. Nasrallah, A. K. E. Kak, L. A. Ismaiil, R. R. Nasr, and W. T. Bawab. Prevalence of accident occurrence among scientific laboratory workers of the public university in lebanon and the impact of safety measures. *Saf. Health Work*, 13(2):155–162, 2022. doi:10.1016/j.shaw.2022.02.001. URL https://doi.org/10.1016/j.shaw.2022.02.001.

[49] Y. Ouyang, Y. Cai, and H. Guo. Visualization and analysis of mapping knowledge domains for food waste studies. *Int. J. Environ. Res. Public Health*, 18(10):5143, 2021. ISSN 1660-4601. doi:10.3390/ijerph18105143. URL https://doi.org/10.3390/ijerph18105143.

[50] K. Pardini, J. J. Rodrigues, O. Diallo, A. K. Das, V. H. C. de Albuquerque, and S. A. Kozlov. A smart waste management solution geared towards citizens. *Sensors*, 20 (8):2380, 2020. doi:10.3390/s20082380. URL https://doi.org/10.3390/s20082380.

[51] L. Pascazio, S. D. Rihm, A. Naseri, S. Mosbach, J. Akroyd, and M. Kraft. Chemical species ontology for data integration and knowledge discovery. *J. Chem. Inf. Model.*, 63(21):6569–6586, 2023. doi:10.1021/acs.jcim.3c00820. URL https://doi.org/10.1021/acs.jcim.3c00820.

[52] P. J. Prasad and G. L. Bodhe. Trends in laboratory information management system. *Chemom. Intell. Lab. Syst.*, 118:187–192, 2012. ISSN 0169-7439. doi:10.1016/j.chemolab.2012.07.001. URL https://doi.org/10.1016/j.chemolab.2012.07.001.

[53] Protokol. Digital product passport (dpp): The complete guide. Web page, 2025. Available at: https://www.protokol.com/insights/digital-product-passport-complete-guide/. Accessed: 2025-11-18.

[54] M. W. Rahman, R. Islam, A. Hasan, N. I. Bithi, M. M. Hasan, and M. M. Rahman. Intelligent waste management system using deep learning with iot. *J. King Saud Univ. Comput. Inf. Sci.*, 34(5):2072–2087, 2022. ISSN 1319-1578. doi:10.1016/j.jksuci.2020.08.016. URL https://doi.org/10.1016/j.jksuci.2020.08.016.

[55] S. D. Rihm, J. Bai, A. Kondinski, S. Mosbach, J. Akroyd, and M. Kraft. Transforming research laboratories with connected digital twins. *Nexus*, 1(1):100004, 2024. ISSN 2950-1601. doi:10.1016/j.ynexs.2024.100004. URL https://doi.org/10.1016/j.ynexs.2024.100004.

[56] S. D. Rihm, Y. R. Tan, W. Ang, H. Y. Quek, X. Deng, M. T. Laksana, J. Bai, S. Mosbach, J. Akroyd, and M. Kraft. The digital lab facility manager: Automating operations of research laboratories through "The World Avatar". *Nexus*, 1(3):100031, 2024. ISSN 2950-1601. doi:10.1016/j.ynexs.2024.100031. URL https://doi.org/10.1016/j.ynexs.2024.100031.

[57] S. D. Rihm, A. Kondinski, and M. Kraft. Product design, synthesis, and lab automation with The World Avatar. *Curr. Opin. Chem. Eng.*, 51:101203, 2026. doi:10.1016/j.coche.2025.101203.

[58] P. Scheibe and J. Schor. Ai-driven science communication: Leveraging llms and knowledge graphs for seamless knowledge exchange (ecotoxfred). bioRxiv preprint, 2025. Available at: https://www.biorxiv.org/content/10.1101/2025.07.04.663152v1. Accessed: 2026-03-02.

[59] C. Seyler, C. Capello, S. Hellweg, C. Bruder, D. Bayne, A. Huwiler, and K. Hungerbühler. Waste-solvent management as an element of green chemistry: A comprehensive study on the swiss chemical industry. *Ind. Eng. Chem. Res.*, 45(22):7700–7709, 2006. doi:10.1021/ie060525l. URL https://doi.org/10.1021/ie060525l.

[60] Sigma-Aldrich. The true cost of inefficient inventory management. Report (web), 2024. Available at: https://www.sigmaaldrich.com/deepweb/assets/sigmaaldrich/marketing/global/documents/436/060/chemical-waste-report-mk.pdf. Accessed: 2025-11-18.

[61] G. Stavropoulou, K. Tsitseklis, L. Mavraidi, K.-I. Chang, A. Zafeiropoulos, V. Karyotis, and S. Papavassiliou. Digital twin meets knowledge graph for intelligent manufacturing processes. *Sensors*, 24(8):2618, 2024. ISSN 1424-8220. doi:10.3390/s24082618. URL https://doi.org/10.3390/s24082618.

