Accelerating Advanced Aluminium Alloy Design via Ontology-Driven Semantic Modelling and Large Language Models

Anton Bakuteev[1], Jaijith Sreekantan[2], Karim Houni[3] and Khuram Pervez[4]
1. Principal Data Scientist
2. Senior Principal Data Scientist
3. Head of Digital Innovation,
4. Director Data Science and AI
Emirates Global Aluminium (EGA) - Industry 4.0, United Arab Emirates
Corresponding authors: abakuteev@ega.ae, jsreekantan@ega.ae
https://doi.org/10.71659/icsoba2025-ch009

Abstract

One of the major challenges in advancing discovery of new aluminium alloys is the heterogeneous nature of the process, experimental and simulation datasets. This data is often fragmented with inconsistent terminologies necessitating an integrated semantic modelling framework for robust data harmonisation to derive meaningful insights. In our work, we develop a multi-layer semantic integration architecture that employs standardised ontologies implemented through frameworks such as Resource Description Framework (RDF) and Web Ontology Language (OWL) to systematically encode critical variables including alloy composition, processing parameters, and performance metrics into interoperable semantic entities. This structured approach facilitates precise data aggregation, automated inference, and advanced query capabilities.

To ensure standardization and compatibility with broader material science data, we align with key standards relevant to materials modelling in general and the aluminium industry in specific, such as relevant ISO standards (ISO 3522 and ISO 7722) for aluminium and aluminium alloys casting, European Materials Modelling Ontology (EMMO) for materials modelling, the Materials Design Ontology (MDO) focusing on material structures and compositions, Materials Mechanics Ontology (MMO) capturing mechanical properties, and ChEBI (Chemical Entities of Biological Interest) providing standardised terminology for chemical elements and compounds. This structured approach facilitates precise data aggregation, automated inference, and advanced query capabilities.

Integral to our methodology is the incorporation of a domain-specific large language model (LLM) that operates within these rigorously defined ontologies. The integrated semantic layer enables language models to better interpret complex experimental protocols mitigating hallucinations and improving reliability and enables users to query multiple data sources with natural language. In this work, we present a comprehensive semantic modelling framework that combines standardised ontologies with LLM-driven natural language querying to accelerate and enhance the design of novel aluminium alloys.

Keywords: Large Language Model (LLM), Semantic modelling, Ontology, Aluminium alloys, Web Ontology Language (OWL).

1. Introduction

The development of novel aluminium alloys is a time-consuming, resource-intensive process, traditionally relying on empirical approaches and expert knowledge. The complexity of alloy design stems from the high-dimensional compositional space, intricate processing-structure-property relationships, and the heterogeneity of available data. Accelerating this process requires advanced computational techniques that can integrate disparate data sources while maintaining scientific rigor [1].

In the aluminium industry, especially for cast alloys, data fragmentation and inconsistent terminologies present significant challenges. Information about composition, microstructure, mechanical and thermal properties, processing parameters, and application-specific performance exists in various formats across research papers, industrial reports, internal databases, and simulation outputs. This fragmentation makes it difficult to establish comprehensive relationships between composition, processing, structure, and properties [2].

Traditional data integration approaches often struggle with the semantic complexity of materials science concepts. For example, a property like "strength" can refer to yield, tensile, compressive, or fatigue strength, each measured under specific conditions. Without proper semantic context, data integration leads to erroneous comparisons and conclusions that hinder alloy development [3].

Recent advances in semantic modelling and large language models (LLMs) offer promising solutions. Ontologies provide formal, explicit specifications of shared conceptualisations, enabling precise definition of entities, properties, and relationships in a domain. LLMs excel at natural language understanding and generation, enabling more intuitive interactions with complex knowledge bases. However, their application in scientific domains is often hindered by hallucination, limited domain knowledge, and a lack of specific reasoning capabilities [4].

