

ORIGINAL RESEARCH

Open access

The Problem of Competing Scientific Ontologies in Materials Knowledge Representation

Ivan Petrov¹, Olga Ivanova^{1*}, Dmitry Smirnov²

Abstract

Competing scientific ontologies represent a pervasive yet under-analyzed failure mode in artificial intelligence applications for materials science. Different classification systems for the same materials, structures, properties, and processes create fundamental incompatibilities that cause AI models to fail in ways that are difficult to diagnose through conventional performance metrics. This paper defines the ontology problem as the inherent challenge of representing material knowledge when multiple, partially incompatible ontologies coexist within the domain, each encoding distinct conceptual boundaries and relational assumptions. It articulates four primary types of ontological competition—category boundary differences, naming conflicts, relationship differences, and granularity differences—that arise repeatedly in materials informatics. These competitions trigger specific failure modes, including transfer failures, evaluation incompatibilities, data integration failures, and communication breakdowns between research communities. Detection relies on explicit ontology audits and cross-ontology testing, while mitigation centers on mapping strategies, ontology-agnostic representations, and community harmonization efforts. By framing ontology competition as a distinct failure mode, the analysis draws on existing literature to propose an ontology-aware framework that strengthens semantic interoperability and model robustness in materials AI. Ultimately, acknowledging and managing competing ontologies is essential for translating data-driven discoveries into reliable, reproducible knowledge.

Keywords Materials informatics, Failure mode analysis, Competing ontologies, Materials knowledge representation, Ontology mismatch, Semantic interoperability

*Correspondence:

Olga Ivanova

olga.ivanova@outlook.com

¹ Department of Materials Science and AI Systems, Moscow State University, Moscow, Russia

² Department of Smart Materials Engineering, Saint Petersburg State University, Saint Petersburg, Russia

Introduction

Materials AI relies on ontologies—formal ways of categorizing materials, structures, properties, and processes. But materials science has multiple, sometimes competing ontologies. The same material may be classified differently across databases. These mismatches cause AI models to fail in subtle ways. This paper analyzes ontology competition as a failure mode.

The reliance on ontologies is not incidental but foundational to contemporary materials informatics. As machine learning

systems ingest vast repositories of material data, they implicitly or explicitly adopt the categorical structures provided by the underlying knowledge representation schemes [1-4]. Yet the domain of materials science is characterized by a proliferation of parallel classification systems that evolved independently within different research communities, funding initiatives, and data infrastructures. A single compound might be labeled as belonging to one crystal system under the conventions of the NOMAD laboratory [5] while receiving an alternative structural descriptor in the OPTIMADE federation [6-9]. Such discrepancies are not mere terminological

preferences; they encode divergent epistemological commitments about what constitutes a meaningful category in materials space.

These ontological divergences become particularly problematic when AI models are trained on data organized under one scheme and then deployed on data structured according to another. The resulting failures are often invisible to standard validation protocols because the models appear to perform adequately within their native ontological environment [2]. For instance, a model trained to predict electronic properties using a taxonomy that treats certain hybrid organic-inorganic frameworks as a single broad class may produce systematically biased predictions when confronted with datasets that subdivide the same chemical space along different compositional or topological boundaries. The problem is compounded by the fact that many prominent materials databases, including those supporting large-scale machine learning campaigns, operate with partially overlapping but non-identical ontological commitments [1, 6].

Moreover, the rapid expansion of materials data science has amplified the visibility of these issues without necessarily resolving them. Efforts to harmonize ontologies, such as the PMD Core Ontology [8] or the materials data science ontology (MDS-Onto) [6], represent important steps toward semantic interoperability; however, they coexist alongside legacy classification systems and domain-specific taxonomies that continue to shape experimental and computational workflows. The consequence is a fragmented knowledge ecosystem in which AI models risk inheriting hidden inconsistencies that propagate through downstream predictions, database integrations, and even scientific communication. This paper, therefore, treats ontological competition not as a solvable technical inconvenience but as a structural failure mode intrinsic to materials knowledge representation. By dissecting its manifestations, mechanisms, and consequences, the analysis aims to equip the community with conceptual tools for building more resilient AI systems [2, 7].

Figure 1 presents a hierarchical ontology-failure architecture showing how competing classification systems generate specific mechanisms of mismatch, produce observable failure modes in materials AI, and motivate corresponding detection and mitigation responses.

