Evaluating the Basic Formal Ontology as a top-level framework for the interoperability of Pistoia Alliance Ontologies

Giacomo De Colle[1,2], James Egan[2], Andrey Seleznev[3] and Asiyah Yu Lin[2,4*]

1 University at Buffalo, Department of Philosophy, Buffalo NY, USA
2 National Center for Ontological Research, Buffalo NY, USA
3 University of Pittsburgh. Pittsburg PA, USA
4 Pistoia Alliance, Inc. Wakefield MA, USA

Abstract

Pistoia Alliance, a not-for-profit consortium in life sciences and healthcare, encourages collaborative efforts to deliver data-driven value by leading ontology projects aimed at ensuring alignment and interoperability across the pharmaceutical, life-science and healthcare industries. The Identification of Medicinal Products Ontology (IDMP-O) was created to ontologically represent the IDMP ISO standards. More ontologies are currently under development by Pistoia Alliance members. To ensure the utility of those ontologies, we hypothesized that the Basic Formal Ontology (BFO) could enable interoperability across them. After examinations of IDMP-O terms, its mid-level Object Management Group (OMG) terms, and BFO terms, we initiated efforts of mapping IDMP-O/OMG to BFO. Preliminary results and issues encountered are discussed here. Challenges reside in the nature and purpose of IDMP-O, which is mostly representing informational specifications for regulatory purposes, as well as the use of OMG:Role, which is far broader than BFO:role in semantics. Pistoia Alliance is committed to continuously evaluate BFO and to establish a connection between BFO and IDMP-O to enable interoperability between its ontologies.

Keywords

Identification of Medicinal Products Ontology (IDMP-O), Basic Formal Ontology (BFO), ISO standards, Object Management Group (OMG), interoperability, Pistoia Alliance

1. Introduction

In this paper, we share first results in evaluating the adoption of the Basic Formal Ontology (BFO) [1] as an upper-level ontology for the purpose of increasing semantic interoperability between ontologies created and implemented by Pistoia Alliance. Pistoia Alliance is a not-for-profit, pre-competitive industry consortium fostering collaboration in life sciences and healthcare. One of its strategic priorities is to deliver data driven value for the life science, health care, and pharmaceutical industries. As part of this effort, Pistoia Alliance leads multiple

Proceedings of the Joint Ontology Workshops (JOWO) - Episode X: The Tukker Zomer of Ontology, and satellite events co-located with the 14th International Conference on Formal Ontology in Information Systems (FOIS 2024), July 15-19, 2024, Enschede, The Netherlands.
* Corresponding author
gdecolle@buffalo.edu (G. De Colle); jamesjegan3@gmail.com (J. Egan); ibs15@pitt.edu (A.Seleznev); asiyah.yulin@pistoiaalliance.org (A.Y. Lin)
0000-0002-3600-6506 (G. De Colle); 0000-0002-3600-6506 (J. Egan); 0009-0002-6904-6229 (A. Seleznev); 0000-0003-2620-0345 (A.Y. Lin)

ontology projects targeting semantic data with the aim of providing a holistic view and ensuring that deliverables are aligned and interoperable. For example, the Identification of Medicinal Products Ontology (IDMP-O) aims to enable semantic interoperability based on FAIR principles to enhance and augment the existing ISO IDMP standards [2] for use in regulatory reporting. The IDMP standards, initially published by ISO in 2012, provide a framework to uniquely identify and describe medicinal products with consistent documentation and terminologies to provide for the reliable exchange of product information between global regulators, manufacturers, suppliers, and distributors [3]. The IDMP standards were originally developed based on regulatory information exchange requirements to facilitate consistent pharmacovigilance and the safety of medications throughout the world, as well as the global supply chain's integrity [3, 4]. They include 5 ISO standards and technical implementation guides: (1) ISO 11615: Data elements and structures for the unique identification and exchange of regulated medicinal product information; (2) ISO 11616: Data elements and structures for the unique identification and exchange of regulated pharmaceutical product information; (3) ISO 11238:Data elements and structures for the unique identification and exchange of regulated information on substances; (4) ISO 11239: Data elements and structures for the unique identification and exchange of regulated information on pharmaceutical dose forms, units of presentation, routes of administration and packaging; and (5) ISO 11240: Data elements and structures for unique identification and exchange of units of measurement [4, 5]. An initial 1.0 release of IDMP-O was published in early January 2024, with a 1.1 release published at the end of Q1 2024 and further updates planned. Related emerging Pistoia ontology projects include the Pharmaceutical CMC Process ontology [6] aiming to standardize Chemistry, Manufacturing, and Control (CMC) recipe data describing chemical and biological production processes, and the Clinical Operations (ClinOps) ontology [7] that enhances clinical research efficiency through standardized terminologies around clinical trial execution and the collection, analysis and exchange of clinical operations data from the clinical protocol.

