January 2023

A method to improve alignments between domain and foundational ontologies

César BERNABÉ [a], C. Maria KEET [‡ b,1], Zubeida C. KHAN[‡ c] and Zola MAHLAZA [b]
[a] *Leiden UMC, The Netherlands*
[b] *University of Cape Town, South Africa*
[c] *CSIR, Pretoria, South Africa*
ORCiD ID: César Bernabé https://orcid.org/0000-0003-1795-5930, C. Maria Keet https://orcid.org/0000-0002-8281-0853, Zubeida C. Khan https://orcid.org/0000-0002-1081-9322, Zola Mahlaza https://orcid.org/0000-0001-9829-1480

Abstract. Foundational ontologies can be used to enable semantic interoperability in modern information systems. Aligning a domain ontology to a foundational ontology is perceived difficult, however. Reasons include confusing underlying concepts, understanding the philosophical ideologies of foundational ontologies, and lack of alignment guidance. For BFO, there is a BFO Classifier tool for alignment, but users still face challenges. To uncover some of these user challenges, an experiment was performed using 10 BFO-aligned domain ontologies. The alignment of domain entities were analysed, revealing seven different types of mistakes in the alignments. To avoid them, the BFO classifier tool was altered to improve the questions and explanations for the core principles of BFO. Thereafter, the BFO classifier tool was evaluated to measure the effect on alignment with a use-case based approach, using the GORO and AWO ontologies. The evaluation revealed that alterations facilitated alignment, as users felt more confident in their results given the improved understanding of the questions and possible answers.

Keywords. Foundational ontologies, BFO, Decision diagram, Methods and Tools

1. Introduction

Foundational ontologies are used to improve ontology development by providing a rich axiomatization for underlying classes on which domains can be modelled. The drawbacks to using foundational ontologies include increased start-up costs to learn about the ontology and lack of guidance [1]. The need for supporting methods to use foundational ontologies, particularly for the BFO foundational ontology, is demonstrated in a survey study [2], where it was found that BFO experts had conflicting opinions on alignment, concluding that a methodological framework is required for guidance. In addition, a recent literature review on the use of foundational ontologies in Bioinformatics observed that that most ontologies aligned to BFO are done at a shallow level, where elements are

[1]Corresponding Author: C. Maria Keet, University of Cape Town, South Africa, mkeet@cs.uct.ac.za.
‡ joint main authors

grounded only on bfo:Continuant or bfo:Occurrent, of which some argue that most of the authors are afraid of mistakes and prefer to 'play safe' [3]. This practice does not facilitate foundational ontology-mediated interoperability as much as it possibly could be. There are two main approaches to that: trying to align automatically and post-check the candidate alignments or guide the modeller from the start in a semi-automatic approach. We focus on the latter in this paper. To address this problem, the BFO Classifier was developed [4] as a decision diagram wrapped in a tool that allows the user to traverse through a series of questions in order to align a domain entity to the BFO hierarchy. The initial evaluation by its developers revealed that while it did assist in guiding the alignment process, there were several challenges and uncertainties regarding domain entity alignment and questions on diagram design more broadly. These include the wording of questions, the structure of questioning, whether a domain ontology developer could use it off-the-shelf without assistance, and quality of the output.

We aim to solve these issues in this paper and lay a more robust foundation for development of methods for foundational ontologies, and decision diagrams in particular. To inform the decision diagram design not only with explanatory material from its developers but also from the praxis perspective, we selected ten domain ontologies that were linked to BFO to extract insights on alignment. Seven root causes of evident misalignments were identified, that, to a greater or lesser extent, can be traced back to that core principles of BFO must have been not well understood by domain ontology developers. To avoid such issues, the decision diagram was improved by re-phrasing most questions and providing illustrative examples, to make it easier for ontology developers to align. This was implemented in a version 2 of the BFO Classifier tool. Both version 1 and version 2 were evaluated with two use cases by two ontology developers with ontologies in different subject domains. In both use cases, the experiments conducted using version 2 resulted in alignments to inner BFO classes when compared to version 1, as well as more consistent results when comparing with the original alignments of the ontologies. More fine-grained alignments of domain ontologies, in turn, may assist better with ontology quality and interoperability.