We present an integrated approach that combines ontology-driven semantic modelling with domain specific LLMs to accelerate aluminium alloy design. Our approach focuses on aluminium cast alloys, critical for aerospace, automotive, and general engineering sectors [5]. We use the Web Ontology Language (OWL) for formalisation and integrate standards such as the Resource Description Framework (RDF), European Materials Modelling Ontology (EMMO) [6], Materials Design Ontology (MDO) [7], Chemical Entities of Biological Interest (ChEBI) [8], and Quantities, Units, Dimensions, and Types (QUDT) [9]. We also adhere to industry standards for aluminium alloys, including ISO 3522 for chemical composition and mechanical properties [10] and ISO 7722 for the global designation system of castings [11]. For all subsequent mentions, we use only the abbreviations.

1.1 Ontologies and LLMs

Ontologies provide a formal, explicit way to define shared conceptualisations, enabling precise definition of entities, properties, and relationships in a domain. LLMs excel at natural language understanding and generation, enabling more intuitive interactions with complex knowledge bases.

1.2 Integrated Approach

The integrated approach combines ontology-driven semantic modelling with domain-specific LLMs to accelerate aluminium alloy design. Our approach focuses specifically on aluminium cast alloys, which are critical for applications in aerospace, automotive, and general engineering sectors.

2. Methodology

Our approach to accelerating aluminium alloy design is built on a multi-layer semantic integration architecture that combines standardised ontologies, graph-based knowledge representation, and LLM-powered natural language interfaces.

The architecture consists of four main layers: (1) data source, (2) ontology, (3) knowledge graph, and (4) application, with a fifth component (5) MCP server [12], as illustrated in Figure 1. Data sources include experimental databases, literature, industrial specifications, simulations, and expert knowledge, all using inconsistent terminologies and formats. The ontology layer provides formal semantic models, extending existing materials science ontologies (EMMO, MSO, CHMO, QUDT) to serve as a semantic bridge. The knowledge graph layer implements the ontology as a property graph, enabling efficient storage, querying, and reasoning [13]. The application layer provides user interfaces and services, including natural language query interfaces, recommendation systems, visualisation tools, and computational design workflows, with LLMs enabling intuitive querying and interaction.

The Model Context Protocol (MCP) server is a critical component that enables LLMs to access and query the knowledge graph directly.

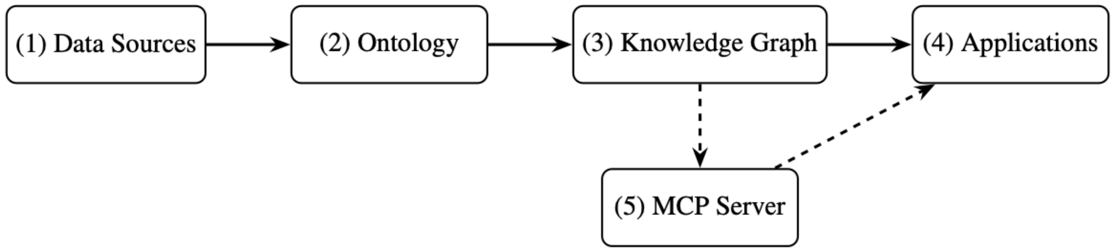

Figure 1. Multi-layer semantic integration architecture for aluminium alloy design.

3. Ontology Design for Aluminium Cast Alloys

The aluminium cast alloys ontology forms the semantic foundation of our approach. It is structured as a hierarchical taxonomy with interconnected modules, including classes for alloys, series, properties, temper, and casting methods. The ontology integrates key external standards (EMMO, QUDT, ChEBI, MDO, MMO) for interoperability.

Key alloy instances (e.g., A201-T7, A356-T6, A319-T6, A413-F) are characterised by comprehensive property sets according to industry standards. The ontology includes over 80 classes, multiple properties, and dozens of alloy instances, providing a robust foundation for advanced querying and knowledge discovery in alloy design.

3.1 Ontology Structure and Hierarchy

The ontology is structured as a hierarchical taxonomy with multiple interconnected modules. Figure 2 illustrates the high-level class hierarchy of the ontology.

