Semantic Multi-Modal Framework for Early Stage Anomaly Detection in Lithium-Ion Batteries

Marwa Zitouni [a,*], Franco Giustozzi [a] and Tedjani Mesbahi [a]
[a] *ICube laboratory UMR 7357 and CNRS, Université de Strasbourg, INSA Strasbourg, France*
E-mail: marwa.zitouni@insa-strasbourg.fr

Abstract. In the era of intelligent energy systems, lithium-ion batteries play a pivotal role in electric mobility and large scale energy storage. Ensuring their safe and reliable operation requires advanced anomaly detection methods capable of interpreting complex and heterogeneous data streams. Conventional monitoring approaches, however, struggle to integrate multimodal information and contextual knowledge in real time, limiting their effectiveness for early detection of safety critical events.

To address this challenge, we present an ontology-based framework for anomaly detection in lithium ion batteries. The framework combines heterogeneous data sources, including time series sensor measurements and thermal imaging, within a modular ontology that encodes battery processes, safety thresholds, and causal dependencies. Stream reasoning is performed through C-SPARQL queries, enabling continuous analysis of evolving data and the identification of anomalies when semantic constraints are violated or precursors of hazardous conditions are observed.

The proposed approach is validated on case studies of overheating and thermal imbalance, demonstrating improved detection and providing interpretable explanations beyond conventional threshold monitoring. These results highlight the potential of combining symbolic knowledge representation with real time data analytics to support transparent and scalable battery management in safety critical applications.

Keywords: Lithium-ion batteries, anomaly detection, ontology-based reasoning, stream reasoning, C-SPARQL, multimodal data, explainable AI, predictive maintenance

1. Introduction

In the current landscape of energy and automotive industries, the push towards sustainability, electrification, and higher operational efficiency has accelerated the adoption of advanced technologies. The demand for reliable, high performance energy storage solutions has led to the widespread deployment of lithium ion batteries *(LIBs)*, which are now central to applications ranging from medical devices and consumer electronics to large scale battery energy storage systems *(BESS)* that stabilize power from intermittent renewable sources such as solar and wind. In particular, *LIBs* are driving the rapid growth of the electric vehicle (EV) [11] market, serving as the backbone of modern transportation systems that aim to reduce carbon emissions while maintaining high performance and reliability.

To address the emerging challenges of modern applications, battery systems are evolving toward human centric, sustainable, and resilient infrastructures, reflecting the principles of *Industry 5.0* [28]. Advanced *Battery Manage-*

*Corresponding author. E-mail: marwa.zitouni@insa-strasbourg.fr.

ment Systems (BMS), intelligent sensing, and AI assisted analytics enable collaboration between humans and machines for adaptive and responsible control of battery operation. By integrating *IoT*, cyber physical systems, and edge computing within a sustainable framework, *LIBs* are transforming from passive energy storage units into active, cooperative elements of human aware industrial ecosystems. These technologies support continuous monitoring, informed human decision making, and environmentally conscious optimization of battery performance ensuring safety, resilience, and sustainability in accordance with the core vision of *Industry 5.0*.

Despite their pivotal role, *LIBs* face critical challenges regarding operational reliability and safety [6]. The complex electrochemical and thermal processes within battery cells are highly sensitive to usage conditions, manufacturing defects, and aging, making *LIBs* susceptible to anomalies that can compromise performance, lifespan, and safety. Among the most severe risks is thermal runaway, a self amplifying heat generation process that can lead to fires or explosions, posing a significant threat in high density applications like *EVs*. Beyond catastrophic failures, undetected anomalies result in gradual degradation of capacity and power, reducing vehicle range and user confidence. The economic impact is substantial as well, since battery packs constitute the most expensive *EV* component, and premature failures trigger costly replacements and operational downtime.

A key barrier to addressing these challenges lies in exploiting the large volumes of heterogeneous data generated during battery operation. Modern *BMS* continuously capture voltage, current, temperature, and other critical parameters, producing rich time series data that remains largely underutilized. Within the framework of *Industry 5.0*, this data can be transformed into structured knowledge that enables predictive maintenance proactively detecting and preventing failures before they occur. Predictive maintenance relies on integrating sensor data with advanced analytics, providing early warnings of developing faults and optimizing maintenance schedules to reduce costs, extend battery life, and enhance safety.

This work is motivated by the need to move beyond conventional approaches that treat battery data and anomaly detection in isolation. Despite advances in monitoring and predictive maintenance, current methods often fail to integrate heterogeneous data sources, contextual knowledge, and temporal reasoning into a unified framework. As a result, subtle early stage faults remain undetected, limiting the ability to anticipate failures and optimize maintenance proactively. This paper addresses this gap by proposing a hybrid approach that combines time series data from internal sensors, thermal imaging, and ontology based background knowledge. This integration provides a more comprehensive, interpretable, and intelligent foundation for battery management in line with *Industry 5.0* principles.

The remainder of the paper is organized as follows. First, we review existing techniques for anomaly detection and predictive maintenance in lithium ion batteries, highlighting their limitations and opportunities. We then introduce the theoretical backgrounds and formal definitions of our hybrid approach, followed by a detailed presentation of the framework itself. Next, we demonstrate its application through experimental evaluations. Finally, we conclude with insights into the implications of this approach and outline future directions, including its extension to more complex and data rich systems in line with *Industry 5.0* principles.

2. Related work

The safe and reliable operation of lithium ion batteries is critical, motivating extensive research into advanced anomaly detection methods. This section reviews and synthesizes the current state of research along two primary dimensions. First, we examine anomaly detection methods. Second, we analyze detection approaches based on data modalities, distinguishing techniques that rely on time series sensor data, thermal or other imaging data, or a combination of both, emphasizing how different data sources influence detection performance. Throughout, we also discuss the role of ontology and knowledge integration in enhancing the interpretability of anomaly detection techniques.

2.1. Anomaly detection approaches in Li-ion Batteries

Anomaly detection in *LIBs* can be broadly categorized based on the underlying methodology, ranging from physics driven models to purely data driven techniques, knowledge based reasoning, and combinations thereof. Each category offers distinct advantages and challenges, which are critical to consider when developing effective monitoring and diagnostic framework.

2.1.1. Model based approaches:

relies on physical representations of battery processes, where anomalies are inferred as deviations from predicted behavior. *Equivalent Circuit Models (ECMs)* as discussed in [33, 42] approximate the electrochemical dynamics of *LIBs* with lumped electrical components, enabling real time estimation of state of charge *(SOC)*, state of health *(SOH)*, and internal resistance. In [24, 47] *Electrochemical Models (EChMs)*, on the other hand, provide a more mechanistic representation by simulating ion transport, charge transfer, and thermal dynamics . These models are particularly powerful in diagnosing degradation pathways such as lithium plating, electrolyte decomposition, or electrode passivation.

The main strengths of model based approaches are their interpretability and grounding in physical laws, which make them attractive for safety critical applications. However, their limitations are equally evident: *ECMs* may oversimplify internal processes, while *EChMs* are computationally expensive and require precise parameterization that varies with chemistry, aging, and environmental conditions. Furthermore, real world variability in operating profiles often renders purely model based anomaly detection impractical for large scale deployment in electric vehicles or stationary storage.

2.1.2. Data driven approaches:

Data driven methods treat anomaly detection as a problem of pattern recognition in sensor data. With the widespread adoption of machine learning *(ML)* and deep learning *(DL)*, these approaches have gained dominance in recent years . Commonly employed algorithms include *Support Vector Machines(SVMs)*, *k-Nearest Neighbors(k-NN)*, random forests, and clustering methods such as k-means or *DBSCAN*. Deep architectures such as long short term memory *(LSTM)* networks, autoencoders, and convolutional neural networks *(CNNs)* have been used to capture complex temporal and spatial features in voltage, current, and temperature time series.

These methods are praised for their ability to handle large datasets, uncover nonlinear relationships, and deliver high predictive performance. However, they are inherently black box approaches, often lacking the interpretability necessary for trust and deployment in safety critical contexts. Moreover, data driven methods require extensive labeled datasets, which are scarce in the anomaly detection domain, since failures are rare and costly to reproduce. Their performance may degrade when applied across chemistries, cell formats, or usage conditions different from the training data.

2.1.3. Knowledge based approaches:

Knowledge based approaches seek to overcome the limitations of black box *ML* by embedding domain knowledge into the detection process. Expert systems and rule based reasoning frameworks were among the earliest efforts, where anomalies were defined by thresholds or heuristics derived from empirical knowledge . More recently, ontologies have emerged as a powerful tool for formalizing *LIB* knowledge. Ontologies enable explicit representation of entities (cells, electrodes, sensors), relations (causality, temporal precedence), and constraints (safety thresholds, degradation rules). This allows for semantic reasoning, where anomalies can be not only detected but also explained in context for example, identifying that an observed temperature rise is temporally aligned with current surges and exceeds safety limits.

Knowledge based approaches thus offer interpretability and integration of heterogeneous data sources. However, their coverage is limited by the scope of encoded knowledge, and most existing ontologies in the battery domain focus on chemistry and degradation modeling rather than anomaly detection. Importantly, their application to hybrid sensor image contexts remains underexplored.

2.1.4. Hybrid approaches:

Recognizing the strengths and weaknesses of individual methods, recent works have attempted to integrate multiple paradigms into hybrid approaches. For example, combining model based methods with machine learning allows bridging physical interpretability alongside data driven adaptability. Similarly, the integration of ontologies with *ML* has been explored to add a semantic reasoning layer for improved transparency. Beyond methodological hybrids, some studies also explore sensor fusion, integrating multiple data modalities such as electrical signals and thermal images.

Despite their promise, hybrid approaches are often fragmented: many prioritize predictive accuracy while neglecting semantic reasoning, or they integrate multiple data modalities without an explicit framework for knowl-

edge driven interpretability. This leaves a gap for frameworks that can truly unify multi modal data with domain knowledge in a structured, explainable manner.

2.2. Anomaly detection according to data type

2.2.1. Time Series Data:

Time series data remain the most widely used input for *LIB* anomaly detection. Voltage, current, and temperature signals are routinely monitored, and derived quantities such as *SOC*, *SOH*, and internal resistance are used as diagnostic indicators. Classical statistical methods and ML/DL architectures (e.g., *LSTMs*, *ARIMA*, *Kalman* filters) have been applied to detect irregular patterns. These methods are attractive because time series data are abundant, inexpensive to collect, and directly linked to performance metrics.

However, the limitation lies in their lack of spatial resolution: localized phenomena such as uneven heating or internal short circuits may remain invisible until they manifest globally. Additionally, anomalies that evolve subtly over time, such as lithium plating, may be confounded with normal aging processes, leading to false positives or delayed detection.