The remainder of this paper is structured as follows. Related work is described in Section 2. The design of the new decision diagram is reported on in Section 3, which is evaluated in Section 4. We discuss in Section 5 and conclude in Section 6.

2. Related work

Several works on methods and tools for domain ontology alignment to a foundational ontology exist. The first that was proposed was the decision tree D3 [5] to align ontologies to the DOLCE foundational ontology [6], which was also integrated with the Moki Semantic Wiki ontology development tool to assist the process of alignment. Moki is not available anymore, and therefore also not the D3 plugin. There are two for BFO. One is a theory-based classification that provides insight on how processes can be classified to BFO occurrents (i.e., those things that unfold over time), demonstrating that they can be used to for scientific data deriving from the measurement of processes of different types (e.g., cardiac events and running) [7]. We recently designed a decision diagram like D3 for BFO and implemented it in the stand-alone tool called "BFO classifier" [4], which is shown in Fig. 1. Its evaluation showed both promise to improve alignments, but also

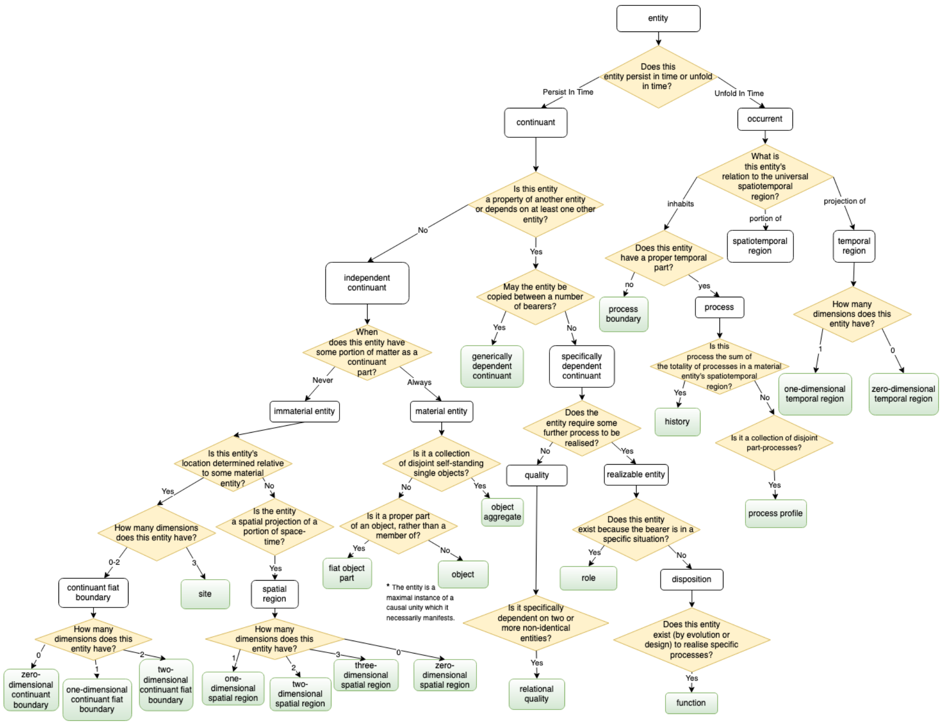

Figure 1. Original decision diagram (see supplementary material for a larger version); Source: based on [4].

challenges were remaining, specifically that the questions are still not easy even for ourselves and it is thus also not clear how well the diagram works for domain experts. Also, it was designed mainly by having it based on the descriptions in BFO supporting material only, but other source material is feasible, as well as possibly a different rearrangement of the decision tree and the potential effects of that on the alignments.

A different approach to solve the problem may be through indirect means: take the alignments between DOLCE and other FOs, including to BFO, and use the D3 decision tree to re-map it, or vice versa, take the BFO Classifier to align to BFO and from the BFO-to-DOLCE alignments obtain those. It is possible to do so with the ROMULUS ontology repository [8] and an algorithm to swap the ontologies [9]. However, it will be insufficient for practical purposes, because there are only 17 alignments between them and it is for BFO v1.0 rather than the current v2.0.