[62] C. Steinbeck, O. Koepler, F. Bach, S. Herres-Pawlis, N. Jung, M. Razum, C. Baldauf, F. Biedermann, T. Bocklitz, F. Kriegstein, B. Krüger, L. Laplanche, F. Müller, S. Neumann, L. Pahnke, A. Paulisch, A. Porzel, U. Schröder, A. Streit, A. Supady, R. Tuderman, M. Westermann, M. Wilkinson, and B. Zdrazil. Nfdi4chem – towards a national research data infrastructure for chemistry in germany. *Res. Ideas Outcomes*, 6:e55852, 2020. doi:10.3897/rio.6.e55852. URL https://doi.org/10.3897/rio.6.e55852.

[63] P. Strömert, J. Hunold, A. Castro, S. Neumann, and O. Koepler. Ontologies4chem: the landscape of ontologies in chemistry. *Pure Appl. Chem.*, 94(6):605–622, 2022. doi:10.1515/pac-2021-2007. URL https://doi.org/10.1515/pac-2021-2007.

[64] D. Sun, L. Wu, and G. Fan. Laboratory information management system for biosafety laboratory: Safety and efficiency. *J. Biosaf. Biosecur.*, 3(1):28–34, 2021. ISSN 2588-9338. doi:10.1016/j.jobb.2021.03.001. URL https://doi.org/10.1016/j.jobb.2021.03.001.

[65] G. Tom, S. P. Schmid, S. G. Baird, Y. Cao, K. Darvish, H. Hao, S. Lo, S. Pablo-García, E. M. Rajaonson, M. Skreta, et al. Self-driving laboratories for chemistry and materials science. *Chem. Rev.*, 124(16):9633–9732, 2024. doi:10.1021/acs.chemrev.4c00055. URL https://doi.org/10.1021/acs.chemrev.4c00055.

[66] M. C. van Rossum, L. B. Vlaskamp, L. M. Posthuma, M. J. Visscher, M. J. Breteler, H. J. Hermens, C. J. Kalkman, and B. Preckel. Adaptive threshold-based alarm strategies for continuous vital signs monitoring. *J. Clin. Monit. Comput.*, 36(2):407–417, 2022. doi:10.1007/s10877-021-00666-4. URL https://doi.org/10.1007/s10877-021-00666-4.

[67] F. Voutsas, J. Violos, and A. Leivadeas. Mitigating alert fatigue in cloud monitoring systems: A machine learning perspective. *Comput. Netw.*, 250:110543, 2024. ISSN 1389-1286. doi:10.1016/j.comnet.2024.110543. URL https://doi.org/10.1016/j.comnet.2024.110543.

[68] A.-J. Wang, H. Li, Z. He, Y. Tao, H. Wang, M. Yang, D. Savic, G. T. Daigger, and N. Ren. Digital twins for wastewater treatment: A technical review. *Engineering*, 36(5):23–39, 2024. ISSN 2095-8099. doi:10.1016/j.eng.2024.04.012. URL https://doi.org/10.1016/j.eng.2024.04.012.

[69] Y. Wu and G. Wang. Machine learning based toxicity prediction: From chemical structural description to transcriptome analysis. *Int. J. Mol. Sci.*, 19(8):2358, 2018. ISSN 1422-0067. doi:10.3390/ijms19082358. URL https://doi.org/10.3390/ijms19082358.

[70] C. Yan, C. Fyfe, L. Minty, H. Barrington, C. Jamieson, and M. Reid. Computer vision as a new paradigm for monitoring of solution and solid phase peptide synthesis. *Chem. Sci.*, 14(42):11872–11880, 2023. doi:10.1039/D3SC01383A. URL https://pubs.rsc.org/en/content/articlelanding/2023/sc/d3sc01383a.

[71] Y. Yang and S.-T. Lin. Physics-embedded machine learning model for phase equilibrium prediction in multicomponent systems. *J. Chem. Inf. Model.*, 65(19):10180–10193, 2025. doi:10.1021/acs.jcim.5c01804. URL https://doi.org/10.1021/acs.jcim.5c01804.

[72] Y. Zhang, C. Sun, W. Shan, J. Cai, L. Jing, and W. Shao. Systems approach for the safety and security of hazardous chemicals. *Marit. Policy Manag.*, 47(4):500–522, 2020. doi:10.1080/03088839.2019.1710612. URL https://doi.org/10.1080/03088839.2019.1710612.

[73] X. Zheng, B. Wang, Y. Zhao, S. Mao, and Y. Tang. A knowledge graph method for hazardous chemical management: Ontology design and entity identification. *Neurocomputing*, 430:104–111, 2021. ISSN 0925-2312. doi:10.1016/j.neucom.2020.10.095. URL https://doi.org/10.1016/j.neucom.2020.10.095.