2.2.2. Thermal Imaging:

Thermal imaging, particularly via infrared (IR) cameras, has emerged as a complementary technique for anomaly detection. Unlike point sensors, thermal imaging provides a two dimensional map of surface temperatures, enabling detection of spatial heterogeneities such as hotspots, thermal gradients, or uneven cell heating in battery packs. Such anomalies are strong precursors to catastrophic events like thermal runaway.

The advantages of thermal imaging lie in its ability to detect spatial anomalies that are invisible to time series sensors. However, it is sensitive to acquisition and environmental conditions, such as external airflow, emissivity variations, or camera placement, which can significantly influence thermal readings. Moreover, thermal images alone lack temporal resolution, making it difficult to distinguish transient fluctuations from genuine anomalies without complementary time-series data.

The reviewed approaches are further synthesized in Table 1, which highlights their main characteristics across methodological categories. Beyond distinguishing which data types are supported, the table emphasizes the balance between interpretability, and the degree of knowledge integration. This comparative perspective illustrates the trade offs faced by existing methods and points to opportunities for hybrid approaches that combine multimodal data with ontology-based reasoning.

The interpretability and explainability levels reported in Table 1 reflect the extent to which each anomaly detection technique provides understandable and actionable insights. Techniques labeled High allow direct tracing of outputs to battery domain knowledge or explicit rules, as in Equivalent Circuit Models, Electrochemical Models, and ontology-based reasoning, where each prediction can be clearly linked to physical parameters or logical rules. Methods rated Moderate or Partial offer limited interpretability, often through intermediate representations such as latent states in Kalman Filters or saliency maps in visual anomaly models (e.g., DRAEM, GAnomaly), providing some insight but requiring expert knowledge for full comprehension. Techniques marked Low are typically black-box models, including deep learning architectures (LSTM, CNN, Autoencoders) and classical ML algorithms (SVM, Random Forest), which provide accurate anomaly detection but little transparency regarding the underlying reasoning. These classifications were assigned based on the degree of correspondence between the model outputs and understandable battery concepts or rules, balancing both technical rigor and practical explainability.

2.2.3. Research Gap:

While extensive research has explored anomaly detection in lithium ion batteries, major gaps persist. Model-based approaches are interpretable but lack scalability, and data-driven methods achieve high accuracy yet remain opaque and hard to generalize. Knowledge-based ontologies improve explainability but typically focus on electrochemical structures rather than dynamic or multimodal behavior. Moreover, the integration of thermal imaging with time series sensor data has rarely been formalized within a semantic framework; most existing sensor fusion studies rely on statistical or machine learning correlations without explicit causal or spatial reasoning. Ontology based systems also tend to be static, limiting their capacity for real time analysis of evolving battery states. To address these limitations, this work introduces the *Battery Anomaly Ontology (BAO)* a modular, extensible, and multimodal

framework that combines static knowledge and stream reasoning to unify diverse data sources, including sensor, temporal, spatial, and visual information for interpretable anomaly detection in lithium-ion batteries.

Table 1
State of Art of Anomaly Detection Techniques in Lithium-Ion Batteries (LIBs).

Ref.	Approach	Technique	Time Series	Thermal Images	Explainable	Interpretability	Knowledge Integration	Remarks
[33, 42]	**Model-Based**	Equivalent Circuit Models (ECM)	✓	✗	✓	✓	✗	Fully interpretable, physics-based, equation-driven model
[24, 47]		Electrochemical Models (EChM)	✓	✗	✓	✓	✗	Physics-based model, fully explainable
[1, 32]		Kalman Filters (EKF, UKF, PF)	✓	✗	✓	✗	✗	Partially explainable, interpretability limited to filter state updates
[3, 20, 23, 37, 38]	**Data-Driven**	Support Vector Machines (SVM)	✓	✗	✗	✗	✗	Black-box, no inherent explanation or interpretability
[4, 14, 18, 38]		Random Forest / k-NN	✓	✗	✗	✗	✗	Decision-making not transparent for large ensembles
[22, 30]		LSTM / RNN	✓	✗	✗	✗	✗	Black-box recurrent models
[16, 21, 26]		CNN	✗	✓	✗	✗	✗	Black-box visual models
[27]		Autoencoders	✓	✗	✗	✗	✗	Black-box latent representation
[29]		PatchCore (memory-bank embeddings)	✗	✓	✗	✗	✗	Black-box embedding-based model
[8]		PaDiM (patch distribution modeling)	✗	✓	✗	✗	✗	Black-box patch-based model
[43]		FastFlow	✗	✓	✗	✗	✗	Black-box normalizing flow model
[41]		STFPM	✗	✓	✗	✗	✗	Black-box visual model
[44]		DRAEM	✗	✓	✓	✗	✗	Post-hoc visual explanations provided
[31]		FUAUAD	✗	✓	✓	✗	✗	Partial explanation through reconstruction errors
[2]		GAnomaly	✗	✓	✓	✗	✗	Partial visual explanations
[15, 17]	**Knowledge-Based**	Rule-Based / Expert Systems	✓	✗	✓	✓	✗	Fully interpretable rules, logic-based
[5, 12, 13, 35]		Ontology-Based Reasoning	✓	✗	✓	✓	✓	Transparent reasoning, fully explainable
[36, 45]		Knowledge Graphs	✓	✓	✓	✓	✓	Transparent knowledge graph reasoning
[19, 39]	**Hybrid**	Physics-Informed Machine Learning	✓	✗	✓	✗	✗	Partially explainable, combines physics with black-box ML
[10, 40]		Sensor Fusion (time-series + images)	✓	✓	✗	✗	✗	Black-box fusion, no inherent interpretability
[7, 9, 25]		Ontology + ML Reasoning	✓	✓	✓	✗	✓	ML provides detection, ontology provides partial explanation

3. Theoretical Foundations

This section introduces the theoretical foundations of the proposed framework. We first define the role of ontologies in knowledge representation, then describe the principles of stream reasoning with C-SPARQL, and finally present the battery specific state indicators that are central to anomaly detection.

Definition 1 (Ontology). An ontology is formally defined as a quadruple:

$$O = \langle C, P, Sub, Applic \rangle$$

where:

- C is the set of *classes* representing the concepts of a domain,
- P is the set of *properties* (object and data properties) describing the relationships and attributes of individuals belonging to classes in C,

- $Sub : C \rightarrow 2^C$ is the *subsumption function* such that for $c \in C$, $Sub(c)$ denotes the set of subclasses; a class c_1 subsumes c_2 iff $\forall x \in c_2, x \in c_1$,
- $Applic : C \rightarrow 2^P$ associates each class with the set of properties that apply to its individuals.

Ontologies provide a *formal and shared representation of knowledge*, enabling semantic interoperability and logical reasoning. In this work, the ontology is expressed in OWL 2, grounded in Description Logics, and extended with modules to represent batteries, sensors, temporal constraints, and anomalies.

Definition 2 (Stream Reasoning). Stream reasoning integrates logical inference with continuous data streams, allowing knowledge bases to evolve over time.

Definition 3 (C-SPARQL Query). A C-SPARQL query is defined as:

$$Q = \langle S, W, q \rangle$$

where S is an RDF stream, W is a sliding window operator that extracts finite subsequences of S, and q is a SPARQL query executed over the windowed data combined with background ontologies.

C-SPARQL provides the theoretical foundation for combining *static background knowledge* (ontology) with *dynamic observations* (sensor streams), supporting the specification of anomaly patterns that depend on both semantic constraints and temporal conditions.

Definition 4 (State of Health SOH). SOH measures the condition of a battery relative to its initial state:

$$SOH(\%) = \frac{C_{\text{current}}}{C_{\text{original}}} \times 100$$

where C_{current} is the current maximum capacity and C_{original} is the rated beginning-of-life capacity. A declining SOH indicates degradation and reduced performance.

Definition 5 (State of Charge SOC). SOC represents the available charge relative to the nominal capacity:

$$SOC(\%) = \frac{Q_{\text{remaining}}}{Q_{\text{nominal}}} \times 100$$

where $Q_{\text{remaining}}$ is the remaining charge and Q_{nominal} is the nominal full capacity. SOC acts as an instantaneous "fuel gauge" for the battery.

Both indicators are crucial for monitoring safety and reliability, and deviations from their expected behavior are strong indicators of potential anomalies.

4. Methodology

The proposed methodology integrates heterogeneous data streams within a semantic reasoning framework to enable anomaly detection and enhance interpretability in *LIBs*. The general architecture of the approach is illustrated in Figure 1. The workflow follows a sequential process: raw measurements from sensors and thermal cameras are first collected and preprocessed to ensure quality and consistency. These heterogeneous data are then semantically translated into *RDF* streams, aligning them with the concepts and properties defined in the ontology. The ontology serves as the semantic backbone of the system, providing both domain knowledge and formal structure for reasoning. Continuous queries are executed over the *RDF* streams to detect abnormal temporal or spatial patterns, and identified anomalies are instantiated in the ontology.

4.1. Data Layer

The Data Layer is responsible for the acquisition and preprocessing of raw battery data, which includes both time series sensor readings and thermal imaging streams. Its objective is to convert heterogeneous signals into semantically enriched, time stamped representations suitable for reasoning.

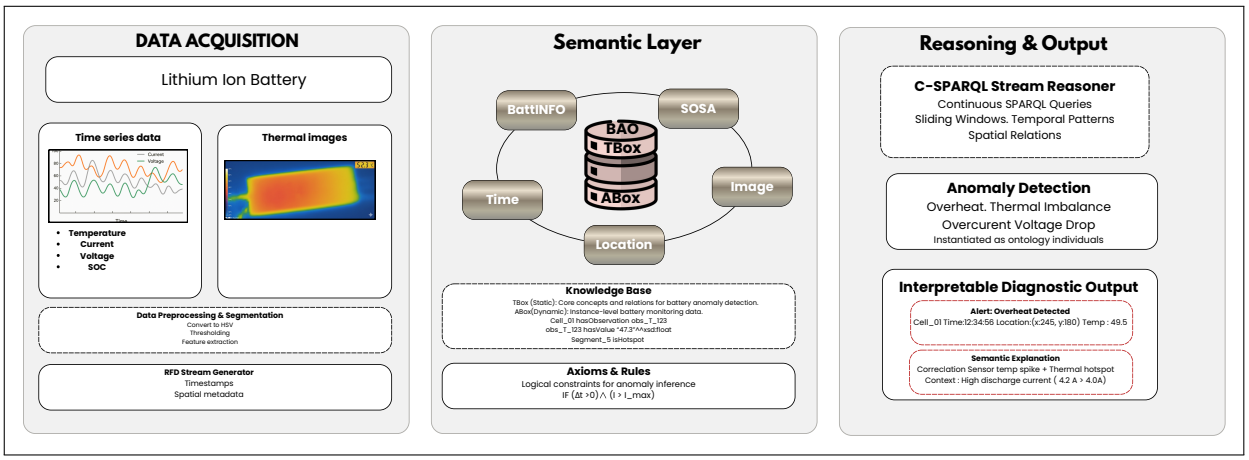

Fig. 1. An overview of the proposed architecture.