The third kind of approach to tackle the problem is with deep learning, to automatically classify terms into into high-level foundational ontology categories [10]. This work focuses on DOLCE and very high-level categories rather than a whole foundational ontology, and, at present, thus also lacks sufficient coverage for aligning a domain ontology to any particular foundational ontology.

The most widely used approach of alignment is still purely manual, with or without guidance of the original foundational ontology developers. This is done across domains and for multiple foundational ontology, such as in geology [11], health [12], biology [13], cybersecurity [14] and data mining [15], and all ontologies within or related to the OBO Foundry [16]. At present, any verification as to whether those alignments are correct—

aside from 'ask the foundational ontology developers'—is not available, with the closest, probably, D3 for DOLCE and the BFO Classifier for BFO. There are, of course, more FOs and many more ontologies linked to those FOs.

3. Decision diagram design

The overarching design process is that of combining several inputs from different angles, divided into three main categories. First, there are our previous experiences with designing and evaluating the first diagram [4]. That decision diagram was based on the respective entity descriptions in BFO documentation, the .owl file, and BFO teaching materials. The decision diagram design decisions, as described in detail in [4], concerned extracting salient distinction among the entities, simplifying some philosophy terminology, the 'shape' of the diagram when the sub-entities are more than two, and what to do with lone leaf entities that may give the sense of an incomplete answer when there is only 'yes' and 'none of the above' to terminate the sequence.

Second, the evaluation of version 1 provided insights that inform avenues to improve the diagram, being to further reduce BFO-specific ontology terminology with informal wording and to add examples. The BFO file has examples for a majority of the entities, but not all, and mostly in biology, but if it is to be understood and used outside that specialised domain, it may serve to devise examples that 'non-biology people', or who are in a different discipline within biology, may also understand; e.g., to avoid example entities with terms like meiosis.

Third, since our aim is to simplify and improve alignments and improve the quality of the alignments, we need to know what is aligned to what and where mistakes are made that should be avoided. Such insights may then be built into to decision diagram as a prevention mechanism.

3.1. Assessment of domain ontology alignments

The third strategy for designing or enhancing a decision diagram, i.e., making it informed by current alignment practices, is described in this section, using BFO and ten arbitrarily selected BFO-aligned domain ontologies.

3.1.1. Materials and methods

We followed the following procedure:

1. Search BioPortal for BFO entities and randomly select 10 ontologies that are aligned to it (i.e., are in the query answer).
2. In a spreadsheet, for each selected BFO-aligned domain ontology, put the entity that has a direct alignment to a BFO entity in that column of the BFO entity.
3. Examine alignments on both basic descriptive statistics and on any issues, such as evidently incorrect alignments, entities that have no alignments, etc. If that aligned entity is actually not of that domain ontology, but one that it reuses from another ontology, note that in brackets afterward.
4. For the unaligned and the mis-aligned, determine why and how that may be turned around into a prevention guideline.

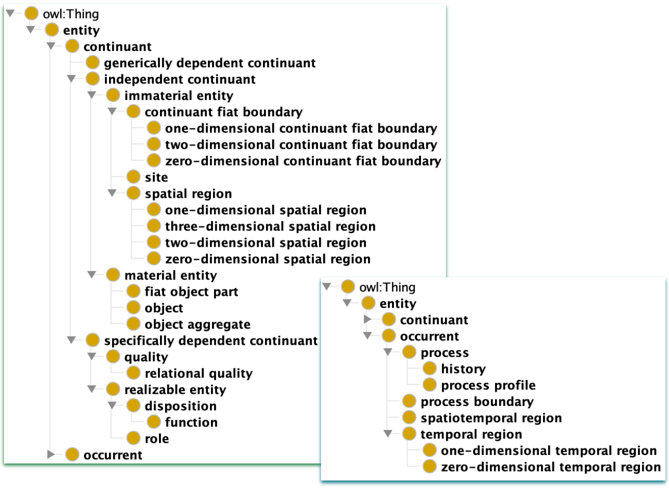

Figure 2. BFO 2.0; left: Continuant expanded; bottom-right: Occurrent expanded.