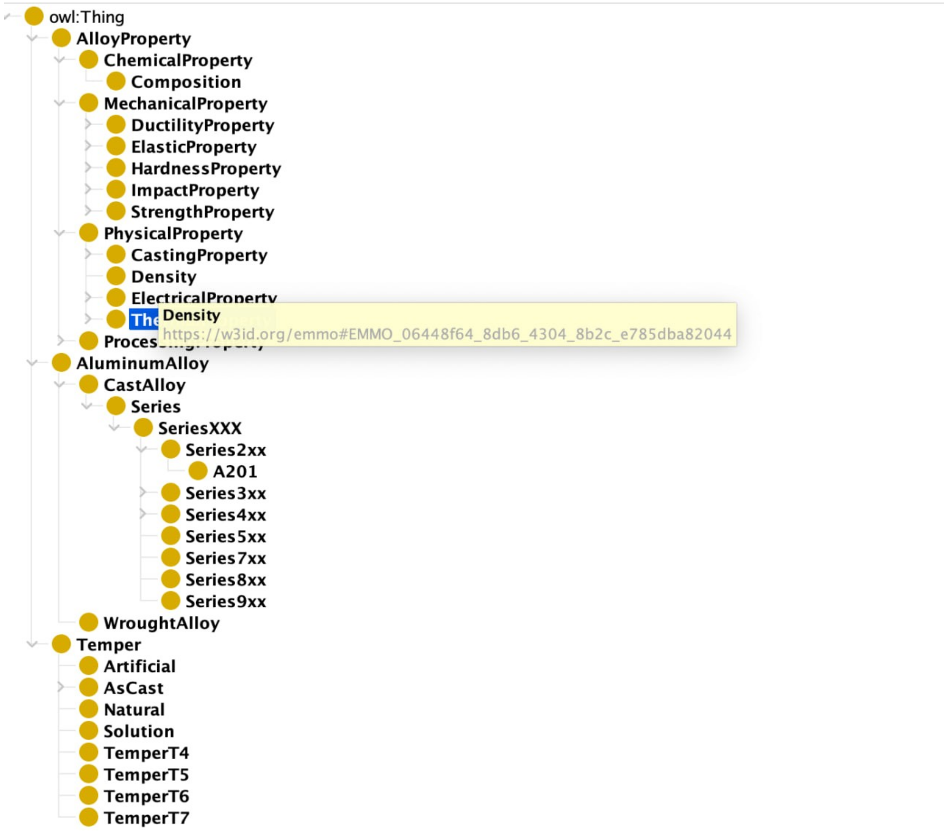

Figure 2. Class hierarchy of the aluminium cast alloys ontology.

The top-level class structure includes:

- **AluminiumAlloy:** The root class for all aluminium alloys, with subclasses:
  - **CastAlloy:** Alloys designed for casting processes
  - **WroughtAlloy:** Alloys designed for wrought processing
- **Series:** Classification of alloys by their primary alloying elements:
  - **Series2xx:** Copper as principal alloying element
  - **Series3xx:** Silicon with copper and/or magnesium
  - **Series4xx:** Silicon as principal alloying element
  - **Series5xx-9xx:** Other principal alloying elements
- **AlloyProperty:** Base class for all properties of alloys:
  - **PhysicalProperty:** Properties like density, thermal conductivity
  - **MechanicalProperty:** Properties like strength, hardness, ductility
  - **ChemicalProperty:** Composition and chemical characteristics
  - **ProcessingProperty:** Properties related to manufacturing
- **Temper:** Heat treatment conditions affecting properties:
  - **AsCast:** As-cast condition (F temper)
  - **Solution:** Solution heat-treated
  - **Artificial:** Artificially aged
  - **Natural:** Naturally aged
- **CastingMethod:** Different casting techniques:
  - **SandCasting:** Using sand moulds
  - **DieCasting:** Using metal moulds under pressure
  - **PermanentMoldCasting:** Using reusable metal moulds
  - **InvestmentCasting:** Using wax patterns

3.2 Specific Alloy Instances

The ontology includes detailed instances of commercially important aluminium cast alloys. Each instance is characterised by a complete set of properties according to industry standards. Figure 3 shows a conceptual representation of an alloy instance (A201-T7) and its associated properties.