Here, we focus on analyzing the IDMP Ontology® (IDMP-O) [2]. The aim of IDMP-O is to create a shared semantic background to mitigate costs and issues in sharing data related to IDMP regulations compliance. Compliance with IDMP ISO standards has been mandated by the European Medicines Agency [4], and the U.S. Food and Drug Administration (FDA) has published related implementation guidance [5]. IDMP-O will play a key role to ensure consistency across implementations. IDMP-O extends from a number of Object Management Group (OMG) ontologies [8], including the Commons Ontology Library, Multiple Vocabulary Facility (MVF) and Languages, Countries and Codes (LCC), and reuses some of the Semantic Publishing and Referencing Ontologies (SPAR) ontologies [9]. IDMP-O consists of multiple modules representing different parts of various ISO documents, and we use the term "IDMP-O" to refer to all the modules taken together.

BFO is a widely adopted top-level ontology and an ISO standard (ISO/IEC 21838-2), which is used as a top-level ontological architecture for more than 600 projects. As such, BFO serves as a baseline for information sharing practices, and enables coherent interoperability of heterogeneous data. To meet Pistoia's goal of delivering data driven value ontologies such as IDMP-O, Pharmaceutical CMC Process, and Clinical Operation (ClinOps) need to be interoperable. The purpose of BFO, and of other top-level ontologies, is to provide a small, shared set of terms composing a top-level hierarchy that other ontologies can reuse for downward population. In the application considered for this paper, the hypothesis is that BFO

would enable interoperability across Pistoia ontologies. More specifically, a mapping from IDMP-O to resources in the BFO ecosystem would allow for IDMP-O to be semantically integrated with other ontologies which are already part of Pistoia Alliance products and adopt BFO.

At the time of writing, development of other Pistoia ontologies was at its beginning, but version 1.1 of IDMP-O was already released [10]. Our initial steps, which are partially discussed in this paper, include an overall assessment of BFO and IDMP-O with the aim of identifying commonalities between IDMP-O and BFO or BFO-compliant ontologies. Other resources coming from other BFO-based ontologies, such as the Common Core Ontologies (CCO) and Industry Foundry Core (IOF) [11, 12, 13], will be introduced in the appropriate place in the rest of the paper.

After the assessment phase, we started to identify possible relations between classes in IDMP-O and in the BFO ecosystem. Next sections discuss our development of this process, as well as issues encountered. IDMP-O was created with the intent of faithfully representing the terminology created in the IDMP standards, which, as we mentioned earlier, are primarily designed to serve as common standards for improved data sharing during regulatory processes, such as pharmacovigilance, submission, and global supply chain. In the rest of this paper, we motivate a process of mapping of IDMP-O into classes developed in the BFO ecosystem, we describe first steps achieved in identifying such mappings, and we discuss currently open questions and next steps to take.

2. Methods

The project involved members of Pistoia Alliance, developed over several months and included multiple steps of consultation with technology partners of Pistoia Alliance as well as of the BFO ontology communities. As a first step, we reviewed and identified the main classes used in IDMP-O and their use case as motivated by the adopted ISO standards. We compared IDMP-O and BFO terms, both manually and using the Lexical OWL Ontology Matcher (LOOM) algorithm implemented on BioPortal as an open-source software. In the second phase of the project, we started identifying possible mappings between IDMP-O and classes in the BFO community. While developing these mappings, we shared our preliminary results with some of the IDMP-O primary developers. When the issues we faced touched foundational issues in BFO, as for example is in the case of mapping OMG:Role which we describe in the next section, we also reached out to other members of the BFO community.