The following ontologies were selected for assessment: the Exercise Medicine Ontology (EXM) [17], Oral Health and Disease Ontology (OHD) [18], Interaction Network Ontology (INO) [19], Human Interaction Network Ontology (HINO), Ontology for Genetic Interval (OGI) [20], Ontology for Biomedical Investigations (OBI) [21], Ontology of Genes and Genomes (OGG) [22], Ontology of Genetic Susceptibility Factors OGSF [23], Ontology of RNA Sequencing (ORNASEQ) [24], and the Imaging Biomarker Ontology (IBO) [25].

3.1.2. Alignment analysis

Basic counting showed that of the 34 BFO entities, 17 do not have any alignment (50%). The leaf entities with 0 alignments are: two dimensional continuant fiat boundary, zero dimensional continuant fiat boundary, site, one dimensional spatial region, zero dimensional spatial region, fiat object part, relational quality, history, one dimensional temporal region, and zero dimensional temporal region, of which 10 are leaf entities in the BFO taxonomy. The one with the most aligned entities is material_entity with 103 subclass alignments.

Twenty-six other ontologies are used in those 10 BFO-aligned domain ontologies in such a way that they actually were the one(s) providing the alignment to BFO, rather than the domain ontology under evaluation. They are, in their chosen abbreviation: OGMS, IAO, SO, FMA, OGI, GAZ, OBI, CELL, GO, CL, CHEBI, EFO, UBERON, OGG, CARO, REO, VO, Lifo, COB, ENVO, PATO, HP, IDO, QIB, INO, and OMRSE, which are available online from the BioPortal[2].

The analysis of the aligned entities resulted in a list of seven type of mistakes in the alignments. The seven key issues are summarised in Table 1 and we will discuss each in turn. First, several domain entities are clearly too high up in the hierarchy, such as the two aligned to entity: ohd:sign and ohd:symptom (both due to OGMS's alignments, hence, a propagation), which clearly are not top-level entities on par with continuant and

[2] `https://bioportal.bioontology.org/`

Table 1. Summary of the types of alignment issues of the 10 assessed domain ontologies to BFO.

Issue	**Example of problematic alignment**
Too high up in the hierarchy	sign ⊑ entity
Wrong main branch	exercise intervention ⊑ material entity
Incorrect sibling	interrogation_point ⊑ site
Missing content in the ontology	guar gum is a stuff/amount of matter
Confusing the physical entity with the role it plays	enzyme ⊑ material entity
Not deep enough in the hierarchy	cell ⊑ material entity
Imprecision in FO entity (or misunderstanding)	the very diverse entities aligned to quality

occurrent. This also occurs 'halfway' in the taxonomy vs aligning to a leaf entity of the foundational ontology; e.g., hino:cell and obi:cell are aligned to material entity, but they are straightforward examples of its subclass bfo:object.

Second, some entities are aligned in a place that is clearly in the wrong main branch. The exm:exercise intervention cannot be a bfo:material entity, because an intervention is a processual entity that unfolds in time.

Third, the domain entity may be aligned to an entity in the right main branch at least, but it is aligned to another sibling than it should be. For instance, obi:interrogation_point is a point presumably, not a bfo:site that is a 3-dimensional entity, an ogi:author is a bfo:role rather than a bfo:material entity, and exm:exercise equipment with a term such as bfo:object is ambiguous whether what is meant is a single device, like the elliptical machine, or as aggregate, such as all the equipment that, say, the Zone Fitness branch (gym) in Rondebosch village has available for use.

Fourth, arguably, content may actually or perceived to be missing from the foundational ontology, from the perspective of the domain ontology at least. The common example is that BFO does not have stuff or amount of matter as distinguishing it from objects. Bio-ontologies, however, have many such entities, such as obi:environmental material (from ENVO), obi:lactic acid (from CHEBI) and obi:guar gum.