4.1.1. Sensor Data Acquisition

Raw sensor data are continuously collected from battery packs, capturing critical parameters:

- Temperature sensors monitor thermal states of cells or modules.
- Current sensors record charging and discharging profiles.
- Voltage sensors measure cell and pack voltages.

4.1.2. Thermal Imaging and Segmentation

Thermal imaging provides a spatial perspective on battery behavior, capturing two dimensional heat maps that reveal heterogeneity such as hotspots and thermal gradients. Unlike scalar sensor readings, thermal images enable early identification of localized degradation or failure risks.

The segmentation pipeline presented in Figure 2 is proceeds as follows:

- Preprocessing: Thermal images are first cropped to isolate the battery cell. Pixel level temperature values are then estimated by mapping RGB intensities to temperature using the image's color scale. To improve accuracy, a neural network is trained on RGB temperature pairs, achieving a mean absolute error *(MAE)* of 0.0086. The resulting dataset stores each pixel's RGB value alongside its corresponding temperature in a structured DataFrame.
- Segmentation: The preprocessed images are transformed into the HSV color space, which provides a clearer separation of hue and intensity than RGB. Thresholding is applied to highlight potential hotspot regions (e.g., red or white areas), and morphological operations are used to suppress noise and refine region boundaries. Contours are then extracted, filtered by size, and divided into distinct segments.
- Feature Extraction: For each identified segment, a set of descriptive attributes is computed to characterize the thermal and spatial properties of the region. The centroid coordinates (x,y) provide the precise spatial location of the segment within the image, enabling the mapping of thermal anomalies to specific areas of the battery cell. The area of the segment quantifies its size, offering insight into the extent of the hotspot or abnormal region. The dominant color, typically red or white, serves as a qualitative indicator of heat intensity and complements the quantitative measurements. Finally, the average temperature of the segment is calculated by aggregating pixel level values, providing a robust thermal descriptor that captures the overall intensity of heating within the region. Together, these attributes form a comprehensive representation of localized anomalies, supporting subsequent semantic annotation and reasoning tasks.

Fig. 2. Thermal imaging and segmentation processing.

4.1.3. Data integration

The Translation and Stream Generation module unifies heterogeneous inputs time series sensor readings and thermal image segments into a single semantic stream. Raw values are converted into RDF triples, where each observation is represented by subject–predicate–object relations. For example, a temperature sensor reading is modeled as a `sosa:Observation` linked to a `BatteryCell` and annotated with its value and `timestamp`, while a thermal hotspot segment is represented as a `ThermalSegment` with attributes such as centroid, area, dominant color, and average temperature.

This process includes semantic enrichment, mapping raw data to ontology defined concepts and properties so that values are not only standardized but also contextualized (e.g., linking a high current reading to an `OvercurrentEvent`). The resulting RDF streams are then used to dynamically populate the ontology's ABox with new individuals that include sensor IDs, spatial coordinates, and temporal metadata. In this way, the knowledge graph is updated in near real time, reflecting the evolving state of the battery.

4.2. Ontology layer

At the heart of the framework lies the Battery Anomaly Ontology (BAO), developed as an extension of the ontology introduced by [46]. While the original model focused on electrochemical knowledge, sensors, time, and anomaly classes, BAO extends it with a dedicated Image module to integrate thermal imaging data alongside time series sensor information, enabling multimodal reasoning for anomaly detection in lithium-ion batteries.

4.3. The BAO Ontology Building Approach

The development of the Battery Anomaly Ontology (BAO) followed a systematic, modular, and iterative ontology engineering approach as described in Figure 3, inspired by best practices from *the NeOn and methodology frameworks* [34]. The goal was to construct a reusable and extensible ontology capable of formally representing the multifaceted knowledge involved in anomaly detection for lithium-ion batteries.

The process began with a requirement specification and conceptual analysis phase, where the main entities, relationships, and reasoning needs were identified. Competency questions were defined to ensure that the ontology could address essential diagnostic queries such as "Which sensors detected a rapid temperature rise in a specific cell at time t?" or "What type of anomaly can be inferred from correlated current and temperature increases?" This phase relied on domain expertise, literature analysis, and prior work reported in [46].

Following this, the ontology was designed according to a reuse-oriented and modular paradigm. Existing standards were integrated to ensure interoperability and semantic consistency:

- SOSA/SSN for representing sensors and observations.
- W3C Time for temporal aspects.
- Geo for spatial context.
- BattINFO for domain specific electrochemical knowledge.

BAO extends these foundations with a newly introduced Image module, enabling the integration of thermal imaging data alongside time series sensor streams. Together, six interlinked modules `BattINFO`, `Sensor`, `Time`, `Location`, `Image`, and `Anomaly` compose the ontology's architecture. Each module captures a distinct yet complementary perspective of the domain, supporting fine grained reasoning over structural, temporal, spatial, and causal dimensions.

The ontology was implemented in OWL 2 using Protégé and Owlready2. The TBox defines the conceptual schema, including class hierarchies, axioms, and property constraints, while the ABox is dynamically populated at runtime with RDF data streams derived from both numerical sensor readings and segmented thermal images.

The ontology was developed through iterative cycles of design, validation, and evaluation to ensure consistency and alignment with real-world cases. This process led BAO to mature into a robust ontology supporting both static and stream reasoning for anomaly detection.

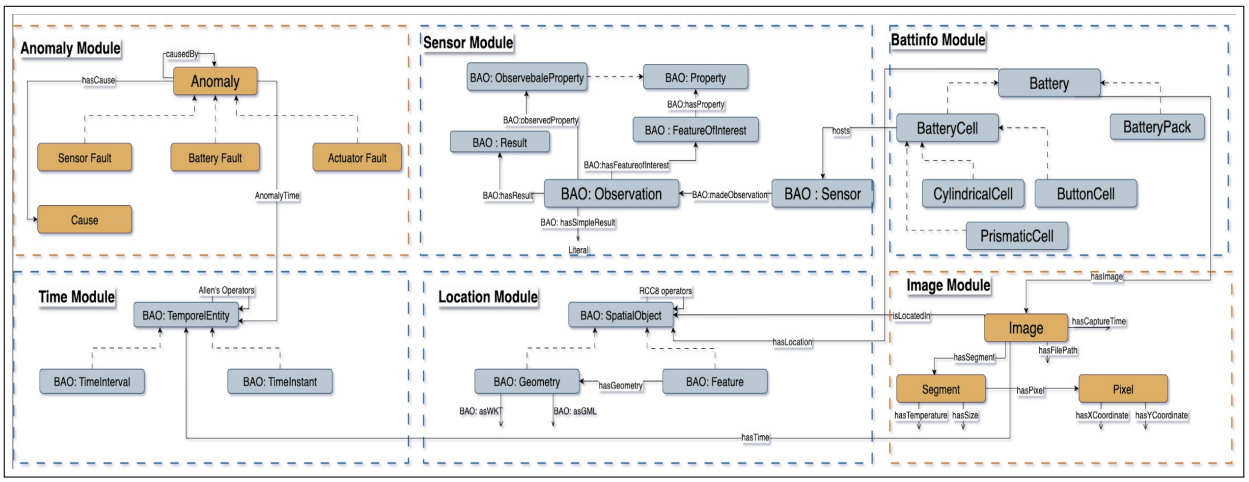

Fig. 3. An overview of the BAO ontology.

4.3.1. *BattINFO Module:*

The `BattINFO` module defines the structural and electrochemical organization of batteries. It introduces classes such as `Battery`, `BatteryPack`, and `BatteryCell`, with subclasses for different formats (e.g., `CylindricalCell`, `PrismaticCell`, `ButtonCell`). Relations like `hasPart` allow hierarchical decomposition of packs into cells, enabling reasoning about anomalies at varying levels of granularity. This module also captures operational states such as SOC, modeled through the class SOC and the data property `SOCValue`. By linking SOC values to specific cells, `BattINFO` provides contextual knowledge crucial for detecting and interpreting abnormal behavior.

`BattINFO` thus acts as the structural backbone of BAO, providing the reference entities to which other modules connect: sensors are hosted by cells, thermal image segments are mapped to cells, temporal events describe their state evolution, and anomalies are classified by their impact on batteries, packs, or cells. In this way, `BattINFO` ensures that all observations and anomalies are grounded in the physical structure of the battery system.

4.3.2. *Time Module:*

The Time module provides a formal representation of temporal aspects essential for reasoning over dynamic battery behavior. It reuses the W3C Time ontology to describe temporal entities such as `time:Instant` and `time:Interval`, enabling precise annotation of sensor observations. Through object properties like `hasTime`, `beginsAt`, and `endsAt`, events and measurements are temporally contextualized, supporting queries such as detecting anomalies within specific time windows or analyzing temperature evolution across cycles. This temporal grounding allows BAO to integrate seamlessly with C-SPARQL queries, where sliding windows operate over instances of `time:Instant` to perform continuous reasoning on incoming RDF streams.

4.3.3. *Location Module:*

The Location module establishes the spatial grounding necessary to link observations with the physical battery structure across multiple scales (pack → module → cell → surface region → image/pixel). It builds upon the OGC GeoSPARQL and W3C Basic Geo ontologies to represent geometric entities and topological relations, enabling BAO to formally specify the spatial context of observations and to align two dimensional image coordinates with real-world battery coordinates. The module introduces the core classes `SpatialEntity` (a generic spatial bearer), `Footprint2D` (a cell or pack surface polygon), and `BoundingBox` and CentroidPoint (compact geometric summaries). Geometries are associated with spatial entities via the `geo:hasGeometry` property and expressed as `geo:asWKT` literals, while topological relations such as `geo:sfWithin`, `geo:sfOverlaps`, and `geo:sfTouches` support reasoning about spatial adjacency and containment (e.g., identifying a hotspot region located within a cell footprint). To bridge image and physical space, the module introduces lightweight alignment terms: `projectsFromImage` links a Segment (Thermal module) to the corresponding `Footprint2D`.

Through hasLocation, every sosa:Observation and thermal Segment can be anchored to a SpatialEntity (typically a BatteryCell footprint from BattINFO), enabling spatio-temporal queries (e.g., "find hotspots overlapping Cell-12 between t1 and t2") and causal rules that depend on where anomalies manifest (edge vs. core regions, inter-cell boundaries, etc.). This spatial layer is thus the glue between SOSA observations, the thermal image regions, and the BattINFO components, supporting precise, interpretable anomaly localization.

