Top level ontologies: desirable characteristics in the context of materials science

Pierluigi Del Nostro[1], Jesper Friis[2], Emanuele Ghedini[3], Gerhard Goldbeck[1], Daniele Toti[1,4] and Francesco Antonio Zaccarini[3]

[1]Goldbeck Consulting Limited, Cambridge, United Kingdom
[2]SINTEF AS, Trondheim, Norway
[3]Alma Mater Studiorum - University of Bologna, Bologna, Italy
[4]Catholic University of the Sacred Heart, Faculty of Mathematical, Physical and Natural Sciences, Brescia, Italy

Abstract
The desiderata for the effective representation of Materials Science and Engineering (MSE) knowledge in Top Level Ontologies (TLOs) are discussed, based on the empirically grounded assumption that different ontologies exhibit different degrees of suitability in different contexts, and with respect to different goals and use-cases. The discussion follows the general requirements for TLOs outlined in ISO/IEC 21838-1, investigating each of them in the context of MSE 's methodological principles and procedural staples. As a result of the analysis, a set of desirable characteristics for TLOs is individuated, providing reasons to favor certain ontology design alternatives. The Elementary Multiperspective Material Ontology (EMMO) is briefly introduced as an example of an ontology engineered to meet the MSE desiderata. While comparing the effectiveness of conceptual frameworks across different contexts remains challenging, the analysis should lead to improvements in knowledge representation for the MSE domain, either directly, or by fostering explicit discussions regarding ontology design choices.

Keywords
Knowledge Representation, Materials Science and Engineering, Science and Industry, Top Level Ontology

1. Introduction

In recent years, the importance of *semantic technologies* has become increasingly evident. These technologies serve as the foundation for efficient digitalization, data sharing and data exploitation, which are essential for driving innovation in industrial contexts [1]. As a result, the development and adoption of semantic technologies have become a priority for stakeholders [2]. National and international actors are actively promoting their use, particularly emphasizing *computational ontologies* [3], due to their crucial role in establishing interoperability.

SeMatS 2024: The 1st International Workshop on Semantic Materials Science co-located with the 20th International Conference on Semantic Systems (SEMANTiCS), September 17-19, Amsterdam, The Netherlands.

✉ pierluigi@goldbeck-consulting.com (P. D. Nostro); jesper.friis@sintef.no (J. Friis); emanuele.ghedini@unibo.it (E. Ghedini); gerhard@goldbeck-consulting.com (G. Goldbeck); daniele.toti@unicatt.it (D. Toti); francesco.zaccarini@unibo.it (F. A. Zaccarini)
🌐 https://materialsmodelling.com (P. D. Nostro); https://www.sintef.no (J. Friis); https://www.unibo.it (E. Ghedini); https://materialsmodelling.com (G. Goldbeck); https://www.unicatt.it (D. Toti); https://www.unibo.it (F. A. Zaccarini)
iD 0000-0002-5174-8508 (P. D. Nostro); 0000-0002-1560-809X (J. Friis); 0000-0003-3805-8761 (E. Ghedini); 0000-0002-4181-2852 (G. Goldbeck); 0000-0002-9668-6961 (D. Toti); 0009-0008-8009-5009 (F. A. Zaccarini)

Computational ontologies can be seen as models of the relevant entities of a system (and the relations among them), or as providing the formal systematization of a domain of interest by focusing on specific concepts and definitions [4, 5]. They are usually classified depending on the generality –understood in terms of domain-invariance – of the concepts they revolve around (Top Level, Middle Level, Domain Level, Application Level), even though the classification is coarse-grained and largely indicative. Alternatively, it is possible to distinguish them based on the richness of the underlying formalization, ranging from Foundational Ontologies to Light-Weight Ontologies) [6, 7, 8]. Domain and Application Level Ontologies (DLOs and ALOs, respectively, in what follows) tend to be light-weight, leaning towards technologies, such as triplestores, that prioritize handling large volumes of instance data. Top- and Middle-Level Ontologies (TLOs and MLOs, respectively, henceforth) are instead usually axiomatized in expressive formal languages, such as FOL and OWL-DL [9]. These ontologies should be an expression of a particular worldview, and are meant to provide a general framework for knowledge representation; this should greatly reduce the possibility of mistakes in conceptualization and ground effective interoperability across multiple domains and actors. TLOs are thus pivotal in industrial environments in which interoperability has to be established across entire value chains, and where pluralistic federated distributed systems seem to be inescapable.