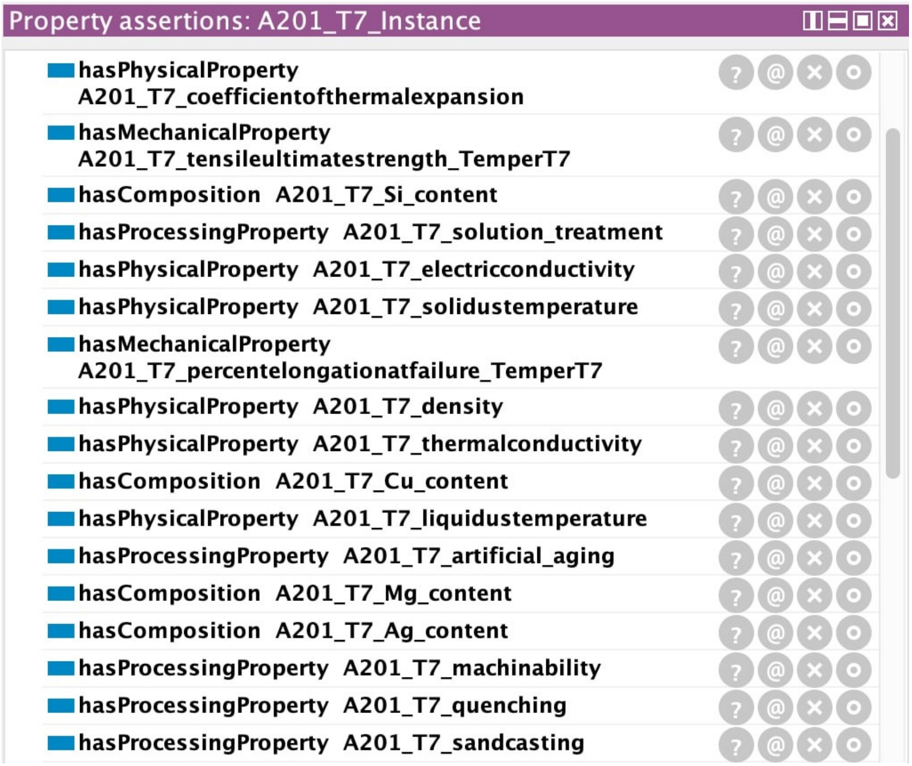

Figure 3. Conceptual representation of the A201-T7 alloy instance and its properties.

Key alloy instances in our ontology include:

- **A201-T7:** High-strength alloy (Al-Cu4.6-Ag-Mn-Ti) used in aerospace applications
- **A356-T6:** Silicon-magnesium alloy (Al-Si7-Mg) with excellent castability
- **A319-T6:** Silicon-copper alloy (Al-Si6-Cu3) used in automotive components
- **A413-F:** High-silicon alloy (Al-Si12) with excellent fluidity

3.3 Integration with External Ontologies

To ensure interoperability with broader materials science data ecosystems, we integrated several external ontologies and standards:

- **EMMO:** Provides foundational concepts for materials modelling, including physical quantities and units.
- **QUDT:** Offers standardised representation of physical quantities and units of measurement.
- **ChEBI:** Provides standardised terminology for chemical elements and compounds.
- **MDO:** Offers concepts related to materials structure and characterisation.
- **MMO:** Provides detailed mechanical properties terminology.

The integration was accomplished through:

1. Importing relevant terms from external ontologies
2. Establishing equivalence and subsumption relationships
3. Mapping properties to standardised definitions
4. Adopting common identification schemes

3.4 Ontology Implementation Approach

We developed a custom Python library for ontology construction to ensure structured, well-documented, and modular code. This approach provides several advantages over traditional ontology development methods:

- LLM-Assisted Code Generation: We leveraged large language models to generate Python code for ontology components. We specifically chose Python due to research showing that LLMs demonstrate superior performance with Python compared to other programming languages. This allowed us to rapidly prototype and iterate on the ontology structure.