The fifth issue ought not to occur at this stage of ontology development anymore, being confusing the physical entity with the role it plays, such as enzyme ⊑ material entity, since enzyme inheres in the bearer that is the protein, and likewise for allergen that is a role that a substance plays in certain settings, but is not inherently so (e.g., pollen is just pollen, as are peanuts simply peanuts and to most people not an allergen), and ogi:author is not a material entity either, but also inheres in a material entity rather.

The penultimate issue we observed was that an entity was aligned in the right main branch, but needs to be aligned deeper down in the hierarchy. For instance, obi:atom (from CELL), hino:cell (from INO and CELL, CL), and ino:organism are all clearly a bfo:object, not just a material entity.

Last, issues may be attributed to imprecision in the foundational ontology or a widespread misunderstanding of the entity. A clear example of that for the BFO-aligned ontologies is bfo:quality, which has vastly different things aligned to it; among the 27, they include race, device setting, disability, genetic variation, fractured, information carrier, and chronic condition. Refinements such as in DOLCE may be of use to consider, such as not only the relational quality BFO has now, but also, e.g., a physical quality and a temporal quality.

3.2. Decision diagram

We shall first devise options for how to avoid alignment issues as described in the previous section, and then address remaining aspects of the design of the decision diagram.

The first issue, of too high up in the hierarchy, may be flipped around for guideline: "do not align to the top entity and if you don't know where else, then don't align it at all", since, in theory at least, those very top entities that are directly subsumed by bfo:entity are assumed to be exhaustive and so it must fit in either of those branches. If either of those does not seem to fit at first glance, then that domain entity may need further ontological investigation first to ascertain its nature. Not deep enough in the hierarchy can be addressed by additional probing of the modeller, both by using the questions in the decision diagram and the entity to be aligned that the decision diagram questions may facilitate in turn.

Avoiding the choice of clearly a wrong main branch may be addressed through a better explanation of core principles and providing examples of obviously different things, such as the object of a ball vs. the process of kicking it. This possible solution applies also to the third issue, of within main branch mistakes, and the fifth one on conflating the role with its bearer. That may be addressed further by pitting a subclass against its siblings rather than relative to its parent entity and, perhaps, using negation (alike with the Advocatus Diaboli tool [26]) or counterfactuals to double check the right answer option is chosen. For instance, in BFO's temporal region vs process (as direct subclasses of occurrent), the former is the duration (part of time) and the latter is that what's happening in that duration of time.

Such probing questions in the decision diagram should also assist with the sixth issue, on trying to align it deeper in the hierarchy and it may be partly achieved by tooling: if a new question immediately pops up once one is answered, one may expect the behaviour that a modeller at least reads and considers it. That is, the very nature of a decision diagram contributes to addressing this issue.

Issues four and seven, on perceived to be missing content anywhere and for bfo:quality in particular cannot of itself be addressed by a decision diagram but only either 1) in a supporting method and tool or 2) by the foundational ontology developers to extend the ontology. For instance, one could ask the modeller in a tool interface to propose an additional entity or to offer the feature to insert a new entity. Candidates for refinements with one or more sub-entities are material entity, site, object aggregate, quality, and disposition. Since some entities in BFO have only one subclass, like quality has only relational quality, at least adding a sibling would help.

Finally, these steps for addressing alignment issues also address the issues observed with the BFO Classifier v1 regarding adding examples. We also have made changes to terminology to try to find easier wording and better-known synonyms. For instance, instead of the mentioning of continuants like in "When does this entity have some portion of matter as a continuant part?", to ask for "Does this entity always have some matter (space and mass) as a part that persists in time?" with as answer options [always/never]. Illustrative examples are also now in the decision diagram; e.g., for the bfo: process-boundary, the example is "the first or last step of a Kizomba dancing session".

The final, revised, decision diagram is shown in Figure 3 (continuant branch) and Figure 4 (occurrent branch). As can be seen when compared to the original one in Figure 1, the structure of the tree is the same, but the questions and answers have been modified.