In fields like Materials Science and Engineering (MSE), characterized by rapid development and innovation cycles, and which need content-oriented representations of systems and processes rather than a simple semantic architecture for software integration, TLOs are arguably indispensable [10, 11]. Machine learning has recently raised awareness of the importance of content-oriented representations, which play a major role in ontologies' varying degrees of suitability with respect to different domains and specific pragmatic goals and use cases [12]. It is therefore pivotal to investigate the pros and cons of different options with respect to the MSE domain, especially for TLOs. However, despite the critical need, the existing literature only offers general principles, and no consensus has been reached on evaluation metrics, largely undermining the possibility of a comparative empirical study [13], which are relatively more approachable for DLOs [14, 15]. The present study aims to address this gap by fostering explicit discussions around ontology design choices, individuating desirable characteristics *specifically for the MSE domain.* While the proposed guidelines have the potential to enhance ontology design on their own, this work will hopefully also lead to more deliberate engineering choices and, derivatively, to improvements in knowledge representation for MSE across the board.

The discussion is organized as follows. In Section 2, the requirements for TLOs outlined in ISO/IEC 21838-1 are taken as a starting point to investigate all the aspects relevant for the MSE domain. Specifically, Section 2.1 examines the pros and cons of different options in view of conceptual coverage requirements, Section 2.2 addresses aspects related to domain neutrality, and Section 2.3 provides a brief overview of implementation aspects. The paper analyzes and deliberates on traditional, conflicting design choices by referring to established scientific methodologies and procedural norms. Based on this analysis, a list of core desirable characteristics for TLOs (relative to the MSE domain) is proposed. The Elementary Multiperspective Material Ontology (EMMO)'s adherence to the supported principles is shown in Section 3, and the results are summarized in Section 4. While alternative conclusions might be drawn from the analysis, it is the authors' belief that the ensuing discussion will significantly enhance knowledge representation in MSE, leading to improvements in the related industrial sectors.

2. Requirements and desirable characteristics for TLOs

ISO/IEC 21838-1[1] is a standard detailing a list of formal, conceptual and bureaucratic requirements for TLOs. In the following discussion, the focus will firstly be on representational requirements (Section 2.1): while these requirements single out topics that ought to be covered by TLOs' conceptual schema, they do not prescribe specific solutions, leaving room for different options with contextually salient pros and cons. The analysis will naturally lead to the discussion of domain neutrality (Section 2.2), briefly touching on matters pertaining to cross-domain interoperability. A few notes on implementation requirements (Section 2.3) will conclude the investigation. The core desiderata are collectively reported in Section 4, which also sums up the proposed guidelines for MSE-focused TLOs.

2.1. Adapting general conceptual requirements to MSE

ISO/IEC 21838-1 outlines the representational requirements for TLOs. These include **(1)** *Space and Time*, **(2)** *Classification*, **(3)** *Actuality and Possibility*, **(4)** *Time and Change*, and **(5)** *Causality*. TLOs should also handle **(6)** *Parthood*, **(7)** *Location*, **(8)** *Constitution*, **(9)** *Scale and Granularity*, **(10)** *Qualities*, **(11)** *Quantities*, **(12)** *Mathematical*, **(13)** *Informational*, **(14)** *Social*, **(15)** *Mental Entities*, and **(16)** *Events and Processes*. Notably, the standard provides a list of topics to be addressed, without prescribing the adoption of a given approach or imposing the endorsement of any specific ontological commitments. ISO-certified ontologies can endorse different approaches (*e.g.*, BFO [8] and DOLCE's [6] different stance on parthood and co-location) or strong stances (*e.g.*, BFO's original rejection of possibilia, ISO 15926's adoption of quadridimensionalism across the board [16]) on specific topics; as such, the standard leaves room for different options, warranting an analysis of their adequacy with respect to specific domains.

Starting from (1) Space and Time, MSE applications often require the representation of systems at high levels of detail, grounded in scientific rigor. Representations should be based on state-of-the-art scientific theories, ensuring compatibility with relativity. A unitary treatment of space and time, often referred to as 4-dimensional or 4D, aligns better with current scientific theories and avoids issues related to simultaneity and ternary relations in Web Languages [17, 18, 19, 16]. Furthermore, separating worldly constraints from observational claims is crucial, as measurements are prone to errors and uncertainty. This separation allows representing inconsistent attributions relative to specific observations [20], adhering to FAIR principles in metadata collection and documentation [21].