4.3.4. SOSA/Sensor Module:

The SOSA (Sensor, Observation, Sample, and Actuator) module captures the sensing infrastructure and observation process. It models the entities involved in data collection, including sosa:Sensor, sosa:Observation, and sosa:ObservableProperty, providing a standard vocabulary for representing real world sensor readings. In BAO, physical sensors such as TemperatureSensor, CurrentSensor, and VoltageSensor are subclasses of sosa:Sensor, each linked to their corresponding observations via the sosa:madeObservation relation. Each observation records its result sosa:hasSimpleResult, unit, and timestamp, and is connected to the observed battery cell through sosa:observedProperty. This alignment with SOSA ensures semantic reasoning across heterogeneous sensor networks.

4.3.5. Anomaly Module:

The Anomaly module formalizes the taxonomy of abnormal conditions that can occur in lithium-ion batteries. It defines the class hierarchy under bao:Anomaly, with subclasses such as Overheat, ThermalImbalance, and Overcurrent, representing distinct anomaly categories derived from safety constraints and operational limits. Logical restrictions are encoded to associate anomalies with their causes, symptoms, and affected components, using properties like hasCause, affectsCell, and riggeredByObservation. For instance, an Overheat anomaly may be inferred when a temperature observation exceeds a defined threshold and is temporally aligned with a current surge. This module provides the reasoning backbone for anomaly detection, allowing both rule-based and data-driven evidence to be represented and explained within a unified semantic framework.

4.3.6. Thermal Image Module:

The Thermal Image module represents the major extension introduced in this work, enabling BAO to integrate thermal imaging data alongside time-series sensor observations. While previous ontology versions focused primarily on electrochemical and temporal knowledge, this new module introduces spatially rich visual concepts that allow the ontology to capture and reason over thermal patterns observed in infrared (IR) images.

Figure 4 presents an overview of the Thermal Image module within BAO, illustrating its main classes, subclass hierarchy, and semantic relations. At the top level, the class ThermalImage specializes the more general Image concept and is associated with its acquisition context through the properties acquiredBy linking to ThermalCamera and acquiredDuring linking to ImageAcquisition. The classes ColorMap and ColorBar describe the image palette used to map pixel intensities to temperature values, ensuring consistency between visual color scales and quantitative thermal data.

Each image is decomposed into multiple Segment instances through the hasSegment property, allowing localized analysis of thermal regions. Segments are characterized by spatial and thermal descriptors centroidX, centroidY, areaPx, meanTemp, and maxTemp which quantify the location, size, and temperature distribution of the region. The subclass HotspotSegment represents regions whose temperature exceeds a safety threshold, serving as semantic indicators of potential overheating or imbalance events. More specifically, each Pixel object stores its (x, y) location and temperature (pixelTemp)of individual image points, enabling pixel level anomaly analysis.

The module also defines the class PixelTemperatureObservation, which aligns with the SOSA pattern for observations. This ensures that both sensor based and image based measurements can be uniformly described and queried. The property observedPixel links each observation to the corresponding pixel, while SurfaceTemperature defines the observable property under investigation.

By connecting Segment and Pixel entities to the physical components defined in the BattINFO module (e.g., specific cells or modules), the Thermal Image module bridges spatially rich image data with the physical battery model. This linkage allows BAO to support spatio temporal reasoning for instance, detecting that a

`HotspotSegment` on a particular cell corresponds to an overheat event captured in the same time window as a high current observation.

Overall, this module transforms BAO into a multimodal ontology capable of representing, aligning, and reasoning over both numerical sensor data and thermal imagery, providing a powerful foundation for interpretable anomaly detection in lithium-ion batteries.

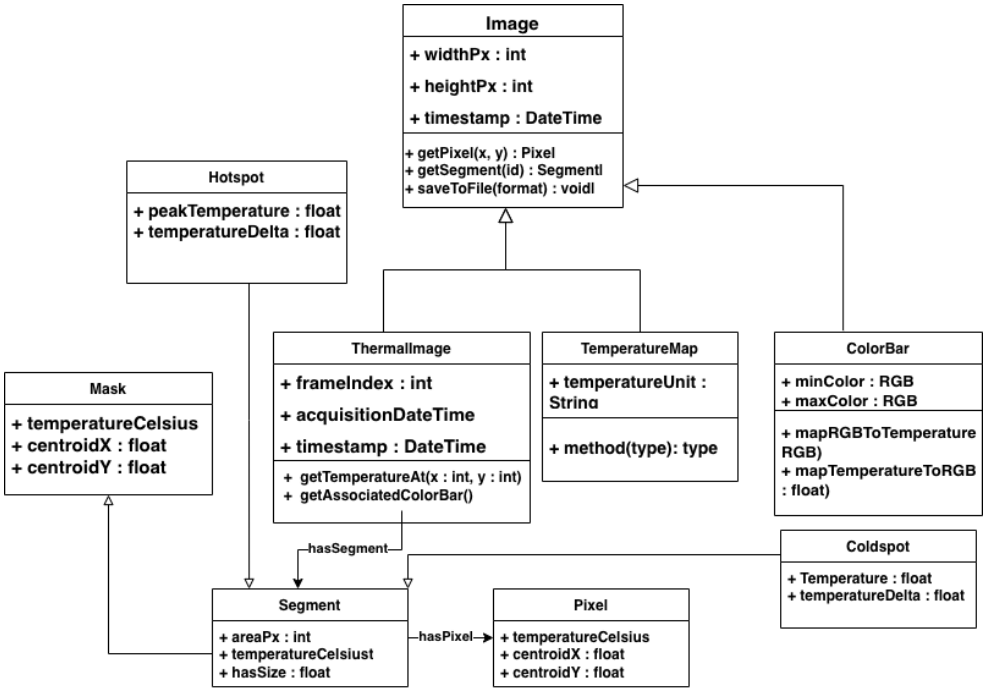

Fig. 4. Part of the thermal image module.

4.4. Stream Reasoning Layer

The Stream Reasoning Layer is the operational core of the framework, where continuous semantic inference is performed over real time sensor and image data using the Battery Anomaly Ontology (BAO) as a background knowledge base. Building upon the static conceptual knowledge defined in the Ontology Layer, this component transforms dynamic sensor readings and segmented thermal images into structured RDF streams, enabling temporal and spatial reasoning within a unified semantic context.

The process begins with the raw data stream generator, which converts time series measurements such as temperature, current, and voltage from individual sensors into RDF triples. Each observation is annotated with its timestamp, the corresponding sensor, and the physical entity being observed. For instance, a temperature sensor attached to a given cell may report a value that is represented as a `sosa:Observation` linked to a `TemperatureSensor` via `sosa:madeObservation`, with its result encoded through `sosa:hasSimpleResult` and the measurement time represented as a `time:Instant`. In parallel, thermal images captured by a `ThermalCamera` are processed and segmented, producing instances of `ThermalImage`, `Segment`, and `HotspotSegment`, each enriched with spatial and statistical attributes such as centroid coordinates, area, mean temperature, and maximum temperature. These image derived entities are streamed into the reasoning layer using the same RDF structure, ensuring interoperability between numeric and visual data streams.

Once the RDF streams are active, the C-SPARQL stream reasoner continuously evaluates persistent queries over sliding time windows. Each query encodes a semantic temporal pattern derived from BAO's axioms and rules, allowing the system to detect deviations that correspond to known anomaly types. For instance, a query may identify

a temperature increase greater than 5 °C within 10 s in any segment classified as a `HotspotSegment`, or detect a simultaneous voltage drop and temperature spike across correlated sensors. By combining temporal evolution (how sensor readings change over time) with spatial relationships (how thermal segments map to specific battery cells), the reasoner detects events such as `Overheat`, `ThermalImbalance`, or `Overcurrent`.

The BAO ontology provides the semantic grounding necessary to interpret these events. It encodes causal and contextual relationships linking `TemperatureSensor` and `ThermalCamera` observations to battery components (`BatteryCell`, Module, Pack), and associating observed behaviors with specific anomaly classes and safety thresholds. Thus, when the stream reasoner identifies a rule match, it can instantiate an anomaly individual (e.g., `bao:Overheat`) enriched with contextual links, including the affected cell, triggering observation, and temporal interval of occurrence.

Crucially, this layer performs multimodal reasoning by synchronizing the time-series and image streams using shared temporal `time:Instant` and spatial `hasLocation` anchors. This enables the system to correlate localized thermal hotspots with concurrent electrical deviations, providing a comprehensive, contextualized assessment of battery health. For example, a temperature rise detected in both a cell level sensor and a neighboring thermal segment may indicate a developing overheat condition, whereas an increasing spatial temperature gradient across segments may signal a propagation risk or thermal imbalance.

By continuously integrating numerical and visual data with ontological knowledge, the stream reasoning layer transforms BAO into a adynamic diagnostic framework. It not only detects anomalies as they emerge but also enriches the knowledge base with detailed semantic descriptions of each event, paving the way for interpretability, trend analysis, and predictive diagnostics in future iterations of the system.

5. Proof of Concept

To assess the effectiveness of our framework, we perform a comprehensive proof of concept evaluation focusing on three distinct yet potentially interrelated anomaly types: overheat situation, thermal imbalances, and optical reflection artifacts. Each anomaly type poses unique detection challenges, highlighting different strengths of our semantic reasoning approach. This structured validation demonstrates how BAO enables accurate and interpretable anomaly detection across diverse battery failure modes and environmental conditions.

5.1. Experimental Setup and Data Source

The experimental validation was conducted on commercial lithium-ion pouch cells (model SPIM11309102-GL40) designed for plug-in hybrid electric vehicle (PHEV) applications. All tests were performed on beginning of life (BOL) cells under controlled laboratory conditions using a multi-channel battery testing system placed inside a temperature regulated chamber maintained at 25 ± 2 °C, in accordance with the manufacturer's standard test conditions.

Each cell was subjected to standard charge–discharge cycling consisting of constant-current/constant-voltage (CC–CV) charging and constant-current (CC) discharging within the operating voltage window of 2.5–4.2 V. Electrical measurements, including cell voltage, current, and temperature, were continuously recorded at a 3 Hz sampling rate to capture the dynamic electro-thermal behavior during operation.

Time series temperature measurements were obtained using three contact temperature sensing points distributed along the pouch cell. Sensors were positioned at the positive terminal region, the negative terminal region, and the geometric center of the cell. This configuration enabled the monitoring of longitudinal temperature gradients and localized thermal deviations during charge–discharge cycling. All temperature, voltage, and current time series data were acquired synchronously at a 3 Hz sampling rate under controlled environmental conditions.