During the second phase of development, we considered two possible methods to adopt in order to connect IDMP-O and BFO. The first includes a complete alignment of IDMP-O to BFO, which would include a subsumption of IDMP-O classes under terms in the BFO hierarchy, as well as a reconstruction of the definitions introduced by IDMP-O to reflect the new hierarchy. The second possible method is the creation of mapping triples from BFO and IDMP-O. A mapping triple is a <s, p, o> triple, where the subject ‘s’ and predicate ‘p’ are terms from two different ontologies, and the predicate ‘p’ is a relation appropriately connecting the two [14, 15].

Among the different strategies for creating mappings, we evaluated the adoption of the method of translation definitions, already employed by Michael Gruninger and associates as part of the COLORE project [16]. A translation definition is a set of relations between one or

more terms in one ontology and a formula in another ontology, such that the union of the first set of terms with the translation definition entails the corresponding formula in the other ontology. For example, as we will discuss below, IDMP-O:Matter is synonymous with BFO:material entity, and as such can be connected by a predicate like owl:equivalentclass. This means that if something is an instance of IDMP-O:Matter, then that instance is also an instance of BFO:material entity, and vice versa.

3. Results

Table 1 shows some core terms introduced in IDMP-O. As discussed above, the purpose of IDMP-O is to represent ISO IDMP standards in order to provide a shared and clear semantic background when stakeholders implement standards to facilitate data exchange specified in regulatory documents. As such, IDMP-O is developed to closely match the terminology and use cases identified by the ISO IDMP standards. The definitions of IDMP-O terms are often the exact quote from the standards or adapted from them.

Table 1
Core IDMP-O and OMG terms and their definitions.

IDMP-O term or OMG term	Definition	Related ISO standards
role (OMG)	named specific behavior of something participating in a particular context	Adapted from ISO/IEC 19763-8:2015(en), Information technology - Metamodel framework for interoperability (MFI) - Part 8, clause 3.1.7
specification (OMG)	explicit requirement or set of requirements to be satisfied by something, such as a product, material, model, process or system	Adapted from ISO 6707-2:2017 Buildings and civil engineering works - Vocabulary - Part 2, clause 3.2.22
substance	specification for matter of defined composition that has discrete existence, whose origin may be biological, mineral or chemical	Adapted from ISO 11238:2018 Health informatics - Identification of medicinal products (IDMP) - clause 3.84
ingredient	role of a substance that is specifically part of or used in the preparation of some manufactured item, pharmaceutical product, medication, or drug	Resourced from ISO 11615:2017 Health informatics - Identification of medicinal products

pharmaceutical product	specification for the qualitative and quantitative composition of a medicinal product in the dose form approved for administration in line with the regulated product information	(IDMP) - clauses 3.1.28 and 9.7, Figure 12 Resourced from ISO 11239:2012 Health informatics - Identification of medicinal products (IDMP) - clause 3.1.21
medicinal product	specification for a pharmaceutical product or combination of pharmaceutical products that may be administered to human beings (or animals) for treating or preventing disease, with the aim/purpose of making a medical diagnosis or to restore, correct or modify physiological functions	Resourced from ISO 11239:2012 Health informatics - Identification of medicinal products (IDMP) - clause 3.1.17
matter	something that has mass and occupies space by virtue of having volume	Resourced from ISO 11238:2018 Health informatics - Identification of medicinal products (IDMP) - clause 3.41
process	structured set of activities involving various enterprise entities designed and organized for a given purpose	Resourced from ISO 18629-11:2005(en) Industrial automation systems and integration - Process specification language - clause 3.1.21