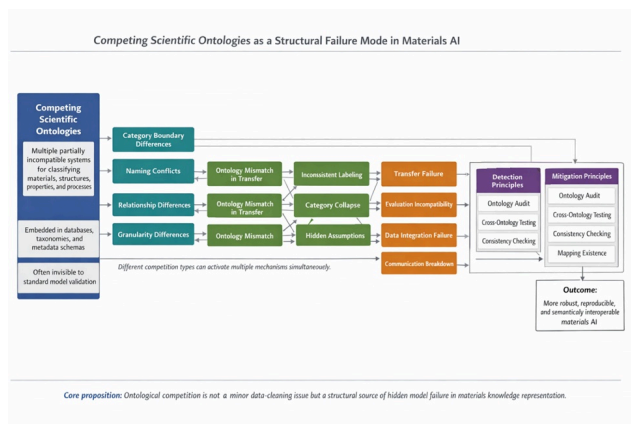


Figure 1. A hierarchical ontology-failure architecture showing how competing classification systems generate specific mechanisms of mismatch, produce observable failure modes in materials AI, and motivate corresponding detection and mitigation responses.

Defining the Ontology Problem

Definition 1: An ontology is a formal specification of concepts, categories, and relationships within a domain [1]. In the context of materials science, an ontology provides the explicit conceptual scaffolding that defines what counts as a distinct material class, how properties relate to structures, and which relationships are considered scientifically meaningful.

Definition 2: Ontological competition is the existence of multiple, partially incompatible ontologies for the same domain, leading to inconsistencies in knowledge representation [2]. When these ontologies are applied to the same underlying physical reality, they produce divergent categorizations that cannot be reconciled without loss of information or the introduction of arbitrary mapping rules.

The ontology problem must be carefully distinguished from related but distinct challenges in materials informatics. Data heterogeneity refers to differences in file formats, storage schemas, or numerical precision while preserving the same categorical framework; ontology competition, by contrast, concerns incompatible categorical frameworks themselves. Measurement error involves inaccuracies or noise in the recorded values of properties, whereas ontological competition concerns the very definitions of which properties are recorded and how they are grouped. A third

related issue—simple terminological synonymy—occurs when different labels are used for identical concepts within a shared ontology; ontological competition goes deeper, involving genuinely divergent conceptual boundaries.

In materials science, the ontology problem manifests with particular acuity because the domain spans multiple scales and disciplinary traditions. Crystal systems, space groups, and property definitions provide classic illustrations. One ontology might treat all materials with a tolerance factor above a certain threshold as perovskites, while another imposes stricter symmetry or compositional constraints [1, 10-17]. These differences are not superficial; they shape the very feature spaces that machine learning algorithms learn from. When models trained under one ontological regime encounter data annotated under another, the learned representations become misaligned at the categorical level, producing failures that cannot be attributed to data quality or model architecture alone. The problem is therefore epistemic rather than merely technical: it concerns the compatibility of the scientific worldviews embedded in the knowledge representations themselves [7].

Types of Ontological Competition

Table 1 consolidates the four forms of ontological competition by linking each to its underlying scientific assumption, its characteristic AI vulnerability, and the most appropriate diagnostic and mitigation priority.

Table 1. Analytical crosswalk of ontological competition types, encoded assumptions, AI vulnerabilities, and diagnostic priorities

Ontological competition type	What differs across ontologies	Embedded scientific assumption	Typical vulnerability introduced
Category boundary differences	Inclusion/exclusion criteria for a class (for example, what counts as a perovskite or high-entropy alloy)	Scientific categories are defined by different threshold rules, symmetry criteria, or	Learned boundaries become ontology-specific through physical generalization

		compositional boundaries	
Naming conflicts	The same concept receives different labels, or the same label refers to different concepts	Terminology is assumed to be semantically stable when it is not	Label collisions superlearn corrupted benchmarks
Relationship differences	Property-to-structure, property-to-process, or hierarchy relations are modeled differently	Causal or explanatory dependencies are encoded differently across communities	Model incompleteness related to structural affect prediction logic explained
Granularity differences	One ontology classifies at a coarse level, while another subdivides the same space in detail	The “right” level of abstraction differs by community or use case	Overgeneralization under labels spur distinct under labels

Type 1: Category boundary differences

Different ontologies draw category boundaries differently. In materials science, the question of what counts as a “perovskite” or a “high-entropy alloy” receives inconsistent answers across classification systems. One ontology may include distorted structures based on geometric tolerance factors, while another insists on ideal cubic symmetry. The consequence is that training datasets constructed under competing ontologies contain systematically different examples of the same nominal class, leading AI models to internalize boundary artifacts that do not generalize.