This approach can be extended to the representation of (10) Qualities and (11) Quantities. Providing means to represent the subjectivity of observation should be a priority for any TLO intended for scientific and industrial contexts [22]. Tracking the multiple ways a quantity or quality can be determined, including through observation and modeling, is pivotal in MSE (see ISO 10303-235:2019). It is highly recommended that a property be identified with a specific measuring or calculating procedure, in accordance with ISO 10303-45:2019, as this information is crucial in application scenarios [23, 24, 25, 26, 27]. Discordance can often be accounted for by the use of different measuring or calculating techniques or instruments. Grounding qualities

---

[1] See https://www.iso.org/standard/71954.html. Publicly available at https://standards.iso.org/ittf/PubliclyAvailableStandards/index.html.

and quantities in specific measurements or calculation procedures helps establish clear identity criteria and aligns with best practices in metrology [28].

Clear, empirically-grounded identity criteria [29, 30] are essential for TLOs in scientific contexts. Following [31], it is possible to distinguish between *extensional* and *intensional* identity criteria. Extensional criteria, based on parts and relations, are generally clearer and easier to establish. Intensional criteria, based on nature or essence, are more difficult to establish unambiguously but allow for more informative distinctions. Specifically, adhering to a scientific worldview, no intrinsic intensional criteria –independent of observation processes and theoretical assumptions– should be attributed to posited entities, emphasizing the importance of compliance with principles of metrology. Given the importance of precision and lack of ambiguity in MSE, stringent extensional criteria should be endorsed, with intensional criteria playing a complementary role within the imposed constraints. While this might appear too stringent a limitation given TLOs' expressive needs, empirically-grounded distinctions allow to largely recover scientific taxonomies following the scientific methodology. This choice is also motivated by the increasing importance of AI agents and sensors for MSE, and industrial/scientific contexts in general [32, 33, 34], since introducing identity criteria based on aspects that are not (directly) empirical could undermine the full interoperability between humans and machines.

Regarding (2) Classifications, (12) Mathematical, (13) Informational, (14) Social, and (15) Mental Entities, it thus seems advisable to limit, and heavily constrain, if not outright avoid, the inclusion of abstract entities in MSE-oriented TLOs. Expressiveness can be recovered through other means, or abstract entities can be reduced to, or otherwise grounded in, non-abstract ones. As shown by the cases cited in the opening of the section, all the aforementioned options can equally satisfy the breadth of coverage requirements.

When it comes to (2) Classifications, whose identity criteria are historically dubious [35], only resources for representing distributive predicates seem useful in MSE, and this can be achieved without using outright sets or classes [36]. (12) Mathematical and (13) Informational entities are salient in MSE: the former are essential for quantitative representation, and the latter for modeling [37]. It is crucial not to confuse a model with what is modeled and to understand data as carrying intrinsic meaning only through interpretation [38]. Reducing these entities to interpreted symbols or physical substrata, while separating syntactic and semantic content, can enhance clarity, although this should be evaluated case by case. (14) Social and (15) Mental Entities are less relevant in MSE itself but can be relevant for the knowledge value chain. A modest stance on commitments should be taken unless strictly necessary.

The maxim on identity criteria should also guide engineering choices regarding (6) Parthood, (7) Location, (8) Constitution, (4) Time and Change, and (16) Events and Processes. For (6) Parthood, adopting extensional mereology (stronger than Classical Extensional Mereology [39]) allows distinguishing entities by their parts, providing simple identity criteria for material entities. Considering boundaries as parts of entities, distinguishing between interior parts and entities can be practical in MSE, though it is not a standard position in mereology [40, 41].

Avoiding a distinction between entities and their location through supersubstantivalism or a relational theory of space-time offers a clean approach to (7) Location, aligning with empiricism and avoiding duplication of entities and parthood relations [42]. This is especially beneficial in contexts where all entities can be located in the environment, such as smart factories and Digital Twins.

A non-tensed representation of time and a unitary approach to location support a strict stance on (4) Time and Change, adopting a broad perdurantist approach to persistence [18]. Specifically, stage-perdurantism, viewing entities as sequences of stages with defined characteristics, avoids ambiguities in property attributions. While endurantism is often assumed in natural language for objects, the difference between the two may not substantially affect expressiveness [17]. However, distinguishing between objects and (16) Events and Processes in ontology challenges the application of extensional criteria across the board. Moreover, in some specific contexts, it is necessary to represent the same entity both as an object and as a process – for instance, in life-cycle management there is a need to track the continuity in product-development and material/processes inter-dependencies [43, 44]. Although it is possible to do so while endorsing the distinction as substantial, this weighs on the cost-benefit analysis. Likewise, for (8) Constitution, property attributions can recover expressiveness without duplicating entities, especially in conjunction with perdurantism [45, 18]. Although there are trade-offs between options, this approach need not result in expressive limitations.