To faithfully represent the mentioned IDMP standards, the main classes used in IDMP-O refer to specifications of medicinal products based on the IDMP standards' purpose of "facilitating the reliable exchange of medicinal product information" [4], rather than physical products themselves. In Table 1, substance, ingredient, pharmaceutical product, and medicinal product are all subtypes of "specification". The OMG class "specification" refers to a set of requirements to be satisfied, say, in a process of manufacturing, or by a product which is the result of one such process. "Substance" in IDMP-O is a subclass of "specification", and it "specifies" a "physical substance", which is a subclass of "matter". The class "matter" is used to refer to physical material entities, and many specification classes point back to physical things by using a "specifies" relation. Now, a BFO:material entity is an independent continuant which has some portion of matter as part [1], and IDMP-O:Matter is defined as something that has mass and volume. Given that only matter can be said to have mass, IDMP-O:Matter is either equivalent to, or at the very least a subclass of, BFO:material entity. But notice that all instances

of BFO:material entity seem to also have mass and occupy space, satisfying the IDMP-O:Matter definition and giving strong support to the two classes being equivalent.

A preliminary comparison of BFO and IDMP-O reveals four overlapping terms with the same label but different semantics: process, role, site, and object.

Table 2
Overlapping IDMP-O and BFO 2020 terms and their definitions.

Terms	BFO 2020 definition	IDMP-O definition or OMG definition
Process	p is a process means: p is an occurrent that has some temporal proper part and for some time t, p has some material entity as participant at t	structured set of activities involving various enterprise entities designed and organized for a given purpose
Role	b is a role means: b is a realizable entity & b exists because there is some single bearer that is in some special physical, social, or institutional set of circumstances in which this bearer does not have to be & b is not such that, if it ceases to exist, then the physical make-up of the bearer is thereby changed.	named specific behavior of something participating in a particular context* （OMG term)
Site	b is a site means: b is a three-dimensional immaterial entity whose boundaries either (1) (partially or wholly) coincide with the boundaries of one or more material entities or (2) have locations determined in relation to some entity	place, setting, or context in which something is situated or to which something is, or may be, bound* （OMG term)
Object	an object is a material entity which manifests causal unity & is of a type instances of which are maximal relative to the sort of causal unity manifested	anything perceivable or conceivable.* (OMG term)

The comparison of definitions, axioms and editor notes of the above four terms revealed that IDMP-O:process is related (SKOS:closematch) to OBI:planned process[2], which is itself a subclass of BFO:process. Therefore, IDMP-O:process is a subclass (or SKOS:narrowerMatch) of BFO:process. The OMG term object is adapted from the ISO 1087 Terminology work and terminology science standard clause 3.1.1. The class OMG:object is defined so broadly that it is almost equivalent to OWL:thing. An OMG:Site is a subclass of OMG:Role, while a BFO:site is

[2] OBI refers to Ontology for Biomedical Investigations [17]. The definition of planned process in OBI is “A process that realizes a plan which is the concretization of a plan specification”.

not a BFO:role. In this paper, we focus on discussing OMG:Role as a first example in aligning BFO and OMG/IDMP-O.

OMG:Role, which refers to the specific behavior that entities have in a particular context when they play such a role, provides contextual information for resources in the ontology. “Ingredient” is an example of an OMG role. Relevantly for our discussion, a substance can be one type of ingredient in one medicinal product, but another in a different product. Therefore, IDMP-O decided to model ingredients as roles played by substances. These roles, in the IDMP-O case, are specified in a regulatory approved specification of the products. As such, “being an ingredient” appears to be an intimately contextual feature, not directly belonging to the physical features of the entity which happens to be an ingredient.

The debate over roles in formal ontology is wide and has produced discussions over many different semantic understandings of roles [18–20] and specifically over BFO:role [19]. OMG:Role has subclasses such as “agent role”, “functional role”, “process role”, “site”, “structural role”, and “undergoer”. IDMP-O mostly expanded from “functional role”, by adding for example, “ingredient”, “medication”, “moiety role” and “product role”. Not all of these subclasses seem to fit under the umbrella of BFO:role. For example, it is not clear whether OMG:site should be mapped into a BFO:site, which is not a subclass of BFO:role. The best possible solution is to implement a disjunctive mapping for the different subclasses of OMG:role, where each of them is mapped into a different class in the BFO or CCO hierarchies. However, the usage and the defined relations within OMG may complicate the mappings.