In parallel, Thermal imaging data were collected using a FLIR E96 infrared camera (640 × 480 pixels) positioned above the test cell. Infrared images were captured at regular intervals during cycling to visualize temperature distribution and hotspot evolution. A fixed color scale ranging from 15 °C to 60 °C was applied to all thermal images to ensure consistent comparison across operating conditions.

The resulting dataset combines synchronized electrical time series data (voltage, current, and temperature) with infrared thermal imagery, and includes three classes of anomalies: overheating events, thermal imbalance along the cell length, and artifacts caused by spatial tape reflection. This dataset forms the experimental basis for validating the proposed framework.

5.2. Detection of an Overheat Event

Overheat detection was evaluated using the experimental dataset described in the previous section. The analysis focuses on identifying abnormal thermal behavior during charge-discharge cycling by exploiting synchronized electrical time series data and infrared thermal imagery. The proposed detection mechanism relies on semantic reasoning within the Battery Anomaly Ontology (BAO) to associate observed temperature elevations with an `bao:OverheatEvent`.

Table 2 provides a summary of the experimental elements relevant to overheat detection and their corresponding semantic representations within BAO.

Table 2
Experimental Setup and BAO Mapping

Parameter	**Experimental Detail**	**BAO Semantic Class & Identifier**
Cell Under Test	Lithium-ion pouch cell (SPIM11309102-GL40, 40 Ah, 3.65 V nominal)	`bao:PouchCell :Cell_40Ah_01`
Test Configuration	CC–CV charge and CC discharge between 2.5–4.2 V at 25 °C	`bao:TestProcedure :CyclingTest_01`
Time-Series Sensors	Voltage, current, and temperature measurements (3 Hz)	`sosa:VoltageSensor :VS_01`
		`sosa:CurrentSensor :CS_01`
		`sosa:TemperatureSensor :TS_01`
Thermal Imaging	Infrared images (FLIR E96, 640×480 px, fixed 15–60 °C scale)	`bao:ThermalCamera :IR_Cam_01`
Detected Anomalies	Overheating, thermal imbalance, optical reflection artifacts	`bao:OverheatEvent`
		`bao:ThermalImbalanceEvent`
		`bao:MeasurementArtifact`

An overheat situation is formally defined as a condition where elevated temperature and current levels persist beyond safe operational limits, with spatial confirmation through thermal imaging. Mathematically, we define:

- T(t): Cell temperature at time t.
- I(t): Cell current at time t.
- $T_{\text{threshold}}$: Temperature safety threshold.
- $I_{\text{threshold}}$: Current safety threshold.
- $\Delta_{\text{persistence}}$: Minimum persistence interval.
- $P_{\text{thermal}}(t)$: Thermal propagation metric from imaging.

The overheat condition occurs when :

$$t \in [t_0, t_0 + \Delta t_{\text{persistence}}] : T(t) > T_{\text{threshold}} \wedge I(t) > I_{\text{threshold}} \wedge P_{\text{thermal}}(t) > P_{\text{threshold}}$$

To validate whether the proposed framework is capable of detecting such an overheating situation, we formulate a set of Framework Validation Questions (FVQs). These questions are not intended to assess the ontology in isolation, but rather to verify that the entire framework—including data representation, temporal reasoning, spatial interpretation, and inference logic can reliably identify and explain overheating events.

Table 3 summarizes the validation questions, the ontology concepts mobilized by the framework, and how each question is addressed in the overheat detection case study.

Table 3
Framework Validation Questions for Overheat Detection

FVQ	Framework Validation Question	Ontology Support	Detection Logic (Query Excerpt)	Model Output (C-SPARQL Result)
FVQ1	Can the framework formally define an *Overheat* event as a conjunction of temperature, current, and thermal propagation constraints?	`bao:OverheatEvent`, `bao:hasIndicator`	`FILTER(?T > T_thr && ?I > I_thr && ?T > T_thr)`	`(cell_01, OverheatEvent_17)`
FVQ2	Can the model represent simultaneous observations of temperature, current, SOC, and thermal imagery for the same cell?	`sosa:Observation`, `bao:affectsCell`	`JOIN TempObs, CurrObs, SOCObs, ThermalObs`	`(cell_01, T=38.6, I=4.5, SOC=0.72, T=6.1)`
FVQ3	Can the model represent the persistence of an anomalous condition over a time window?	`time:Interval`	`RANGE 60s STEP 5s`	`(cell_01, t_start=1240s, t_end=1300s)`
FVQ4	Can the model distinguish sustained overheating from transient thermal spikes?	`bao:OverheatEvent`	`HAVING(COUNT(?t) > n)`	`(cell_01, duration=60s)`
FVQ5	Can overheating be distinguished from normal thermal gradients by exceeding a defined temperature threshold?	`bao:TemperatureThreshold`	`FILTER(?T > T_thr)`	`(T=38.6 > 35.0)`
FVQ6	Can overheating be spatially validated using thermal imagery rather than point sensing alone?	`bao:ThermalImageObservation`, `bao:HotspotRegion`	`FILTER(?T > T_thr)`	`(T=6.1)`
FVQ7	Can the framework report all causal observations contributing to overheating detection?	`bao:triggeredByObs`	`SELECT ?obs`	`(TempObs_21, CurrObs_21, ImgObs_08)`
FVQ8	Can the detected overheating event be temporally localized with start and end times?	`time:hasBeginning`, `time:hasEnd`	`BIND window start/end`	`(t_start=1240s, t_end=1300s)`

The evaluation demonstrates that BAO satisfies all competency questions required for multimodal, real time overheat detection. The ontology is capable of: (i) integrating heterogeneous sensor streams, (ii) encoding temporal and spatial constraints, (iii) distinguishing transient from persistent anomalies, and (iv) producing interpretable and causally grounded anomaly events. The case study validates that BAO is structurally and semantically aligned with the requirements of thermal–electrical anomaly detection in lithium-ion batteries.

The detection logic is encoded as a continuous C-SPARQL query. This query subscribes to the fused data streams and performs stream reasoning over a sliding window whose duration matches $\Delta t_{persist}$ (5 seconds). The query is structured to join observations from all three modalities (temperature, current, and thermal area) pertaining to the same cell.

The C-SPARQL query for detecting $A_{overheat}$ is shown in Listing 1 as follows:

Listing 1: C-SPARQL query for detecting *Overheat* events

```
REGISTER QUERY Overheat AS

PREFIX anomaly: <http://example.org/ns#>
PREFIX sosa:    <http://www.w3.org/ns/sosa/>

SELECT ?c ?vTemp1 ?vThermal ?vCurrent ?vSoc
FROM STREAM <Stream_TempSensor1>     [RANGE 60s STEP 5s]
FROM STREAM <Stream_ThermalSensor_1> [RANGE 60s STEP 5s]
FROM STREAM <Stream_CurrentSensor>   [RANGE 60s STEP 5s]
FROM STREAM <Stream_Soc1>            [RANGE 60s STEP 5s]
FROM <http://example.org/ontology>

WHERE {
  ?c anomaly:hosts anomaly:TempSensor1 ;
```

```
    anomaly:hosts anomaly:ThermalSensor_1 ;
    anomaly:hosts anomaly:CurrentSensor ;
    anomaly:hosts anomaly:Soc1 .

  anomaly:TempSensor1 anomaly:madeObservation ?oTemp1 .
  ?oTemp1 anomaly:hasSimpleResult ?vTemp1 .

  anomaly:ThermalSensor_1 anomaly:madeObservation ?oThermal .
  ?oThermal anomaly:hasSimpleResult ?vThermal .

  anomaly:CurrentSensor anomaly:madeObservation ?oCurrent .
  ?oCurrent anomaly:hasSimpleResult ?vCurrent .

  anomaly:Soc1 anomaly:madeObservation ?oSoc .
  ?oSoc anomaly:hasSimpleResult ?vSoc .

  FILTER (
    (?vCurrent > 4.2) &&
    (?vTemp1   > 45.0) &&
    (?vThermal > 45.0) &&
    (?vSoc     > 0.5)
  )
}
```

When all constraints were satisfied within the defined temporal window, the reasoning layer instantiated a new RDF individual representing the detected bao:Overheat event. The generated instance captured the causal observations, measured indicato rs, and temporal bounds of the event, as illustrated in Listing 2.

Listing 2: RDF instance generated for an Overheat event

```
:Overheat_01
    rdf:type bao:Overheat ;
    bao:affectsCell :Cell_01 ;
    bao:triggeredByObservation (
        :Obs_TempSensor1_t1
        :Obs_ThermalSensor_1_t1
        :Obs_CurrentSensor_t1
        :Obs_Soc1_t1
    ) ;
    bao:hasIndicators (
        :TempSensor1_Value_01
        :ThermalSensor_1_Value_01
        :Current_Value_01
        :Soc1_Value_01
    ) ;
    bao:hasTimeInterval [
        time:hasBeginning "2025-11-07T10:15:00Z"^^xsd:dateTime ;
        time:hasEnd       "2025-11-07T10:16:00Z"^^xsd:dateTime
    ] ;
    bao:hasCondition [
        bao:tempSensor1Value   "33.2"^^xsd:float ;
        bao:thermalTempValue   "34.1"^^xsd:float ;
        bao:currentValue       "4.3"^^xsd:float ;
        bao:socValue           "0.67"^^xsd:float
    ] .
```

Table 4 presents representative outcomes of overheat detection for the NMC pouch cell. The results demonstrate that the ontology based framework accurately identifies thermal anomalies by combining electrical measurements and thermal imaging data within a semantically rich model.

Table 4
Result of an overheat detection.

Timestamp (s)	Temp_1 (°C)	Temp_2 (°C)	Temp_3 (°C)	I (A)	V (V)	$T_{thermal}$ (°C)	Image
15.0	46.8	47.1	46.8	4.5	3.7	47.0	
25.0	47.1	47.8	47.3	4.9	3.6	47.8	
30.0	50.6	50.2	50.4	5.3	3.5	50.8	

At early timestamps (e.g., 15 s), although the cell exhibits moderate temperature increases, the ontology does not trigger an overheat event, reflecting its ability to consider both the magnitude and persistence of anomalies. As the test progresses (30 s and 40 s), the system detects a sustained overheat condition, fully instantiating an active `bao:Overheat` event when elevated temperature, high current, and significant hotspot area persist concurrently.

The incorporation of thermal imagery is particularly valuable, providing spatial context through hotspot area measurements. This enables the ontology to detect localized overheating that could otherwise be missed by electrical data alone. Additionally, the temporal reasoning implemented via sliding windows ensures that only persistent anomalies are flagged, preventing transient fluctuations from generating false detections.

Overall, the results confirm that the ontology-driven approach provides robust, interpretable, and early detection of overheat events. By integrating multi modal observations with temporal reasoning, the system captures both the underlying causes and the duration of anomalies, providing a reliable and actionable basis for proactive battery monitoring.