For (3) Actuality and Possibility [46], no specific option appears to be clearly superior for MSE. Possibilities can be represented as actual or by exploiting disjunctive constraints in the terminological box; as such, possible world semantics or encoding should be introduced only if they are necessary to enhance the expressiveness and clarity of the representational framework [47, 48]. If discourse on possibilities is allowed, restricting the modal space to the realm of the scientifically possible, based on the best scientific theories, is advisable. Interventionist or manipulation accounts seem especially suitable for MSE [49].

(5) Causality and (9) Scale and Granularity are particularly salient for MSE. Causality is central to scientific knowledge and industrial workflows, making ingrained support for causal discourse essential. While causal talk in science has been debated [50, 51], it remains ubiquitous and crucial [52, 53, 54]. TLOs should support causal discourse either through relations or by facilitating the introduction of axiomatic constraints with an intended causal or law-based interpretation. Both counterfactual [55] and productive [56] notions of causality are used in sciences, but productive notions, being more robust and physically interpretable, are preferable for MSE. Given the rise of Causal AI, productive notions might be more beneficial for distinguishing actual causal relations from correlations.

Representing entities across different (9) Scales and Granularities is vital for MSE, as innovation often involves analyzing materials across various disciplines, from particle physics to chemistry. A reductionist approach for specific systems avoids ambiguities and sets clear identity criteria. However, scientific pluralism and issues like multiple realizability should be acknowledged and accounted for [57, 58, 59].

2.2. Domain-neutrality and usability

Given the discussion up until this point, it appears that there are limits to the domain neutrality of a TLO's concepts. This is to be expected given that TLOs are the expression of (different) worldviews. For instance, even if *e.g.*, "causality" might be a ubiquitous term, different variants are more or less adequate in given contexts. Problematically, many of the alternatives outlined in Section 2.1 are mutually incompatible and competing, leaving no neutral ground to retreat to for the sake of generality. Such observations can raise doubts about the efficacy of TLOs in fostering interoperability across different domains – or value-chains, for the MSE domain. It's important to note, however, that the broad applicability of TLOs is generally not compromised by specific design decisions, which only impact their suitability in particular situations. Nevertheless, these considerations are crucial in ontology engineering, especially when developing domain-specific hubs. It is worth adding that domain neutrality may be outright detrimental, insofar as adopting concepts shared by the communities employing the ontologies is pivotal to prevent misuse.

Indeed, one of the core impediments to the effective exploitation of TLOs concerns the accessibility of their terminology.[2] While ontologies should be based on shared conceptualizations [60], stakeholders are often required to take up what is, to all effects, a foreign jargon to exploit the semantic artifacts. To be accessible to MSE practitioners, an ontology should use concepts in line with state-of-the-art science, related to recognized gold standards for the scientific community. This can prove challenging due to the idiosyncrasies among scientific disciplines, which are reflected at the level of notions and underlying concepts (often leading to incompatible standards).

2.3. Implementation aspects

ISO/IEC 21838-1 also provides guidelines concerning axiomatizations in expressive formal languages (*e.g.*, First Order Logic, or Common Logic) and machine-readable implementations in OWL 2. While neuro-symbolic AI may soon go ways towards addressing, at least partially, the tension between expressiveness and computational efficiency [61, 62], the latter is currently a core limitation of semantic technologies [63]. Since MSE is characterized by large datasets, as well as fast innovation and development cycles, lightweight versions compatible with less expressive OWL profiles (*e.g.*, OWL 2 EL/RL/QL) ought to be supported, and an efficient, modular architecture is paramount. However, a rigorous conceptualization in expressive formal languages is pivotal in order to avoid conceptual mistakes. While this topic is beyond the scope of this paper, as well as of ISO/IEC 21838-1, this also has consequences related to versioning and maintenance, given the ultimate aim of ontologies of grounding interoperability. Related points might be explored in more detail in a future publication.

3. EMMO's position & MSE desiderata

The Elementary Multiperspective Material Ontology was engineered by taking into account the points discussed in the previous sections. The ontology has been developed by MSE practitioners, in close collaboration with analytical philosophers, within a number of European projects under the umbrella of the European Materials Modeling Council (EMMC)[3]. The innovative features of EMMO, compared to standard taxonomies and other ontologies (both foundational and lightweight), include three main aspects: (i) the influence of natural sciences in its framework, (ii) its unique architecture with a common core and multiple modular *perspectives*, and (iii) its pragmatic stance concerning commitments. These features enhance its usability, formal robustness, and expressive capacity in knowledge representation. Regarding the specific points

---

[2] See for instance https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5035160b3&appId=PPGMS.

[3] https://emmc.eu/.

discussed in the previous sections, EMMO's traits are briefly listed in the following points; a more detailed exposition of EMMO is beyond the scope of this paper.