The differences between OMG:Role and BFO:role are exemplified by the fact that an OMG:Role doesn’t have any restriction on the type of resources it can be OMG:isPlayedBy, while a BFO:role is restricted by the BFO:inheres in relation that only allows a BFO:independent continuant (non-spatial region) as its range. For example, a BFO:role could be OMG:played by subclasses of OMG:Matter. On the other hand, specifications in IDMP-O can play roles. But a specification, as we have just discussed, is not a subclass of OMG:matter. A specification rather seems to be, in BFO terms, a type of generically dependent continuant (GDC), and more specifically a type of information content entity (ICE), as introduced in CCO [21]. CCO is a family of ontologies composing a mid-level architecture [13] extending the BFO, already widely adopted in defense and intelligence analysis domains. ICEs are defined in CCO as being in a relation of aboutness to some BFO:entity, in the same way in which an OMG:Specification specifies a physical entity.

The issue can be solved in two ways. The first is to allow a GDC to bear a BFO:role. The second is to introduce a cognate notion of roles to be adopted for representing contextual information about GDCs. Roles for GDCs seem to have a broader use even outside of the IDMP-O use case described in this paper. For example, a certain ICE could have the role of being a credential for a certain institution, or a password could have the role of being accepted to provide access to a certain system but not to another. Such addition would then be beneficial to broader use cases.

IDMP-O:substance also seems to fall under the scope of ICEs. An IDMP-O:substance seems to be a type of ICE which specifies some IDMP-O:Matter. As already argued, IDMP-O:Matter is equivalent to BFO:material entity. This would mean that an IDMP-O:substance is equivalent to a CCO:information content entity which CCO:is about some BFO:material entity. Table 2 shows the first tentative mappings achieved as a result of part 2 of the project. IDMP-O:process is

mapped into a subclass of BFO:process, even though the definition mentions "structured set of activities", which could make it equivalent to a CCO:act or CCO:planned act.

Table 3
First tentative mappings between terms in IDMP-O and terms in BFO or CCO.

IDMP-O/OMG term	Mapping predicate	BFO/CCO term
Role (OMG)	skos:relatedMatch	Role
Specification (OMG)	rdfs:subClassOf	Information Content Entity
Substance	owl:equivalentClass	ICE and is about some Material Entity
Matter	owl:equivalentClass	Material Entity
Process	rdfs:subClassOf	Process

Again, the evaluation of a more precise mapping in this case relies on further discussion with Pistoia Alliance stakeholders and with a closer investigation of the use of IDMP-O:process subclasses. Notice that, pending the issues described above, an OMG:Role is for the moment only a SKOS:relatedMatch with a BFO:role. Although our preference is for creating mappings which are able to connect terms through logical axioms in OWL, we used SKOS mappings as a provisional tool for providing human-readable relations between them, as well as for guiding future developments. We are currently also evaluating other efforts in automating or semi-automating mappings steps. At the time of writing, we were not able to identify any suitable match for subclasses of OMG:Role such as IDMP-O:ingredient.

4. Discussion

In the section above we identified preliminary potential mappings between core classes in IDMP-O, OMG, and classes in the BFO ecosystem in an attempt to evaluate BFO as a framework to enable the interoperability of IDMP-O and other Pistoia ontologies. This preliminary effort also identified standing issues in creating such mappings. One unresolved problem we have identified is the axiomatization of BFO:role, which doesn't allow for GDCs to bear roles. This means that, if an OMG:Specification is indeed a GDC, then it is not allowed to have a BFO:role, which impedes an equivalence mapping between OMG:Role and BFO:role. Moreover, some of the subclasses of OMG:role may not be subclasses of BFO:role at all.