5.3. Detection of Thermal Imbalance Event

To further illustrate how the proposed ontology supports anomaly detection in lithium ion batteries, we apply it to a representative thermal imbalance scenario. This case study demonstrates how synchronized thermal and electrical observations can be interpreted to identify early deviations from expected behaviour.

5.3.1. Experimental Setup and Data Source

The thermal imbalance study was conducted using the same experimental platform and instrumentation described for the overheat analysis, including the FLIR infrared imaging system and the standard electrochemical cycling protocol. In this case, however, the setup was intentionally modified to introduce controlled spatial non uniformities on the cell surface.

To create regions with different emissivity, most of the pouch cell was coated with matte black paint, while a small area was left masked to form a reflective patch. This design induces artificial thermal irregularities across the surface, enabling the observation of lateral temperature gradients and localized hotspot behaviour. The infrared camera, positioned above the cell as in the previous tests, captured time lapse thermal images throughout cycling, allowing continuous monitoring of these spatial variations.

Unlike the overheat scenario which focuses on global temperature rise this configuration emphasizes relative differences across the cell surface. The resulting thermal maps provide the spatial detail needed to detect imbalance patterns indicative of non uniform heat dissipation or potential early stage failure mechanisms.

A thermal imbalance is formally defined as a condition in which spatial temperature differences across the cell exceed safe operational limits, potentially leading to uneven aging and localized stress. Let us define:

- $T_i(t)$: Temperature of the i-th cell region at time t.
- $\Delta T_{ij}(t) = |T_i(t) - T_j(t)|$: Temperature difference between regions i and j.
- $T_{\text{diff,thresh}}$: Maximum allowable temperature differential.
- $\Delta_{\text{persistence}}$: Minimum persistence interval.

The thermal imbalance condition occurs when:
$t \in [t_0, t_0 + \Delta_{\text{persistence}}] : \Delta T_{ij}(t) > T_{\text{diff,thresh}}$

This condition is monitored using the ontology driven framework, where time series measurements are instantiated as `TempObservation` individuals, and thermal imaging processing generates `ThermalImageObservation` individuals with `hasHotspotArea` and `hasTemperatureGradient` properties.

5.3.2. Ontology Evaluation Through Competency Questions for Thermal Imbalance

To evaluate BAO's capability for thermal imbalance detection, we define a set of competency questions directly linked to the definition and detection of thermal imbalance. These questions focus on spatial temperature differences, gradients, affected regions, persistence, and event instantiation.

Table 5: Framework Validation Questions for Thermal imbalance Detection

FVQ	Competency Question	Ontology Concepts Involved	Support in Case Study
FVQ1	Can BAO represent temperature differences between multiple regions of the same cell?	`bao:TempObservation`, `bao:CellRegion`, `bao:hasTemperatureGradient`	Yes. Each cell region is instantiated and linked to its measurements; gradients are computed as properties.
FVQ2	Can BAO detect thermal imbalance only when temperature differences persist over a defined interval?	`bao:ThermalImbalance`, `time:Interval`, `bao:TempObservation`	Yes. Sliding windows enforce persistence criteria before instantiating an event.
FVQ3	Can BAO specify which cell regions are involved in a detected thermal imbalance?	`bao:affectsCell`, `bao:CellRegion`, `bao:ThermalImbalance`	Yes. The ThermalImbalance instance references all affected regions.
FVQ4	Can BAO capture the magnitude of lateral temperature gradients as part of the ThermalImbalance event?	`bao:ThermalImbalance`, `bao:hasTemperatureGradient`, `bao:TempObservation`	Yes. Gradient values are stored as properties of the event instance.
FVQ5	Can BAO detect localized thermal imbalance even if global cell temperature is nominal?	`bao:ThermalImbalance`, `bao:CellRegion`, `bao:TempObservation`	Yes. Gradients above the threshold in specific regions trigger events early.
FVQ6	Can BAO integrate temperature observations from multiple regions to compute overall imbalance?	`bao:TempObservation`, `bao:ThermalImbalance`, `bao:hasTemperatureGradient`	Yes. Observations from all regions are fused to determine the event.
FVQ7	Can BAO instantiate a ThermalImbalance event in RDF when all conditions (gradient, persistence, regions affected) are met?	`bao:ThermalImbalance`, `bao:hasIndicators`, `bao:triggeredByObservation`	Yes. The event captures affected regions, gradient magnitude, and causal observations.
FVQ8	Can BAO distinguish thermal imbalance caused by persistent gradients from short-lived local hotspots?	`bao:ThermalImbalance`, `bao:HotspotRegion`, `bao:hasTemperatureGradient`	Yes. Persistent lateral differences trigger the event; transient local spikes do not.
FVQ9	Can BAO support continuous detection of thermal imbalance from streaming temperature data?	`bao:TempObservation`, `bao:ThermalImageObservation`, `bao:ThermalImbalance`	Yes. C-SPARQL queries over fused streams detect imbalance continuously.

This evaluation confirms that BAO fully supports the definition and detection of thermal imbalance. It enables reasoning over lateral gradients, persistence, affected regions, and event instantiation, allowing early detection of spatial anomalies even when global cell temperature remains within safe limits.

The detection logic is encoded as a continuous C-SPARQL query in Listing 3, subscribing to fused temperature streams and thermal imaging streams. The reasoning is performed over a sliding temporal window matching $\Delta_{\text{persistence}}$ (5 seconds), and temperature differences between cell regions are computed.

Listing 3: C-SPARQL query for detecting *Thermal Imbalance* events

```
REGISTER QUERY ThermalImbalance AS

PREFIX anomaly: <http://example.org/ns#>
PREFIX sosa:    <http://www.w3.org/ns/sosa/>

SELECT ?c ?vTemp1 ?vTemp2 ?vTemp3 ?vThermal ?vGrad
FROM STREAM <Stream_TempSensor1>       [RANGE 60s STEP 5s]
FROM STREAM <Stream_TempSensor2>       [RANGE 60s STEP 5s]
FROM STREAM <Stream_TempSensor3>       [RANGE 60s STEP 5s]
FROM STREAM <Stream_ThermalSensor_1>   [RANGE 5s STEP 2s]
FROM <http://example.org/ontology>

WHERE {
  ?c anomaly:hosts anomaly:TempSensor1 ;
     anomaly:hosts anomaly:TempSensor2 ;
     anomaly:hosts anomaly:TempSensor3 ;
     anomaly:hosts anomaly:ThermalSensor_1 .

  anomaly:TempSensor1 anomaly:madeObservation ?o1 .
  ?o1 anomaly:hasSimpleResult ?vTemp1 .

  anomaly:TempSensor2 anomaly:madeObservation ?o2 .
  ?o2 anomaly:hasSimpleResult ?vTemp2 .

  anomaly:TempSensor3 anomaly:madeObservation ?o3 .
  ?o3 anomaly:hasSimpleResult ?vTemp3 .

  anomaly:ThermalSensor_1 anomaly:madeObservation ?oThermal .
  ?oThermal anomaly:hasSimpleResult ?vThermal .

  ?seg :hasTemperatureGradient ?vGrad .

  FILTER (
    ((ABS(?vTemp1 - ?vTemp2) > 3.0) || (ABS(?vTemp1 - ?vTemp3) > 3.0)) &&
    (?vThermal > 45.0)
  )
}
```

Once all constraints are satisfied within the temporal window, a new RDF individual representing a `bao:ThermalImbalance` event is instantiated, capturing both the measurements and affected cell regions As shown in Listing 4.

Listing 4: RDF instance generated for a Thermal Imbalance event

```
:ThermalImbalance_01 rdf:type bao:ThermalImbalance ;
    bao:affectsCell :Cell_01 ;
    bao:triggeredByObservation (
        :Obs_TempSensor1_t1
        :Obs_TempSensor2_t1
        :Obs_TempSensor3_t1
        :Obs_ThermalSensor_1_t1
    ) ;
    bao:hasIndicators (
        :TempSensor1_Value_01
        :TempSensor2_Value_01
        :TempSensor3_Value_01
        :ThermalSensor_1_Value_01
        :Thermal_Gradient_01
    ) ;
    bao:hasTimeInterval [
        time:hasBeginning "2025-11-07T10:15:00Z"^^xsd:dateTime ;
        time:hasEnd       "2025-11-07T10:16:00Z"^^xsd:dateTime
```

```
    ] ;
    bao:hasCondition [
        bao:tempSensor1Value "35.2"^^xsd:float ;
        bao:tempSensor2Value "28.7"^^xsd:float ;
        bao:tempSensor3Value "30.1"^^xsd:float ;
        bao:thermalTempValue "33.5"^^xsd:float ;
        bao:temperatureGradient "6.3"^^xsd:float
    ] .
```

Table 6 presents representative examples of thermal imbalance detection. The ontology-driven framework successfully captures lateral temperature disparities, avoiding false alarms from transient or localized fluctuations. Early snapshots with moderate gradients (e.g., 15 s) do not trigger an event, whereas persistent deviations above the threshold (e.g., 30 s and 40 s) correctly instantiate `bao:ThermalImbalance` events.

Table 6
Result of a thermal imbalance detection.

Timestamp (s)	Temp_1 (°C)	Temp_2 (°C)	Temp_3 (°C)	I (A)	V (V)	$T_{thermal}$ (°C)	Image
1.0	26.2	28.0	22.1	4.2	3.0	22.2	
28.0	42.0	48.1	45.5	4.2	3.7	45.0	
28.0	36.2	48.0	52.1	4.5	3.8	46.2	
59.0	32.5	49.1	49.8	5.0	3.9	49.5	

6. Results and Evaluation

A rigorous evaluation is essential to validate the claims of semantic integration and multimodal anomaly detection presented in this work. To assess the conceptual robustness and practical effectiveness of the proposed Battery Anomaly Ontology, we employ the *TOMM* framework, a methodology for evaluating engineering ontologies across structural, semantic, and functional dimensions. Unlike metric only assessments, *TOMM* examines ontology modules individually, enabling a detailed diagnosis of their internal coherence, interaction patterns, and modular boundaries. This analysis is particularly important for *BAO*, whose architecture combines heterogeneous knowledge domains electrochemical structure, sensing, temporal reasoning, spatial alignment, and thermal imaging into a multi module ontology. Applying *TOMM* at the module level allows us to uncover deficiencies in modularity, such as overlapping semantic roles, insufficient axiom isolation, or dependencies that compromise reuse and maintainability.