Type 2: Naming conflicts

The same concept may receive different names, or different concepts may share the same name. A “2D material” in one taxonomy might refer exclusively to van der Waals layered compounds, while in another it encompasses any atomically thin structure regardless of bonding type. Conversely, the label “metal-organic framework” might denote a narrow topological class in one ontology and a broader compositional family in another. These conflicts create label collisions that confuse supervised learning pipelines and render cross-database comparisons unreliable.

Type 3: Relationship differences

Ontologies specify different relationships between concepts. One system might treat thermal conductivity as a property directly subordinated to crystal structure, while another routes it through microstructure or defect density hierarchies. Such relational divergences alter the inferred causal graphs that machine learning models implicitly learn, producing predictions that encode incompatible assumptions about physical dependencies.

Type 4: Granularity differences

Different ontologies operate at different levels of classificatory detail. A coarse-grained ontology might group all oxides, while a fine-grained counterpart distinguishes subclasses based on oxidation state, coordination environment, or electronic character. When models trained at one level of granularity are applied to data annotated at another, they encounter either over-generalization or spurious distinctions that degrade performance.

Conceptually, ontology competition can be visualized as a set of overlapping but non-congruent circles in a Venn diagram. Each circle represents the categorical extent of a given ontology. Materials located in the intersection regions are assigned to different parent classes depending on which ontology is consulted. In contrast, materials in the non-overlapping segments exist in one classification but have no direct counterpart in others. The resulting diagram illustrates how partial overlap creates zones of ambiguity and incompatibility that no single ontology can resolve without external mapping rules [1, 2, 7].

Mechanisms of Ontology Failure

Ontology mismatch in transfer

A model trained on data organized under one ontology encounters systematic failure when applied to data structured according to another. The learned decision boundaries, feature importances, and latent representations become misaligned with the new categorical scheme, even when the underlying physical measurements are identical [3, 5].

Inconsistent labeling

The same material receives different categorical labels across datasets. Supervised models therefore learn from contradictory supervisory signals, producing unstable or averaged representations that capture neither ontology faithfully. This mechanism is especially pronounced when datasets are merged without ontological alignment [8, 9].

Category collapse

Models learn distinctions that exist in one ontology but are collapsed or absent in another. The result is either overfitting to ontology-specific artifacts or loss of predictive power when those distinctions prove irrelevant under the competing scheme [4, 6].

Hidden assumptions

Every ontology encodes community-specific assumptions about relevance, causality, and granularity. When these assumptions remain implicit, models inherit unstated priors that conflict with those embedded in alternative ontologies, manifesting as silent failures during deployment or reinterpretation [2, 7].

A Typology of Ontology Failure Modes

Transfer failure

A model trained under one ontological framework fails to generalize to data annotated under another. Mechanism: Direct mismatch between training and inference category structures. Materials example: A property-prediction model trained on perovskite data from the Materials Project ontology collapses when tested against NOMAD-classified structures that employ stricter symmetry criteria [5]. Detection signature: Abrupt performance degradation on

cross-ontology validation sets despite identical chemical compositions.

Evaluation incompatibility

Performance metrics cannot be meaningfully compared across studies that employ different ontologies.

Mechanism: Non-equivalent class definitions render accuracy, precision, or F1 scores incommensurable.

Materials example: Two studies reporting “success rates” for high-entropy alloy discovery cannot be reconciled if one counts compositional complexity while the other emphasizes configurational entropy [10-17]. Detection signature: Published results that appear contradictory despite ostensibly similar tasks.

Data integration failure

Combining datasets from competing ontologies produces logical contradictions or spurious correlations. Mechanism: Merging creates hybrid records that violate the internal consistency of either original ontology. Materials example: Integrating space-group data from OPTIMADE [9] with property annotations from PMD [8] may assign incompatible symmetry labels to the same entry. Detection signature: Integrity violations or unexpected statistical anomalies after data fusion.

Communication breakdown

Researchers cannot interpret or reproduce each other’s results because ontological commitments remain unstated. Mechanism: Implicit reliance on different category systems prevents shared understanding of reported findings.