Take as an example the design pattern of IDMP-O:ingredient [20] to further present the problem with mapping OMG:role. In IDMP-O, a "product composition" defines a "pharmaceutical product", and a "product composition" has as ingredient "ingredient". The IDMP-O "has ingredient" relation is a subclass of OMG:manifests, which "indicates a role that realizes, displays, or shows something, typically in some context". A "pharmaceutical product" comprises some "substance" and a "substance" plays the role of "ingredient". In this way, IDMP-O provides information about the substance playing an ingredient role in the context of a product composition. Figure 1 shows a preliminary rendition in equivalent BFO terms representing the above pattern. IDMP-O:defines is provisionally mapped into CCO:describes, IDMP-O:has ingredient has been provisionally mapped into CCO:prescribes, and IDMP-O:comprises has been provisionally mapped into BFO:has continuant part. The OMG:plays role object property doesn't seem to be equivalent to BFO:inheres in due to the "substance" being a

"specification". This "plays role" relation has been substituted by a new relation, "GDC-has role", that connects GDCs to roles. The exact mappings for these object properties are a subject for future discussions and cannot be exhaustively discussed in this place.

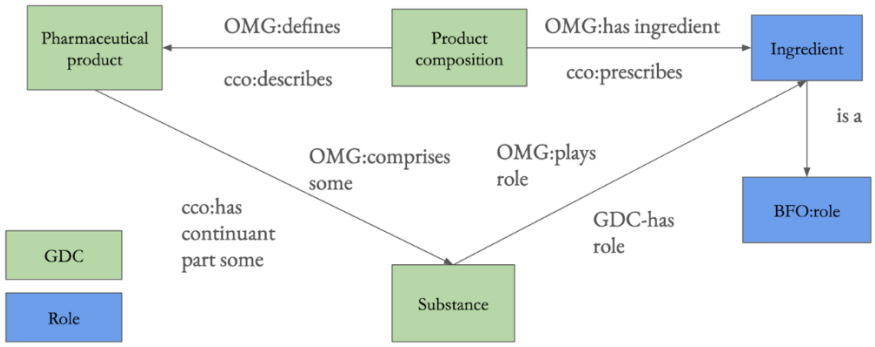

Figure 1: This figure shows a BFO-based approach to equivalently represent a partial IDMP-O pattern of substance and ingredient [19]. Pharmaceutical product, Product composition and Substance are all specifications and are mapped into BFO:GDC. Product composition has as ingredient a certain ingredient, which is a BFO:role. A new ad-hoc relation is created to relate a GDC to a role. Provisional mappings for other object properties are also provided.

Some terms defined in the original IDMP ISO standards are ambiguous or circular. In order to remove any ambiguity and ensure a consistent implementation of the standards across the industry, Pistoia Alliance IDMP-O project members as well as the primary ontology developers from EDMC went through lengthy discussions for the past two years to develop a more precise semantic understanding of those terms and to create the current version of IDMP-O. The agreement reached was to treat most of the terms as referring to the documentation and reports created when assessing regulatory compliance, for example, in the case of pharmacovigilance, or medicinal product traceability for global supply chain integrity [5].

For these reasons, many of the classes in IDMP-O would be subclasses of BFO:GDC, and the large amount of contextual information stored in some of these regulations, as we have seen, requires roles to be related to such entities. The issue can be solved by allowing for a GDC, such as a specification, to play a role. This solution can be implemented as native to IDMP-O or other Pistoia Alliance ontologies, or can be put forward as recommendations to the BFO community. As highlighted in this paper, there are multiple use cases for representing contextual information for a GDC. These range from the IDMP-O use cases discussed in this work and exemplified in Figure 1 to the representation of relations between GDCs or parts of the same GDC, as it is the case for example for parts of a document.

In many cases, these relations may be reduced to other models which can be created using existing BFO resources. Nevertheless, if only as a shortcut relation, allowing GDCs to bear roles seems to be an extremely powerful tool for ontological representation, which greatly simplifies models for use cases which require representing complex information entities. Discussion with members of the BFO development team revealed that this was already an open issue in the community. Introducing a way to connect GDCs to roles, or introducing a new class for

representing the specific types of roles borne by GDCs, would allow for BFO-based ontologies to specify these new relations and to represent a wider array of use cases.