- Modular Architecture: The library implements a modular design pattern where each ontology component (classes, properties, individuals) is defined in separate modules, improving maintainability and enabling collaborative development.

- Automated Data Extraction: We implemented web scraping components that parse industry websites and technical documentation to extract alloy properties and specifications. This automated approach populated our ontology with accurate instance data while significantly reducing manual data entry.

- Validation and Testing: The Python framework includes automated validation routines that check for ontological consistency, proper reference to external ontologies, and adherence to design patterns.

The complete ontology includes 83 classes, 6 object properties, 26 data properties, and instances representing dozens of aluminium cast alloys with their full property specifications. This comprehensive semantic model provides the foundation for advanced querying, reasoning, and knowledge discovery in aluminium alloy design.

4. LLM Integration with Semantic Model

Integrating LLMs with our ontology-based semantic model creates a system that combines structured domain knowledge with natural language capabilities.

4.1 Challenges in Applying LLMs to Materials Science

Key challenges include limited domain expertise, hallucination, outdated knowledge, reasoning limitations, and context-dependent terminology. Our approach grounds LLM responses in a formal ontology and verified knowledge graph to address these issues.

4.2 Integration Approach

We use an agentic [15] framework where the LLM interacts with the knowledge graph via the Model Context Protocol (MCP). Our approach provides the LLM with specialised tools to query the database using both SPARQL and Cypher queries, allowing it to retrieve precise information from the ontology and knowledge graph.

- Schema-Guided Interaction: We extract a minimal schema from the ontology, excluding individuals, to give the LLM an understanding of the overall structure without overwhelming its context window. This schema serves as a map that guides the LLM in formulating appropriate queries and interpreting returned results.

- Multi-Hop Graph Traversal: By using the MCP tools, the LLM can traverse the graph through multiple hops, following relationships between entities to answer complex questions that require integrating information from different parts of the knowledge graph. For example, when asked about the relationship between composition and mechanical properties of a specific alloy, the LLM can first query for the alloy's composition, then for its properties, and finally establish the connections between them.

- Tool-Calling Proficiency: This approach requires LLMs with strong abilities in tool calling and sequential reasoning. The model must understand when and how to use different query tools, interpret the returned results, and determine if additional queries are needed to complete the user's request. This capability is particularly important when dealing with complex materials science questions that cannot be answered with a single graph query.

This integrated methodology enables robust semantic integration of aluminium alloy data while providing intuitive natural language interfaces that accelerate the exploration and design of new alloys.

5. Results and Evaluation

Applying our ontology-driven semantic modelling and LLM integration approach to aluminium alloy design, we evaluated the system using competency questions. The ontology successfully addressed all queries, demonstrating comprehensive knowledge representation and practical utility for alloy design tasks.

5.1 Knowledge Representation Evaluation

We evaluated the ontology's ability to represent aluminium cast alloy knowledge through competency questions - specific queries that the ontology should be able to answer. Table 1 shows examples of competency questions and the system's ability to address them.

Table 1. Competency questions for evaluating knowledge representation.

Competency Question	Supported
Which alloys have yield strength above 300 MPa?	Yes
What is the chemical composition of A201-T7?	Yes
Which alloys are suitable for aerospace applications?	Yes
How does heat treatment affect the properties of A356?	Yes
Which casting methods are appropriate for A413?	Yes
What are the thermal properties of high-silicon alloys?	Yes
How do A201 and A356 compare in mechanical properties?	Yes
What processing parameters are used for T6 treatment of A319?	Yes

The ontology successfully addressed all competency questions, demonstrating its comprehensiveness in representing aluminium cast alloy knowledge.

6. Limitations and Challenges

In this section, we outline the main limitations and challenges identified in our approach:

- Ontology construction and automation: Our semi-automated workflow accelerates ontology development by combining LLM-driven parsing with a Python ontology library. However, manual verification and domain expert oversight remain essential to ensure semantic correctness and prevent errors. The Python-centric outputs can also limit adaptability to other frameworks.