Materials example: A claim about “enhanced stability in 2D materials” may refer to entirely different sets of compounds depending on the source ontology [1]. Detection signature: Persistent disagreement in the literature despite access to the same raw data.

Detection Principles

Detection of ontology-related failures in materials AI requires deliberate, proactive strategies rather than reliance on standard validation pipelines. Because ontological competition operates at the level of conceptual structure rather than numerical accuracy, conventional metrics such as mean absolute error or classification accuracy often remain insensitive to the underlying mismatch. The following principles provide a systematic framework for

surfacing these failures before they propagate into published results or deployed models.

Ontology audit

Every materials AI study must explicitly document the ontology or classification system used for data labeling, feature engineering, and model interpretation [1, 8]. An ontology audit involves listing the precise version of the taxonomy employed—whether the PMD core ontology [8], the OPTIMADE schema [9], or a custom extension of the materials data science ontology [6]—and mapping each category back to its definitional boundaries. This step reveals hidden divergences early; for example, a study claiming to model “perovskite stability” may appear internally consistent until the audit discloses that the training data followed a geometric tolerance-factor definition while the validation set followed a stricter space-group symmetry criterion drawn from NOMAD conventions [5]. Without such documentation, reviewers and downstream users cannot assess whether performance claims are ontology-bound or genuinely generalizable.

Cross-ontology testing

Models must be deliberately tested on datasets annotated under at least one competing ontology different from the training schema [2, 7]. This principle extends beyond ordinary cross-validation by requiring the creation or identification of parallel test sets that preserve identical physical measurements but re-label them according to an alternative categorical framework. In practice, a property-prediction network trained on high-entropy alloy data classified under a compositional-complexity ontology [17-29] should be evaluated against the same compounds re-annotated under a configurational-entropy taxonomy. Performance degradation that appears only under this cross-ontology regime signals ontological mismatch rather than overfitting or data scarcity.

Consistency checking

Automated or semi-automated checks must scan merged datasets for label inconsistencies that arise when records from different ontological sources are combined [6, 9]. Consistency checking algorithms flag entries where the same material identifier receives incompatible category assignments—such as a compound assigned both “metal-organic framework” and “coordination polymer” under competing naming conventions—or where relational assertions (for example, property-to-structure hierarchies)

contradict one another. The Materials Project and NOMAD repositories, when naively fused without prior alignment, frequently trigger these checks, exposing latent conflicts that would otherwise remain invisible within single-ontology silos [4, 5].

Mapping existence

Before any integration or transfer step, researchers must determine whether a formally specified mapping between the source and target ontologies exists and, if so, assess its completeness and fidelity [18-24]. If no mapping is available, or if the mapping covers only a subset of categories (as is common between legacy taxonomies and newer efforts such as MDS-Onto [6]), this absence itself constitutes a detectable risk factor. The principle forces explicit acknowledgment that unmapped regions of category space will produce irreducible ambiguity, thereby preventing the silent propagation of failure.

Collectively, these detection principles shift the burden from post-hoc debugging to preemptive ontological hygiene, ensuring that competing ontologies are not treated as background infrastructure but as active variables that must be monitored.

Mitigation Principles

Mitigation of ontological competition demands more than ad-hoc fixes; it requires embedding ontology awareness into the entire materials AI workflow. The principles below translate the conceptual analysis of failure modes into actionable design and governance strategies.

Ontology selection justification

Authors must provide a reasoned justification for the choice of any particular ontology, explicitly addressing why its category boundaries, naming conventions, and relational structure are appropriate for the scientific question at hand [1, 2]. Justification goes beyond convenience and includes discussion of trade-offs—such as the coarser granularity of one taxonomy versus the finer resolution of another—and acknowledgment of communities whose competing ontologies were deliberately set aside. This practice surfaces assumptions that would otherwise remain implicit and enables reviewers to evaluate whether the chosen ontology introduces unintended biases.

Ontology mapping

Where competing ontologies cannot be replaced by a single standard, explicit, machine-readable mappings must be constructed and published alongside models and datasets [7, 22]. These mappings function as translation layers that record, for every category in the source ontology, its nearest equivalent (or set of equivalents) in the target ontology, along with any residual ambiguity or information loss. In materials science, such mappings have already proven feasible between OPTIMADE and PMD schemas [8, 9]; extending them systematically would allow models to operate in a “mapped” rather than “native” mode, dramatically reducing transfer failures.