6.1. Ontology Evaluation Using TOMM

We evaluated our ontology module using the **Tool for Ontology Module Metrics (TOMM)**, which computes structural, relational, and logical metrics to assess modularity quality. Each metric is briefly defined below, with its corresponding formula:

- **Size ($|M|$):** Total number of entities (classes C, object properties OP, data properties DP, individuals I):

$$\text{Size}(M) = |C| + |OP| + |DP| + |I|$$

- **Appropriateness:** How close the module size is to optimal, based on axioms x:

$$\text{Appropriateness}(x) = \frac{1}{2} - \frac{1}{2}\cos\left(\frac{x\pi}{250}\right)$$

- **Atomic Size:** Average size of interdependent sets of axioms:

$$\text{Atomic Size}(M) = \frac{\sum_i \text{Atom}_i}{|M|}$$

– **Attribute Richness (AR):** Average number of attributes per class:

$$AR(M) = \frac{|\text{attributes}|}{|C|}$$

– **Inheritance Richness (IR):** Average number of subclasses per class:

$$IR(M) = \frac{\sum_{C_i \in C} |HC(C_i)|}{|C|}$$

– **Intra-module Distance (IMD):** Sum of shortest-path distances between entities (Freeman's Farness):

$$\text{IMD}(M) = \sum_{i=1}^{n} \sum_{j=1}^{n} \text{distance}_{ij}$$

– **Cohesion:** Degree of relatedness among entities. We split the definition for clarity in a narrow column:

$$\text{Cohesion}(M) = \frac{\sum_{c_i, c_j \in M} SR(c_i, c_j)}{|M|(|M|-1)},$$

$$\text{where } SR(c_i, c_j) = \begin{cases} \frac{1}{\text{farness}(i)}, & \text{if relation exists} \\ 0, & \text{otherwise} \end{cases}$$

– **Independence:** Module self-containment based on encapsulation and coupling:

$$\text{Independent if Encapsulation}(M) = 1$$

$$\text{and Coupling}(M) = 0$$

Table 7
Ontology module evaluation metrics using TOMM.

Metric	Result	Interpretation
Size ($\|M\|$)	158	Moderate size, manageable
Appropriateness	0.877	Near optimal size
Atomic Size	2.97	Small, manageable groups
Attribute Richness (AR)	1.8	Classes are descriptive
Inheritance Richness (IR)	2.88	Moderate hierarchy depth
Intra-module Distance (IMD)	6544.0	Entities moderately dispersed
Cohesion	0.86	High cohesion, strong interconnections
Independence	High	Fully encapsulated, low coupling

The TOMM results for our ontology module are summarized in Table 7.The ontology module demonstrates a moderate and appropriate size with good attribute richness and a moderate inheritance hierarchy. The high cohesion (0.86) indicates strong interconnections among entities, supporting efficient reasoning and knowledge retrieval. The module also exhibits high independence, being fully encapsulated with minimal coupling, which facilitates modular updates without impacting the rest of the ontology. Overall, these results confirm the module's strong modularity, structural quality, and semantic clarity, making it well suited for downstream anomaly detection applications.

6.2. Performance Assessment with Anomalib and Merlion Baselines

This section presents the evaluation of the proposed ontology-based anomaly detection framework and its comparison with data-driven baselines. We first describe the evaluation methodology and anomaly score derivation, followed by the ground-truth annotation strategy and anomaly categories, and finally the performance metrics used for quantitative assessment.

6.2.1. Evaluation Methodology

The evaluation procedure follows the algorithmic process detailed in Algorithm 1. For each sliding window in the test data stream, C-SPARQL queries are evaluated over the Battery Anomaly Ontology (BAO) to determine whether predefined safety constraints are violated. These ontology-based rules yield discrete violation indicators for temperature, current, and thermal gradient anomalies.

For each window, a continuous anomaly score is computed by aggregating the active violations using expert-defined severity weights and their temporal persistence within the window. The anomaly score is initialized to zero and incremented for each violated constraint according to its corresponding severity weight (0.4 for temperature, 0.3 for current, and 0.3 for thermal gradient) multiplied by the fraction of the window duration during which the violation persists. This procedure implements the following formulation:

$$S_{\text{anomaly}}(w) = \sum_i \text{severity\_weight}_i \cdot \text{violation\_indicator}_i \cdot \text{persistence\_factor}_i \tag{1}$$

The resulting anomaly score is stored alongside the expert-provided ground-truth label associated with each sliding window. After processing the full test stream, binary anomaly predictions are generated by thresholding the continuous anomaly scores at $\tau = 0.5$, yielding detected or non-detected anomaly labels.

6.2.2. Ground Truth and Anomaly Categories

The experimental dataset was annotated by domain experts into three anomaly categories reflecting distinct abnormal operating conditions:

- **Overheating events:** Cell temperature exceeding safety thresholds in conjunction with elevated current and persistent thermal propagation.
- **Thermal imbalance events:** Non-uniform temperature distributions across the cell surface, characterized by lateral thermal gradients exceeding operational limits.
- **Sensor artifacts:** Spurious thermal readings caused by measurement errors such as optical reflections or sensor noise.

For evaluation purposes, all annotated anomaly categories are treated as a single positive class, while normal operating conditions are labeled as the negative class. Sensor artifact cases highlight the sensitivity of the proposed framework: although they do not correspond to true battery health anomalies, they demonstrate the system's ability to detect deviations from expected behavior. Future work will incorporate artifact-specific ontology rules to improve precision.

6.2.3. Evaluation Metrics

Detection performance is quantified using standard classification metrics. Precision, recall (true positive rate), F1-score, and accuracy are computed based on the binary anomaly predictions and expert-provided ground-truth labels. In addition, the Area Under the Receiver Operating Characteristic curve (AUROC) is computed directly from the continuous anomaly scores by varying the detection threshold.

AUROC provides a threshold-independent measure of performance, enabling fair comparison with probabilistic machine learning baselines. The complete evaluation procedure, including anomaly score computation and metric calculation, is summarized in Algorithm 1.

Algorithm 1 Anomaly Score Computation and Evaluation Metrics

Require: Test data stream $\mathcal{W}$, expert labels $\mathcal{Y}$
Ensure: Precision, Recall, F1-score, Accuracy, AUROC

Initialize anomaly score list $\mathcal{S} \leftarrow [\,]$
Initialize ground-truth list $\mathcal{G} \leftarrow [\,]$
for each sliding window $w \in \mathcal{W}$ **do**
  Evaluate C-SPARQL queries on w to obtain violation set V
  Initialize anomaly score $S \leftarrow 0$
  if $V_{\text{temp}} = 1$ **then**
    $S \leftarrow S + 0.4 \times p_{\text{temp}}$
  end if
  if $V_{\text{current}} = 1$ **then**
    $S \leftarrow S + 0.3 \times p_{\text{current}}$
  end if
  if $V_{\text{gradient}} = 1$ **then**
    $S \leftarrow S + 0.3 \times p_{\text{gradient}}$
  end if
  Append S to $\mathcal{S}$
  Append expert label $y_w \in \{0, 1\}$ to $\mathcal{G}$
end for
Generate binary predictions:

$$\hat{y}_i = \begin{cases} 1, & \text{if } S_i \geqslant \tau \\ 0, & \text{otherwise} \end{cases}$$

where $\tau = 0.5$
Compute Precision, Recall, and F1-score using $\mathcal{G}$ and $\hat{\mathcal{Y}}$
Compute Accuracy using $\mathcal{G}$ and $\hat{\mathcal{Y}}$
Compute AUROC using $\mathcal{G}$ and continuous scores $\mathcal{S}$
return Precision, Recall, F1-score, Accuracy, AUROC

For comparison, all baseline models were trained on clean, normal battery data using a 70/30 split: 70% of the normal data was used for training, while the remaining 30% served as part of the test set. The test set additionally included anomalous samples to evaluate model performance under abnormal operating conditions. To further assess robustness, we separately evaluated the baselines on the anomalous data subset.

Time-series baselines (Merlion) were trained on sensor streams including temperature, current, and voltage and evaluated on the test set containing both normal and anomalous time-series.

Thermal imaging baselines (Anomalib: PatchCore, PaDiM) were trained exclusively on normal thermal images and tested with images exhibiting anomalous thermal patterns.

All performance metrics precision, recall, F1-score, accuracy, and AUROC were computed using the same ground-truth labels as the ontology-based framework. This ensures a fair and consistent comparison between the ontology-based and data-driven approaches.

Table 8
Evaluation of Time-Series and Thermal Imaging Baselines vs. OntologyBased Framework.

Model / Configuration	Precision	Recall	F1-Score	AUROC
Time-Series Baselines (Merlion)				
AutoEncoder	0.881	0.852	0.866	0.912
LSTM Detector	0.905	0.873	0.889	0.931
ARIMA	0.768	0.821	0.794	0.951
Thermal Imaging Baselines (Anomalib)				
PatchCore	0.861	0.872	0.866	0.890
PaDiM	0.748	0.835	0.788	0.796
EfficientAD	0.952	0.846	0.896	0.912
Ontology-Based Framework (BAO)				
Multimodal (Time + Thermal)	0.905	0.880	0.892	0.935

6.2.4. Disscusion

The evaluation results reveal complementary strengths between the ontology-based framework and data-driven baselines, reflecting distinct approaches: knowledge-based deductive reasoning versus data-driven inductive learning.

The multimodal ontology approach achieves strong performance through knowledge-based deductive reasoning, which delivers high interpretability and low data dependency. Each detection can be traced back to explicit expert rules providing physically meaningful explanations for instance, "Cell temperature exceeded 45°C while current surpassed 3A, with Segment A showing a 2°C rise within 30 seconds." This interpretability is critical in safety critical applications where understanding why an anomaly was detected enables informed corrective actions.

However, the framework's detection is bounded by its encoded knowledge. Anomalies outside predefined rules such as novel degradation mechanisms or subtle deviations not crossing explicit thresholds may go undetected. For example, gradual temperature drift with unusual temporal dynamics that remains below thresholds would not trigger rule-based constraints. This highlights the fundamental limitation of deductive reasoning: the system can only detect what has been explicitly anticipated and formalized.

In contrast, data-driven baselines, employ inductive learning that offers scalability and adaptability, particularly when expert knowledge is incomplete or when facing previously unseen situations. These models develop implicit representations of normality and detect anomalies as statistical outliers, surfacing novel failures that experts may not have anticipated such as rare thermal patterns or complex multivariate correlations invisible to threshold-based rules. However, this generalization comes at the cost of interpretability and potential false positives from statistically rare but operationally normal variations.

7. Conclusion and Future Work

This work presented an ontology based framework for multimodal battery anomaly detection, integrating time series sensor data with thermal imaging through semantic reasoning. The Battery Anomaly Ontology (BAO) encodes domain expertise and physical constraints, enabling interpretable detection of predefined safety violations. TOMM evaluation confirms strong modularity and semantic clarity, while experimental results demonstrate competitive performance against state of the art baselines.