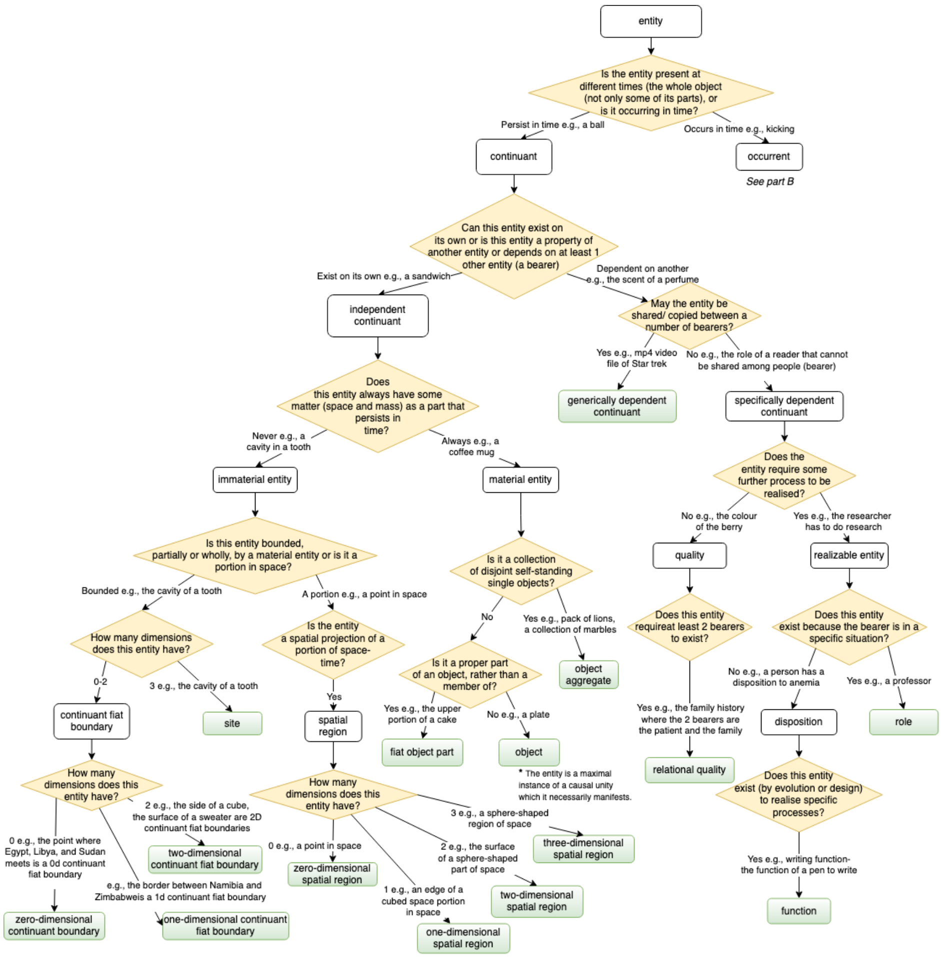

Figure 3. Revised decision diagram - part A - (see supplementary material for a larger version). Green/shaded boxes: leaf nodes in BFO v2.

4. Evaluation

We evaluated the BFO classifier diagram and tool using a use-case validation in two different subject domains. First, we checked the alignment results by comparing the tool's output with previous work from the literature. Second, we contrasted the outcomes of the first and second version of the tool to assess the impact of the changes on the decision diagram. We also took into account the perceptions of the participants while using the tool, even though the tool is more of a vehicle than that it has been subjected to human-computer interaction design principles.