Ontology-agnostic representations

Model architectures should incorporate intermediate representations that are deliberately decoupled from any single ontological scheme—leveraging graph embeddings, unsupervised feature extractors, or physics-informed latent spaces that encode raw structural and compositional data without committing to predefined category labels [3, 4, 6]. By learning directly from atomic coordinates or spectral signatures rather than from ontology-derived features, these representations minimize the risk that downstream predictions inherit the quirks of a particular taxonomy. The Materials Data Science Ontology itself suggests pathways toward such agnostic embeddings by separating domain knowledge from data-science abstractions [6].

Community standards

The materials informatics community should actively support and participate in ongoing harmonization initiatives, treating ontology alignment as a shared infrastructure project rather than an individual research burden [1, 8, 9]. Participation includes contributing to the evolution of PMD, OPTIMADE, and related efforts, as well as adopting emerging top-level ontologies when they reach sufficient maturity. Community standards also encompass the creation of shared repositories of validated ontology mappings, thereby lowering the barrier for individual researchers to implement mitigation strategies.

Ontology versioning

Every trained model and published dataset must record the exact version of the ontology used, together with the version of any mapping layer applied, as part of standard provenance metadata [2, 5]. Ontology versioning parallels software versioning: changes in category definitions or boundary criteria are treated as breaking changes that

require re-evaluation of affected models. This principle ensures reproducibility and allows future researchers to trace performance shifts back to specific ontological updates rather than attributing them vaguely to “data drift.”

Taken together, these mitigation principles transform ontological competition from an intractable source of failure into a manageable dimension of model engineering and scientific practice.

Relation to Other Failure Modes

Ontology competition does not exist in isolation; it intersects with, amplifies, and sometimes masquerades as several other recognized failure modes in materials AI. Understanding these relations clarifies both the distinctiveness of the ontological problem and the ways it can be misdiagnosed.

First, ontological competition constitutes a specific subtype of domain shift. While conventional domain shift refers to changes in data distribution (for example, different temperature ranges or synthesis conditions), ontology-induced domain shift arises from changes in the categorical scaffolding itself [3, 4]. A model experiencing ontology mismatch encounters a shift not merely in feature statistics but in the very meaning of the labels that define the task, rendering standard domain-adaptation techniques insufficient unless they explicitly address categorical realignment.

Second, competing ontologies generate a particularly insidious form of epistemic debt [2]. The assumptions embedded in any chosen ontology—about which distinctions matter, which relationships are causal, and which boundaries are scientifically legitimate—are rarely articulated in model cards or methods sections. Over time, this unacknowledged debt accumulates, manifesting as silent failures when the model is transferred to new communities or new data infrastructures that operate under divergent ontological commitments. Unlike algorithmic bias or data leakage, ontological epistemic debt is structural rather than statistical and therefore resists detection by conventional fairness or robustness audits.

Third, ontologies form part of the scaffolding problem in scientific AI [5, 7]. Scientific infrastructure—databases, metadata schemas, and classification systems—is often

treated as a neutral background, yet it actively shapes what questions can be asked and what answers can be trusted. When multiple, competing scaffolds coexist, the scaffolding problem becomes acute: models built atop one scaffold may appear robust within their local context but collapse when the scaffold is swapped. Recognizing ontological competition, therefore, reframes the scaffolding problem from a general concern about data infrastructure to a specific, diagnosable failure mode rooted in incompatible knowledge representations.

Table 2 sharpens the paper’s theoretical contribution by distinguishing ontology-induced failure from adjacent problems such as data heterogeneity, measurement error, domain shift, leakage, and model mis-specification.

Table 2. Distinguishing ontology-induced failure modes from adjacent failure sources in materials AI

Failure source	Unit of mismatch	What is actually changing?	Why conventional validation may miss it
Ontology competition	Conceptual framework	Category definitions, naming logic, relational assumptions, or classificatory granularity	The Model may perform well inside its native ontology and fail only when labels or conceptual boundaries shift
Data heterogeneity	Data format or schema	File structures, database schema, metadata conventions, or units	Harmonization pipelines may normalize surface differences without revealing conceptual divergence
Measurement error	Observed value	Noise, instrument inaccuracy, missingness, or uncertainty	Aggregate metrics may absorb noise, obscuring whether the error stems

		in recorded properties	from labels of measurement
Conventional domain shift	Distribution	Training and deployment data differ in composition, operating conditions, or sampling regime	Validation may not simulate deployment conditions
Data leakage	Information boundary	Improper overlap between training and test information	Validation appears artificially strong rather than weak
Model mis-specification	Functional form	An architecture or hypothesis class cannot represent the true relationship	Poor performance may look similar to ontology failure when labels are taken for granted

By distinguishing ontology competition from these related modes while acknowledging their overlaps, the present analysis sharpens diagnostic precision: a performance drop that disappears after category realignment is more likely ontological than distributional, more epistemic than statistical, and more scaffold-dependent than architecture-dependent.