- Representational requirements:
  - (1) *Space and Time.* EMMO adopts a causal relational theory of spacetime, consistent with special relativity [64]. Spatiotemporal constraints and measurements are distinct.
  - (10) *Qualities* and (11) *Quantities.* Instead of committing to universals or tropes, EMMO endorses a semiotic approach, based on Peirce's [65], where symbols are connected by an interpreter to an object. Observations are distinguished from other forms of attributions and grounded in causal processes, and syntactic and semantic aspects are kept distinct.
  - (2) *Classifications.* EMMO avoids commitments to sets or classes, recovering the expressiveness via collections (disconnected entities) and semiosis.
  - (12) *Mathematical* and (13) *Informational Entities.* EMMO refers to Floridi's work to deal with data and information as distinctions that make a difference [66]. Mathematical entities are likewise understood structurally.
  - (14) *Social* and (15) *Mental Entities.* EMMO takes distinction based on these aspects as non-substantial, yet allows the attribution of social and mental roles through semiotic processes.
  - (6) *Parthood.* EMMO endorses AGEM, an extensional theory stronger than CEM [67].
  - (7) *Location* and (8) *Constitution.* Endorsing a relational theory of time, and extensional conditions across the board, these relations are not needed in EMMO. The related expressiveness is recovered through perspectives.
  - (4) *Time and Change*, and (16) *Events and Processes.* EMMO endorses a form of perdurantism with respect to change, allowing for a unitary representation of objects and processes. The relevant distinction is understood as non-substantial, and dealt with in a specific perspective, to facilitate common sense knowledge representation, and thus interoperability with other ontologies. A unitary representation of objects and processes is highly relevant to materials, as they are understood to be inherently 'process dependent'.
  - (3) *Actuality and Possibility.* EMMO does not specifically commit to possible worlds or similar ingrained machinery to represent possibility.
  - (5) *Causality.* EMMO's uppermost module is based on a mereocausal theory aligned with productive Conserved Quantities theories of causation and formalizing Feynman's Diagrams [52, 68, 69, 70, 71].
  - (9) *Scale and Granularity.* Two of EMMO's perspectives offer tools to represent entities at different granularities and functional roles in systems and processes. EMMO's reductionistic stance is in line with [72, 59].
  - *Identity Criteria.* All the core identity criteria are extensional, being based on parthood, causal relations (and thus, indirectly, location) as well as properties, yet not property attributions. Intensional criteria are introduced in perspectives to recover expressiveness, yet are constrained by the former and not considered ontologically substantial, allowing for pluralism.

- Usability:
  - EMMO retains general applicability by distributing axiomatic constraints at different levels (*e.g.,* introducing more specific constraints regarding Causality relative to certain kinds, complementary to general principles grounding the representation of spatiotemporal relations).
  - EMMO's core categories are either formal or empirically grounded.
  - EMMO's concepts are connected with gold standards such as the Standard Model, and validated by practitioners.
- Implementation aspects:
  - EMMO is being formalized in First-Order Logic to reduce the risk of errors in conceptualization.
  - The machine-readable version of EMMO, in OWL 2 DL, is triplestore-friendly, including no implicit constraints that are not supported by one of the supported profiles (RL/EL, QL). The same artifact can be employed using profile-specific reasoners to allow direct exploitation by DLOs and ALOs.
  - EMMO is inherently modular, given its multiperspective architecture, exhibiting distinctions based on scientific domains at the lower levels.

4. Core uptakes and conclusive remarks

This work has provided an analysis of the desirable characteristics of TLOs for the representation of MSE knowledge, critical given the varying degrees of suitability of different ontologies to distinct domains, pragmatic goals and use cases. Given the lack of consensus on evaluation metrics in the literature, the investigation has proceeded by referring to methodological tenets of the discipline, taking ISO/IEC 21838-1's requirements as a starting point to ensure comprehensiveness. In summary, the following characteristics have been individuated as pivotal:

- Endorsement of a scientific worldview, grounded in state-of-the-art sciences:
  - focus on empirically-grounded categories, following the scientific methodology.
  - adoption of a terminology accessible by MSE stakeholders, with explicit connections to the gold standards of sciences.
- Use of extensional identity criteria to avoid ambiguities.
- Ingrained support for highly detailed qualitative and quantitative representation.
- Adoption of a stance for multi-scale representation, compatible with scientific pluralism.
- Focus on scientific laws, causal processes, and industrial workflows.
- Endorsement of a stance grounding properties in observational processes or modeling techniques, and allowing and tracking relativization of attributions.
- Endorsement of a stance clearly distinguishing systems and models.
- Lightweight implementations support MSE's fast innovation and development cycles.