- Reasoning and explainability: While the system excel at information retrieval and basic comparative analyses, complex causal reasoning about alloy behaviour demands multi-step inference. Providing transparent explanations and provenance for these reasoning chains remains an open challenge.

- Validation and trust: Ensuring the reliability of system recommendations, especially for high-stakes applications, requires robust validation workflows, transparent auditing of inference steps, and clear provenance tracking.

- Scalability and context limitations: As the ontology schema grows, embedding the entire model in an LLM context window is impractical. Our agentic graph-query approach addresses this but introduces additional overhead in tool calls and multi-hop reasoning. Improving performance, reducing latency, and enhancing tool-use proficiency for large-scale ontology traversal remain critical challenges.

7. Conclusions

We presented an integrated approach to accelerating aluminium alloy design through ontology-driven semantic modelling and LLMs. Our main contributions:

- Development of a comprehensive ontology for aluminium cast alloys using a semi-automated, LLM-assisted workflow.
- Integration of the ontology with a property graph database for advanced semantic querying.
- Enabling natural language querying of the knowledge graph using LLMs.
- Demonstration of the effectiveness with real-world alloy data and use cases.

This integration of semantic technologies with LLMs addresses data integration and knowledge accessibility challenges in materials science, making alloy knowledge more accessible and actionable, and accelerating materials innovation for diverse applications.

8. References

1. Ankit Agrawal and Alok Choudhary, Perspective: materials informatics and big data: realisation of the fourth paradigm of science in materials science, *APL Materials*, 4 (2016) 053208, http://doi.org/10.1063/1.4946894.
2. Toshihiro Ashino, Materials ontology: an infrastructure for exchanging materials information and knowledge, *Data Science Journal*, 9 (2010), 54–61, http://doi.org/10.2481/dsj.008-041.
3. Janna Hastings et al., Chebi in 2016: improved services and an expanding collection of metabolites, *Nucleic Acids Research*, 44 (2016) D1214–D1219, https://www.ebi.ac.uk/chebi/init.do.
4. Emanuele Ghedini et al., EMMO, the European materials and modelling ontology, https://github.com/emmo-repo/EMMO.

5. ISO 3522:2007, Aluminium and aluminium alloys – castings – chemical composition and mechanical properties, *International Organization for Standardization*, 2007.
6. ISO 7722:2000, Aluminium and aluminium alloy castings – global designation system for castings, *International Organization for Standardization*, 2000.
7. J. Gilbert Kaufman, Introduction to aluminium alloys and tempers, *ASM International*, (2000).
8. Youssra El-Korchi et al., The materials design ontology, https://github.com/w3id.org/mdo.
9. Model Context Protocol (MCP), official documentation and introduction, https://modelcontextprotocol.io/introduction.
10. Luca Montanelli et al., High-throughput extraction of phase–property relationships from literature using natural language processing and large language models, *Integrating Materials and Manufacturing Innovation*, 13 (2024), 396–405, http://doi.org/10.1007/s40192-024-00344-8.
11. Ralph Hodgson et al., QUDT – quantities, units, dimensions and types, https://qudt.org.
12. Lukas Twist et al., LLMs love Python: a study of LLMs' bias for programming languages and libraries, arXiv preprint arXiv:2503.17181, 2025, https://arxiv.org/abs/2503.17181.
13. Yanpeng Ye et al., Construction and application of materials knowledge graph in multidisciplinary materials science via large language model, *NeurIPS* 2024, http://doi.org/10.48550/arXiv.2404.03080.
14. Wayne Xin Zhao et al., A survey of large language models, *arXiv* preprint arXiv:2303.18223, (2023), http://doi.org/10.48550/arXiv.2303.18223.
15. Shunyu Yao et al., ReAct: synergizing reasoning and acting in language models, Advances in Neural Information Processing Systems (NeurIPS), 2022, https://arxiv.org/abs/2210.03629.