Implications for Materials AI Practice

The recognition of competing ontologies as a distinct failure mode carries concrete implications for how the field conducts research, reviews manuscripts, and builds communal resources.

For authors

(a) explicitly state the ontology used in every study, (b) discuss the limitations and boundary assumptions of that ontology, and (c) provide cross-ontology test results whenever claims of generality are made. These

requirements raise the evidentiary bar without stifling innovation; they simply ensure that claims about model robustness are not inadvertently ontology-bound [1, 6]. Authors who adopt this practice will produce papers whose results can be more confidently interpreted and extended by other communities.

For reviewers

(a) ask about ontology assumptions in every manuscript, (b) question results that depend on unstated ontology choices, and (c) require explicit mapping or justification when multiple data sources are merged. Reviewers thereby become active guardians of semantic interoperability rather than passive evaluators of numerical performance [2, 8].

For the community

(a) develop and maintain ontology harmonization efforts such as extensions to PMD and OPTIMADE, (b) create and curate open repositories of validated ontology mappings, and (c) establish ontology reporting standards analogous to the FAIR principles for data. These communal actions convert what is currently an individual burden into shared infrastructure, accelerating the transition toward genuinely interoperable materials AI [7, 9].

Adopting these practice changes does not eliminate ontological competition—nor should it, given the legitimate diversity of scientific perspectives—but it renders that competition visible, manageable, and productive.

Conclusion

Competing scientific ontologies constitute a fundamental yet under-recognized failure mode in materials knowledge representation. Different classification systems, naming conventions, and category boundaries generate incompatibilities that cause AI models to fail in ways invisible to standard validation protocols. This paper has defined the ontology problem, articulated its four primary types of competition, identified the mechanisms through which mismatch produces failure, and presented a typology of four specific ontology failure modes. It has further offered detection and mitigation principles grounded in the existing literature on materials ontologies and informatics infrastructures. By treating ontological competition as a structural rather than incidental challenge, the analysis calls for an ontology-aware paradigm in materials AI—one that acknowledges the coexistence of multiple, partially

incompatible knowledge representations and actively manages their incompatibilities through audit, mapping, agnostic representations, versioning, and community harmonization. Only by embedding such practices can the field move from fragmented, ontology-bound predictions toward reliable, reproducible, and semantically interoperable scientific knowledge. The future of materials AI depends not on the elimination of ontological diversity but on the disciplined recognition and navigation of its consequences.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 17 Jul 2024 Revised: 29 Aug 2024 Accepted: 07 Oct 2024
Published online: 18 January 2025

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Sinha PK, Gajbe SB, Debnath S, Sahoo S, Chakraborty K, Mahato SS. A review of data mining ontologies. *Data Technol Appl.* 2022;56(2):172-204.

Hagedorn TJ, Smith B, Krishnamurty S, Grosse I. Interoperability of disparate engineering domain ontologies using basic formal ontology. *J Eng Des.* 2019;30(10-12):625-54.

Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature.* 2018;559(7715):547-55.

Schmidt J, Marques MR, Botti S, Marques MA. Recent advances and applications of machine learning in solid-state materials science. *npj Comput Mater.* 2019;5(1):83.

Draxl C, Scheffler M. The NOMAD laboratory: From data sharing to artificial intelligence. *J Phys Mater.* 2019;2(3):036001.

Rajamohan BP, Bradley AC, Tran VD, Gordon JE, Caldwell HW, Mehdi R, et al. Materials data science ontology (MDS-

ONTO): Unifying domain knowledge in materials and applied data science. *Sci Data.* 2025;12(1):628.

Norouzi E, Waitelonis J, Sack H. The landscape of ontologies in materials science and engineering: A survey and evaluation. *arXiv preprint arXiv:2408.06034.* 2024 Aug 12.

Bayerlein B, Schilling M, Birkholz H, Jung M, Waitelonis J, Mädler L, et al. PMD core ontology: Achieving semantic interoperability in materials science. *Mater Des.* 2024;237:112603.