Notably, EMMO complies with all the aforementioned points. The proposed guidelines are grounded in an explicit analysis and rationale and are meant to foster a productive discussion, leading to improvements in the knowledge representation and the engineering of TLOs for the MSE domain, thereby benefiting the related industrial sectors.

Acknowledgments

This work received funding from the European Commission via the Horizon 2020 projects NanoMECommons (G.A. n. 952869), OntoTrans (G.A. n. 862136) and OpenModel (G.A. n. 953167). We'd like to thank an anonymous reviewer for their helpful comments.

References

[1] L. Himanen, et al., Data-Driven Materials Science: Status, Challenges, and Perspectives, Advanced Science 6 (2019) 1900808. doi:10.1002/advs.201900808.

[2] Z. Feng, Knowledge Graphs, Springer Nature Singapore, Singapore, 2023, pp. 753–785. doi:10.1007/978-981-16-5172-4_12.

[3] S. Staab, R. Studer, Handbook on ontologies, Springer Science & Business Media, 2013.

[4] W3C, Building the Web of Data, 2013. URL: https://www.w3.org/2013/data/.

[5] N. Guarino, D. Oberle, S. Staab, What is an Ontology, Springer, Berlin Heidelberg, 2009.

[6] C. Masolo, et al., WonderWeb Deliverable D18 Ontology Library, Technical Report, IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web, 2003.

[7] S. Borgo, L. Lesmo, The Attractiveness of Foundation Ontologies in Industries, Frontiers in Artificial Intelligence and Applications 174 (2015) 1—-9.

[8] R. Arp, B. Smith, A. D. Spear, 27Kinds of Ontologies and the Role of Taxonomies, in: Building Ontologies With Basic Formal Ontology, The MIT Press, 2015. URL: https://doi.org/10.7551/mitpress/9780262527811.003.0002. doi:10.7551/mitpress/9780262527811.003.0002.

[9] W3C, Web Ontology Language., https://www.w3.org/OWL/, 2012.

[10] B. Bayerlein, et al., A Perspective on Digital Knowledge Representation in Materials Science and Engineering, Advanced Engineering Materials 24 (2022) 2101176. doi:10.1002/adem.202101176.

[11] A. Agrawal, A. Choudhary, Perspective: Materials informatics and big data: Realization of the "fourth paradigm" of science in materials science, Apl Materials 4 (2016).

[12] N. Adamovic, OntoCommons Roadmap V1, 2023. doi:10.5281/zenodo.7544509.

[13] S. Borgo, et al., Foundational ontologies in action: Understanding foundational ontology through examples, Applied Ontology 17 (2022) 1–16. doi:10.3233/AO-220265.

[14] X. Zhang, C. Zhao, X. Wang, A survey on knowledge representation in materials science and engineering: An ontological perspective, Computers in Industry 73 (2015) 8–22. doi:10.1016/j.compind.2015.07.005.

[15] A. De Baas, P. D. Nostro, J. Friis, E. Ghedini, G. Goldbeck, I. M. Paponetti, A. Pozzi, A. Sarkar, L. Yang, F. A. Zaccarini, D. Toti, Review and Alignment of Domain-Level Ontologies for Materials Science, IEEE Access 11 (2023) 120372–120401. doi:10.1109/ACCESS.2023.3327725.

[16] D. Leal, Iso 15926 "life cycle data for process plant": an overview, Oil & Gas Science and Technology-revue De L Institut Francais Du Petrole - OIL GAS SCI TECHNOL 60 (2005) 629–637. doi:10.2516/ogst:2005045.

[17] E. Hirsch, Quantifier Variance and Realism: Essays in Metaontology, Oxford University Press, 2011. doi:10.1093/acprof:osobl/9780199732111.001.0001.

[18] T. Sider, Four Dimensionalism: An Ontology of Persistence and Time, Oxford University Press, 2001.

[19] M. Giunti, G. Sergioli, G. Vivanet, S. Pinna, Representing n-ary relations in the Semantic Web, Logic Journal of the IGPL 29 (2019) 697–717. doi:10.1093/jigpal/jzz047.

[20] J. Bogen, J. Woodward, Saving the phenomena, The Philosophical Review 97 (1988) 303–352. URL: http://www.jstor.org/stable/2185445.

[21] M. D. Wilkinson, et al., The FAIR Guiding Principles for scientific data management and stewardship, Scientific data 3 (2016). doi:10.1038/sdata.2016.18.

[22] C. Masolo, A. Botti Benevides, D. Porello, The interplay between models and observations, Appl. Ontol. 13 (2018) 41–71. URL: https://doi.org/10.3233/AO-180193. doi:10.3233/AO-180193.