The multimodal approach achieved strong detection performance, effectively identifying overheating events, thermal imbalance conditions, and sensor artifacts through rule based reasoning. The framework's key strength lies in its interpretability and grounding in validated physical principles, providing high-confidence detection with clear explanations.

However, comparative analysis revealed complementary detection paradigms. While the ontology framework excels at detecting codified safety violations, data-driven models demonstrate superior capability in generalizing to unseen scenarios. These findings suggest neither paradigm alone provides comprehensive coverage. Future work should explore hybrid architectures combining semantic reasoning for explainable detection of codified risks with machine learning as a complementary screening layer for novel pattern identification. Such architectures would leverage ontology-based reasoning for interpretable detection of known hazards while employing statistical learning to surface emergent anomalies for expert review and ontology refinement.

Code Availability The stream reasoning queries, ontology, and datasets utilized for anomaly detection, along with the experimental results, are available at [https://github.com/Marwa-Zitouni/SAFE-LIB]

References

[1] Rami Ahmad and Eman H Alkhammash. Online adaptive kalman filtering for real-time anomaly detection in wireless sensor networks. *Sensors*, 24(15):5046, 2024.

[2] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Asian conference on computer vision*, pages 622–637. Springer, 2018.

[3] Polyzois Bountzis, Dimitris Kavallieros, Theodora Tsikrika, Stefanos Vrochidis, and Ioannis Kompatsiaris. A deep one-class classifier for network anomaly detection using autoencoders and one-class support vector machines. *Frontiers in Computer Science*, 7:1646679, 2025.

[4] Xiwen Chen, Zhong Yuan, and Shan Feng. Anomaly detection based on improved k-nearest neighbor rough sets. *International Journal of Approximate Reasoning*, 176:109323, 2025.

[5] David Clinton. Ontology-guided anomaly detection in multi-source data integration systems. 2024.

[6] C. Comanescu et al. Ensuring safety and reliability: An overview of lithium-ion batteries. *Batteries*, 11(1):6, 2024.

[7] Gustavo Alexsandro de Lima and Mara Abel. Ontology-enhanced deep learning framework for anomaly detection in oil and gas production plants. 2024.

[8] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International conference on pattern recognition*, pages 475–489. Springer, 2021.

[9] Kyle DeMedeiros, Abdeltawab Hendawi, and Marco Alvarez. A survey of ai-based anomaly detection in iot and sensor networks. *Sensors*, 23(3):1352, 2023.

[10] Laura Erhan, M Ndubuaku, Mario Di Mauro, Wei Song, Min Chen, Giancarlo Fortino, Ovidiu Bagdasar, and Antonio Liotta. Smart anomaly detection in sensor systems: A multi-perspective review. *Information Fusion*, 67:64–79, 2021.

[11] Zhi-Wei Gao et al. Development and commercial application of lithium-ion batteries in electric vehicles: A review. *Processes*, 13(3):756, 2025.

[12] Franco Giustozzi, Julien Saunier, and Cecilia Zanni-Merk. A semantic framework for condition monitoring in industry 4.0 based on evolving knowledge bases. *Semantic Web*, 15(2):583–611, 2024.

[13] Patrick Hammer, Tony Lofthouse, Enzo Fenoglio, Hugo Latapie, and Pei Wang. A reasoning based model for anomaly detection in the smart city domain. In *Proceedings of SAI Intelligent Systems Conference*, pages 144–159. Springer, 2020.

[14] Joshua S Harvey, Joshua Rosaler, Mingshu Li, Dhruv Desai, and Dhagash Mehta. Explainable unsupervised anomaly detection with random forest. *arXiv preprint arXiv:2504.16075*, 2025.

[15] Markus Heinrich, Arwed Gölz, Tolga Arul, and Stefan Katzenbeisser. Rule-based anomaly detection for railway signalling networks. *International Journal of Critical Infrastructure Protection*, 42:100603, 2023.

[16] Haoqi Huang, Ping Wang, Jianhua Pei, Jiacheng Wang, Shahen Alexanian, and Dusit Niyato. Deep learning advancements in anomaly detection: A comprehensive survey. *IEEE Internet of Things Journal*, 2025.

[17] Raihan Ul Islam. *An Improved Belief Rule-Based Expert System with an Enhanced Learning Mechanism*. PhD thesis, Luleå University of Technology, 2020.

[18] Shujun Ji, Kai Liu, Bo Ling, Jiadong Li, Jinteng Wang, and Xun Ma. A study on improved random forest-based anomaly detection of regional tariff data under distributed photovoltaic access. *Results in Engineering*, page 105920, 2025.

[19] Mudita Kohli and Indu Chhabra. A comprehensive survey on techniques, challenges, evaluation metrics and applications of deep learning models for anomaly detection. *Discover Applied Sciences*, 7(7):784, 2025.

[20] Sibo Li, Yongqin Zhou, Ran Li, and Xu Zhao. Online lithium battery fault diagnosis based on least square support vector machine optimized by ant lion algorithm. *International Journal of Performability Engineering*, 16(10):1637, 2020.

[21] Zhong Li, Yuxuan Zhu, and Matthijs Van Leeuwen. A survey on explainable anomaly detection. *ACM Transactions on Knowledge Discovery from Data*, 18(1):1–54, 2023.

[22] Benjamin Lindemann, Benjamin Maschler, Nada Sahlab, and Michael Weyrich. A survey on anomaly detection for technical systems using lstm networks. *Computers in Industry*, 131:103498, 2021.

[23] Haowen Lu. Evaluating the performance of svm, isolation forest, and dbscan for anomaly detection. In *ITM Web of Conferences*, volume 70, page 04012. EDP Sciences, 2025.

[24] Chao Lyu, Yankong Song, Jun Zheng, Weilin Luo, Gareth Hinds, Junfu Li, and Lixin Wang. In situ monitoring of lithium-ion battery degradation using an electrochemical model. *Applied Energy*, 250:685–696, 2019.

[25] Andrew Ma, Abhir Karande, Natalie Dahlquist, Fabien Ferrero, and N Rich Nguyen. Sensor fusion enhances anomaly detection in a flood forecasting system. *Journal of Sensor and Actuator Networks*, 14(2):34, 2025.

[26] Mohsin Munir, Muhammad Ali Chattha, Andreas Dengel, and Sheraz Ahmed. A comparative analysis of traditional and deep learning-based anomaly detection methods for streaming data. In *2019 18th IEEE international conference on machine learning and applications (ICMLA)*, pages 561–566. IEEE, 2019.

[27] Asif Ahmed Neloy and Maxime Turgeon. A comprehensive study of auto-encoders for anomaly detection: Efficiency and trade-offs. *Machine Learning with Applications*, 17:100572, 2024.

[28] Alessandro Neri et al. Fostering lithium-ion battery remanufacturing through industry 5.0. *International Journal on Interactive Design and Manufacturing (IJIDeM)*, pages 1–18, 2025.

[29] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2022.

[30] Nawel Said, Majdi Mansouri, Rami Al Hmouz, and Atef Khedher. Deep learning techniques for fault diagnosis in interconnected systems: A comprehensive review and future directions. *Applied Sciences*, 15(11):6263, 2025.

[31] Abdelrahman Shabayek, Arunkumar Rathinam, Matthieu Ruthven, Djamila Aouada, and Tazdin Amietszajew. Ai-enabled thermal monitoring of commercial (phev) li-ion pouch cells with feature-adapted unsupervised anomaly detection. *Journal of Power Sources*, 629:235982, 2025.

[32] Ming Shi, Lin Sun, Zhongcheng Bi, and Renchu He. A hybrid kalman filter and physics-informed neural network approach for leakage detection and localization in heat exchanger networks. *Computers & Chemical Engineering*, page 109259, 2025.

[33] Yue Song, Jinsong Yu, Jinhan Zhou, Jian Zhang, Diyin Tang, and Zetian Yu. Detection of voltage fault in lithium-ion battery based on equivalent circuit model-informed neural network. *IEEE Transactions on Instrumentation and Measurement*, 73:1–10, 2024.

[34] Mari Carmen Suárez-Figueroa and Asunción Gómez-Pérez. Neon methodology for building ontology networks: a scenario-based methodology. In *Proceedings of the International Conference on Software, Services & Semantic Technologies*, Sofia, Bulgaria, 2009.

[35] Lionel Tailhardat, Yoan Chabot, and Raphaël Troncy. Noria-o: an ontology for anomaly detection and incident management in ict systems. In *European Semantic Web Conference*, pages 21–39. Springer, 2024.

[36] Lionel Tailhardat, Yoan Chabot, and Raphaël Troncy. Anomaly detection using knowledge graphs: A survey for network management and cybersecurity application. 2025.

[37] Ye Tao, Jiongcheng Yan, Enquan Niu, Pengming Zhai, and Shuolin Zhang. An svm-based anomaly detection method for power system security analysis using particle swarm optimization and t-sne for high-dimensional data classification. *Processes*, 13(2):549, 2025.

[38] Phan Thanh Noi and Martin Kappas. Comparison of random forest, k-nearest neighbor, and support vector machine classifiers for land cover classification using sentinel-2 imagery. *Sensors*, 18(1):18, 2017.

[39] Srikanth Thudumu, Philip Branch, Jiong Jin, and Jugdutt Singh. A comprehensive survey of anomaly detection techniques for high dimensional big data. *Journal of big data*, 7(1):42, 2020.

[40] Nitin Tyagi, Sunil S Chavan, Syam Machinathu Parambil Gangadharan, et al. Towards intelligent industrial systems: A comprehensive survey of sensor fusion techniques in iiot. *Measurement: Sensors*, 32:100944, 2024.

[41] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. *arXiv preprint arXiv:2103.04257*, 2021.

[42] Yiming Xu, Xiaohua Ge, Ruohan Guo, and Weixiang Shen. Recent advances in model-based fault diagnosis for lithium-ion batteries: A comprehensive review. *Renewable and Sustainable Energy Reviews*, 207:114922, 2025.

[43] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint arXiv:2111.07677*, 2021.

[44] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8330–8339, 2021.

[45] Andy Zhou, Xiaojun Xu, Ramesh Raghunathan, Alok Lal, Xinze Guan, Bin Yu, and Bo Li. Knowgraph: Knowledge-enabled anomaly detection via logical reasoning on graph data. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 168–182, 2024.

[46] Marwa Zitouni, Franco Giustozzi, Ahmed Samet, and Tedjani Mesbahi. Toward anomaly representation in lithium-ion batteries: An ontology-based approach. *Procedia Computer Science*, 246:1319–1328, 2024.

[47] Bosong Zou, Lisheng Zhang, Xiaoqing Xue, Rui Tan, Pengchang Jiang, Bin Ma, Zehua Song, and Wei Hua. A review on the fault and defect diagnosis of lithium-ion battery for electric vehicles. *Energies*, 16(14):5507, 2023.