5. Conclusion

Although ISO standards for identification of medicinal products provide definitions and terminologies, not all of these are directly usable to develop an ontology. These definitions can be circular, vague, or even missing. A consistent interoperable implementation of the IDMP ISO standards across companies and stakeholders is the motivation behind the creation of the IDMP-O ontology, since an ontological approach can clearly define and remove the ambiguity surrounding the terminology employed in ISO standards. IDMP-O is developed as part of efforts initiated by Pistoia Alliance alongside other life sciences, health care, and pharmaceutical industry ontologies, for example, the Pharmaceutical Process CMC Ontology and Clinical Operations Ontology. Both projects are based on the IOF-core ontology and the BFO ecosystem, therefore evaluating BFO for interoperability across those ontologies is critical. After evaluating the semantics of core IDMP-O terms and the higher level OMG terms, we came to the conclusion that although IDMP-O cannot be easily aligned to BFO, establishing mappings from IDMP-O/OMG to BFO and viceversa is indeed possible. Many of the terms in IDMP-O refer to documentations and prescriptions such as OMG:Specification, which are types of GDC in BFO. The first major obstacle we encountered is that IDMP-O makes use of an OMG:Role that is far broader than BFO:role in semantics. Furthermore, specifications in IDMP-O can play roles, but a GDC in BFO is not allowed to bear roles. The authors are working with developers in the BFO community to establish the IDMP-O use case, in order to potentially expand the use of BFO:role so that GDCs can be connected with roles.

The main purpose of IDMP-O is to faithfully represent IDMP standards for regulatory purposes. A lot of terms in IDMP-O are represented as specifications of the regulated medicinal product information. IDMP-O does connect specifications with physical entities via the "specifies" relation. Pistoia Alliance is furthermore developing other BFO-based ontological tools which focus on the R&D phase in a pharmaceutical product life cycle. The R&D world utilizes many BFO-based OBO foundry ontologies, whose primary focus is to represent physical entities, and not the regulations surrounding them. Connecting R&D data regarding physical processes and entities to regulatory data will only be feasible by adopting a common semantic understanding. Evaluating BFO as a top-level ontology to achieve these objectives and establishing a BFO-IDMP-O connection is a continuous effort to which Pistoia Alliance is committed, whose first steps we have presented in this paper.

Acknowledgements

The authors wish to acknowledge the insights gained from comments received by Pawel Garbacz from EDM Council, Elisa Kendall from Object Management Group, Dr. Christian Baber, the Chief Portfolio Officer from Pistoia Alliance, and John Beverley, assistant professor from University at Buffalo, SUNY and Co-director of the National Center for Ontological Research.

References

[1] R. Arp, B. Smith, A.D. Spear, Building ontologies with Basic Formal Ontology, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2015.

[2] C. Maskell, IDMP Ontology Project, (2022). https://www.pistoiaalliance.org/projects/current-projects/idmp-ontology/ (accessed April 26, 2024).

[3] Center for Drug Evaluation and Research, Identification of Medicinal Products (IDMP), FDA (2022). https://www.fda.gov/industry/fda-data-standards-advisory-board/identification-medicinal-products-idmp (accessed April 26, 2024).

[4] Information Management Division, Introduction to ISO Identification of MedicinalProducts, SPOR programme, (2016). https://www.ema.europa.eu/system/files/documents/other/wc500217406_en.pdf (accessed April 26, 2024).

[5] Food and Drug Administration, Identification of Medicinal Products — Implementation and Use Guidance for Industry, (2023). https://www.fda.gov/media/166736/download (accessed April 26, 2024).

[6] H. Taylor, Pharmaceutical CMC Process Ontology, (2023). https://www.pistoiaalliance.org/projects/pharmaceutical-cmc-process-ontology/ (accessed April 26, 2024).

[7] H. Taylor, Clinical Operations Ontology, (2024). https://www.pistoiaalliance.org/projects/current-projects/clinical-operations-ontology/ (accessed April 26, 2024).