[23] A. McMillan, et al., A review of composite product data interoperability and product life-cycle management challenges in the composites industry, Advanced Manufacturing: Polymer and Composites Science 3 (2017) 130–147. doi:10.1080/20550340.2017.1389047.

[24] N. Swindells, Materials Property Data in STEP, Springer Berlin Heidelberg, Berlin, Heidelberg, 1996, pp. 151–160. doi:10.1007/978-3-642-80286-7_17.

[25] C. Masolo, Founding properties on measurement, in: Formal Ontology in Information Systems, 2010, pp. 89–102. URL: https://api.semanticscholar.org/CorpusID:3230503.

[26] A. Frigerio, A. Giordani, L. Mari, Outline of a general model of measurement, Synthese 175 (2010) 123–149. URL: http://www.jstor.org/stable/40801334.

[27] F. Probst, Observations, measurements and semantic reference spaces, Appl. Ontol. 3 (2008) 63–89.

[28] Joint Committee for Guides in Metrology, International Vocabulary of Metrology, JCGM, 2021. URL: https://www.bipm.org/documents/20126/54295284/VIM4_CD_210111c.pdf/a57419b7-790f-2cca-f7c9-25d54d049bf6.

[29] J. R. M. Nicola, G. Guizzardi, Individual determinacy and identity criteria in ontology-driven information systems, in: Formal Ontology in Information Systems, 2018. URL: https://api.semanticscholar.org/CorpusID:52274494.

[30] N. Guarino, The role of identity conditions in ontology design, in: Proceedings of the International Conference on Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science, COSIT '99, Springer-Verlag, Berlin, Heidelberg, 1999, p. 221–234.

[31] C. Partridge, A. Mitchell, M. West, A survey of Top-Level Ontologies: To inform the ontological choices for a Foundation Data Model, 2020. doi:10.17863/CAM.58311.

[32] J. Schmidt, M. R. Marques, S. Botti, M. A. Marques, Recent advances and applications of machine learning in solid-state materials science, npj computational materials 5 (2019) 83.

[33] M. Glauer, F. Neuhaus, S. Flügel, M. Wosny, T. Mossakowski, A. Memariani, J. Schwerdt, J. Hastings, Chebifier: automating semantic classification in ChEBI to accelerate data-driven discovery, Digital Discovery 2024 (2024). doi:10.1039/D3DD00238A.

[34] C. Partridge, A. Mitchell, M. Loneragan, H. Atkinson, S. de Cesare, M. Khan, Coordinate systems: Level ascending ontological options, in: 2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-

C), 2019, pp. 78–87. doi:10.1109/MODELS-C.2019.00017.

[35] W. V. Quine, Philosophy of Logic (2Nd Edition), Harvard University Press, Cambridge, MA, 1986.

[36] O. Linnebo, Plural Quantification, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Spring 2022 ed., Metaphysics Research Lab, Stanford University, 2022.

[37] E. M. Sanfilippo, Ontologies for information entities: State of the art and open challenges, Appl. Ontol. 16 (2021) 111–135. URL: https://doi.org/10.3233/AO-210246. doi:10.3233/AO-210246.

[38] L. Floridi, Information: A very short introduction, OUP Oxford, 2010. P. 23.

[39] A. Varzi, A. Cotnoir, Mereology, Oxford: Oxford University Press, 2021.

[40] A. C. Varzi, On the Boundary Between Mereology and Topology, in: Philosophy and the Cognitive Sciences. Proceedings of the 16th International Wittgenstein Symposium, Vienna, Hölder-Pichler-Tempsky, 1994, pp. 423–442.

[41] A. C. Varzi, Boundaries, continuity, and contact, Noûs 31 (1997) 26–58.

[42] C. Gilmore, et al., Location and Mereology, in: The Stanford Encyclopedia of Philosophy, Spring 2024 ed., Metaphysics Research Lab, Stanford University, 2024.

[43] M. Ashby, Y. Bréchet, D. Cebon, L. Salvo, Selection strategies for materials and processes, Materials & Design 25 (2004) 51–67. URL: https://www.sciencedirect.com/science/article/pii/S0261306903001596. doi:https://doi.org/10.1016/S0261-3069(03)00159-6.

[44] B. Bayerlein, M. Schilling, H. Birkholz, M. Jung, J. Waitelonis, L. Mädler, H. Sack, Pmd core ontology: Achieving semantic interoperability in materials science, Materials & Design 237 (2024) 112603. URL: https://www.sciencedirect.com/science/article/pii/S0264127523010195. doi:https://doi.org/10.1016/j.matdes.2023.112603.