Andersen CW, Armiento R, Blokhin E, Conduit GJ, Dwaraknath S, Evans ML, et al. OPTIMADE, an API for exchanging materials data. *Sci Data.* 2021;8(1):217.

Zhao X, Greenberg J, McClellan S, Hu YJ, Lopez S, Saikin SK, et al. Knowledge graph-empowered materials discovery. In: 2021 IEEE International Conference on Big Data. Piscataway, NJ: IEEE; 2021:4628-32.

Del Nostro P, Friis J, Ghedini E, Goldbeck G, Toti D, Zaccarini FA. Top level ontologies: Desirable characteristics in the context of materials science. In: CEUR Workshop Proceedings. Aachen: CEUR-WS; 2024:29-40.

Zhang Y, Chen F, Liu Z, Ju Y, Cui D, Zhu J, et al. A materials terminology knowledge graph automatically constructed from text corpus. *Sci Data*. 2024;11(1):600.

Takahashi L, Takahashi K. Visualizing scientists' cognitive representation of materials data through the application of ontology. *J Phys Chem Lett*. 2019;10(23):7482-91.

Otte JN, Kiritsi D, Ali MM, Yang R, Zhang B, Rudnicki R, et al. An ontological approach to representing the product life cycle. *Appl Ontol*. 2019;14(2):179-97.

Solórzano-Requejo W, Vega CA, Martínez RZ, Bodaghi M, Lantada AD. Ontology for smart 4D printed material systems and structures synergically applied with generative artificial intelligence for creativity promotion. *Smart Mater Struct*. 2025;34(1):015045.

Boyce B, Dingreville R, Desai S, Walker E, Shilt T, Bassett KL, et al. Machine learning for materials science: Barriers to broader adoption. *Matter*. 2023;6(5):1320-3.

Lantada AD, Korvink JG, Islam M. Taxonomy for engineered living materials. *Cell Rep Phys Sci*. 2022;3(4):100807.

Nop G, Mundy M, Smith JD, Paudyal D. Faithful novel machine learning for predicting quantum properties. *npj Comput Mater*. 2025;11(1):244.

Huang G, Guo Y, Chen Y, Nie Z. Application of machine learning in material synthesis and property prediction. *Materials*. 2023;16(17):5977.

Greenberg J, Polson K, McClellan S, Zhao X, Kalinowski A, An Y. Enhancing semantic interoperability across materials science with HIVE4MAT. *arXiv preprint arXiv:2411.00676*. 2024 Nov 1.

McClellan S, An Y, Zhao X, Lin X, Greenberg J. Characterizing semantic ambiguity of the materials science ontologies. *arXiv*

preprint *arXiv:2310.00078*. 2023 Sep 29.

Lambrix P, Armiento R, Li H, Hartig O, Abd Nikooie Pour M, Li Y. The materials design ontology. *Semant Web*. 2024;15(2):481-515.

Voigt SP, Kalidindi SR. Materials graph ontology. *Mater Lett*. 2021;295:129836.

De Baas A, Del Nostro P, Friis J, Ghedini E, Goldbeck G, Paponetti IM, et al. Review and alignment of domain-level ontologies for materials science. *IEEE Access*. 2023;11:120372-401.

Liu T, Barnard AS. The emergent role of explainable artificial intelligence in the materials sciences. *Cell Rep Phys Sci*. 2023;4(10):101630.

Broeckmann C, Bührig-Polaczek A, Hallstedt B, Krupp U, Rajaei A, Rom M, et al. Materials within a digitalized production environment. In: *Internet of Production: Fundamentals, methods and applications*. Cham: Springer International Publishing; 2023. pp. 139-52.

Lushnei S, Shumskyi D, Shykula S, Jimenez-Ruiz E, Garcez AD. Large language models as oracles for ontology alignment. In: *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. Stroudsburg, PA: Association for Computational Linguistics; 2026. pp. 2435-49.

Kolyvakis P. Approaching ontology alignment through representation learning to bridge the semantic gap in engineering applications [dissertation]. Lausanne: EPFL.

Khan H, Saqib M, Khattak HA, Ali SI, Lee S. Ontology alignment for accurate ontology matching: A survey. In: *International Conference on Smart Homes and Health Telematics*. Cham: Springer Nature Switzerland; 2023. pp. 338-49.