[8] Object Management Group, THE OMG® SPECIFICATIONS CATALOG, (n.d.). https://www.omg.org/spec/.

[9] S. Peroni, D. Shotton, The SPAR Ontologies, in: D. Vrandečić, K. Bontcheva, M.C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee, E. Simperl (Eds.), Semantic Web – ISWC 2018, Springer International Publishing, Cham, 2018: pp. 119–136. https://doi.org/10.1007/978-3-030-00668-6_8.

[10] https://github.com/edmcouncil/idmp

[11] M.H. Karray, N. Otte, R. Rai, F. Ameri, B. Kulvatunyou, B. Smith, D. Kiritsis, C. Will, R. Arista, The Industrial Ontologies Foundry (IOF) perspectives, NIST (2021). https://www.nist.gov/publications/industrial-ontologies-foundry-iof-perspectives (accessed April 26, 2024).

[12] M. Jensen, G. De Colle, S. Kindya, C. More, A.P. Cox, J. Beverley, The Common Core Ontologies, (2024). http://arxiv.org/abs/2404.17758 (accessed May 1, 2024).

[13] J. Beverley, G. De Colle, M. Jensen, C. Benson, B. Smith, Middle Architecture Criteria, (2024). http://arxiv.org/abs/2404.17757 (accessed May 1, 2024).

[14] Prudhomme Tim, De Colle Giacomo, Liebers Austin, Sculley Alec, Xie (Karl) Peihong, Cohen Sydney and Beverley John, Mapping the Provenance Ontology to Basic Formal Ontology (*forthcoming*).

[15] Euzenat, J. & Shvaiko, P. Ontology Matching, 2nd edition.,Heidelberg: Springer, 2015.

[16] M. Gruninger, C. Chui, M. Katsumi, Upper Ontologies in COLORE, in: CEUR, Free University of Bozen-Bolzano, Italy, 2017. https://ceur-ws.org/Vol-2050/FOUST_paper_2.pdf (accessed May 1, 2024).

[17] A. Bandrowski, R. Brinkman, M. Brochhausen, M.H. Brush, B. Bug, M.C. Chibucos, K. Clancy, M. Courtot, D. Derom, M. Dumontier, L. Fan, J. Fostel, G. Fragoso, F. Gibson, A. Gonzalez-Beltran, M.A. Haendel, Y. He, M. Heiskanen, T. Hernandez-Boussard, M. Jensen, Y. Lin, A.L. Lister, P. Lord, J. Malone, E. Manduchi, M. McGee, N. Morrison, J.A. Overton, H. Parkinson, B. Peters, P. Rocca-Serra, A. Ruttenberg, S.-A. Sansone, R.H. Scheuermann, D. Schober, B. Smith, L.N. Soldatova, C.J. Stoeckert, C.F. Taylor, C. Torniai, J.A. Turner, R. Vita, P.L. Whetzel, J. Zheng, The Ontology for Biomedical Investigations, PLOS ONE 11 (2016) e0154556. https://doi.org/10.1371/journal.pone.0154556.

[18] F. Loebe, Abstract vs. social roles --Towards a general theoretical account of roles, Appl. Ontol. 2 (2007) 127–158.

[19] N. Guarino, BFO and DOLCE: So Far, So Close..., Cosm. TAXIS 4 (2017) 10–18.

[20] R. Mizoguchi, S. Borgo, YAMATO: Yet-another more advanced top-level ontology, Appl. Ontol. 17 (2022) 211–232. https://doi.org/10.3233/AO-210257.

[21] M. Jensen, G. De Colle, S. Kindya, C. More, A.P. Cox, J. Beverley, The Common Core Ontologies, Preprint (n.d.). https://arxiv.org/abs/2404.17758.

[22] E. Kendall, Pattern: Representing Ingredients - APPROVED - Pistoia Alliance IDMP Ontology - Confluence, (2023). https://dil-edmcouncil.atlassian.net/wiki/spaces/IDMP/pages/7246664/Pattern+Representing+Ingredients+-+APPROVED (accessed May 1, 2024).