[45] R. Wasserman, Material Constitution, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Fall 2021 ed., Metaphysics Research Lab, Stanford University, 2021.

[46] N. Hartmann, Possibility and Actuality, De Gruyter, Berlin, Boston, 2013. URL: https://doi.org/10.1515/9783110246681. doi:doi:10.1515/9783110246681.

[47] D. K. Lewis, On the Plurality of Worlds, Wiley-Blackwell, Malden, Mass., 1986.

[48] T. Yagisawa, Possible Objects, in: E. N. Zalta, U. Nodelman (Eds.), The Stanford Encyclopedia of Philosophy, Summer 2023 ed., Metaphysics Research Lab, Stanford University, 2023.

[49] C. Hitchcock, J. Woodward, Explanatory generalizations, part ii: Plumbing explanatory depth, Noûs 37 (2003) 181–199. doi:10.1111/1468-0068.00435.

[50] E. Mach, Die Principien der Wärmelehre. Historisch-kritisch entwickelt, J.A. Barth, Leipzig, 1900.

[51] J. Ladyman, D. Ross, Every Thing Must Go: Metaphysics Naturalized, Oxford University Press, 2007. doi:10.1093/acprof:oso/9780199276196.001.0001.

[52] W. C. Salmon, Causality and Explanation, Oxford University Press, 1998. doi:10.1093/0195108647.001.0001.

[53] N. Cartwright, Causal Powers: What Are They? Why Do We Need Them? What Can Be Done with Them and What Cannot?, 2007.

[54] A. Marmodoro, E. Mayr, Metaphysics: An Introduction to Contemporary Debates and Their History, Oup Usa, New York, 2019.

[55] B. M. Loewer, Mental causation, or something near enough, in: B. P. McLaughlin, J. D.

Cohen (Eds.), Contemporary Debates in Philosophy of Mind, Blackwell, 2007, pp. 243–64.

[56] J. Kim, Causation and mental causation, in: B. P. McLaughlin, J. D. Cohen (Eds.), Contemporary Debates in Philosophy of Mind, Blackwell, 2007, pp. 227–242.

[57] S. H. Kellert, H. E. Longino, C. K. Waters (Eds.), Scientific Pluralism, University of Minnesota Press, Minneapolis, 2006.

[58] J. Bickle, Multiple Realizability, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Summer 2020 ed., Metaphysics Research Lab, Stanford University, 2020.

[59] T. E. Tahko, Unity of Science, Cambridge: Cambridge University Press, 2021.

[60] N. Guarino, D. Oberle, S. Staab, What is an ontology?, in: Handbook on ontologies, Springer, 2009, pp. 1–17. doi:10.1007/978-3-540-92673-3_0.

[61] N. Lazzari, et al., Sandra - A Neuro-Symbolic Reasoner Based On Descriptions And Situations, ArXiv abs/2402.00591 (2024).

[62] B. Zied, et al., From Shallow to Deep Interactions Between Knowledge Representation, Reasoning and Machine Learning, CoRR abs/1912.06612 (2019). arXiv:1912.06612.

[63] H. J. Levesque, R. J. Brachman, Expressiveness and tractability in knowledge representation and reasoning, Computational Intelligence 3 (1987) 78–93. doi:10.1111/j.1467-8640.1987.tb00176.x.

[64] F. A. Zaccarini, C. Masolo, E. Ghedini, S. Borgo, From causation (and parthood) to time: the case of emmo, FOIS 2023 Proceedings (2023) tbd.

[65] A. Atkin, Peirce's Theory of Signs, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Summer 2013 ed., Metaphysics Research Lab, Stanford University, 2013.

[66] L. Floridi, The Philosophy of Information, Oxford University Press, 2011. URL: https://doi.org/10.1093/acprof:oso/9780199232383.001.0001. doi:10.1093/acprof:oso/9780199232383.001.0001.

[67] A. Varzi, Mereology, in: E. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Spring 2019 ed., Metaphysics Research Lab, Stanford University, 2019.

[68] W. C. Salmon, Scientific Explanation and the Causal Structure of the World, Princeton University Press, 1984. URL: http://www.jstor.org/stable/j.ctv173f2gh.

[69] P. Dowe, Causality and conserved quantities: A reply to salmon, Philosophy of Science 62 (1995) 321–333. doi:10.1086/289859.

[70] D. J. Griffiths, Introduction to Quantum Mechanics, 2nd ed., Pearson Prentice Hall, 2004.

[71] G. 't Hooft, M. J. G. Veltman, Diagrammar, NATO Sci. Ser. B 4 (1974) 177–322.

[72] M. Esfeld, C. Sachse, Conservative Reductionism, New York: Routledge, 2011.