Both evaluation cases were performed by inviting authors of domain ontologies to use the BFO classifier tool. The evaluators, one of whom is an author of one of the two ontologies used in this evaluation, were invited to collaborate on this paper at a later stage, to ensure they would not be primed by the decision diagram design process. The

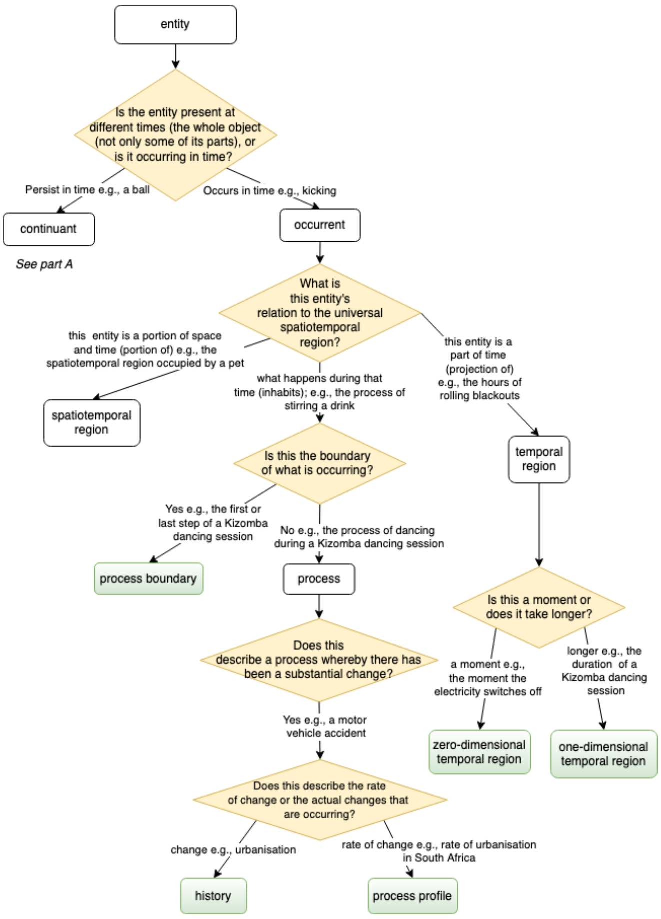

Figure 4. Revised decision diagram - part B - (see supplementary material for a larger version). Green/shaded boxes: leaf nodes in BFO v2.

ontologists have never used the BFO classifier tool before participating in the evaluation experiment.

4.1. Materials and methods

The main aim of the evaluation was to assess the BFO classifier's usability and effectiveness. Here, various permutations of use cases are possible, such as aligning a domain ontology which is not previously grounded on a foundational ontology, or creating a new ontology already aligned to BFO through the use of the tool. We chose two distinct options that would allow us to compare ours to previous results. First, we performed an alignment of the African Wildlife Ontology, which had been previously manually aligned to BFO v1.1. Second, we conducted a realignment of the Goal-Oriented Requirements Ontology (GORO) [27], which is originally grounded on the Unified Foundational Ontology (UFO) [28]. The former would be useful to see if there are any changes to the alignments and the, somewhat subjective, ease of alignment. The latter has to do in part with terminology use, which is different across foundational ontologies, and whether the distinctions BFO makes are understandable from that framework, and, hence, indirectly an aid in learning about BFO.

The African Wildlife Ontology is a tutorial ontology to demonstrate a range of aspects of ontology development and ontology engineering. Its contents focuses on both basic general knowledge of African wildlife, and with branches into the scientific domain as well. For this experiment, AWO version 3b was used.

GORO focuses on consensually describing the Goal-Oriented Requirements Engineering (GORE) field, which is a sub-field of Software Engineering. The ontology was created following the Systematic Approach for Building Ontologies (SABiO) [29] and it incorporates domain knowledge extracted from concepts from eight distinct goal-modelling languages. GORO is part of and extends other ontologies of the Software Engineering Ontology Network (SEON) [30], which uses UFO as its backbone. GORO aims to facilitate the analysis and improvement of existing goal-modelling languages, and to solidify GORE-related concepts.

The evaluation procedure for the two use cases was similar. For AWO, one of the authors (who did not conduct the alignment) removed the previous BFO alignments from the ontologies to then pass it on to the domain ontology developers to perform the alignments using either BFO Classifier version 1 and save the axioms in the OWL file or to do this with the BFO Classifier version 2 and to save the alignment axioms in a new OWL file. For GORO, the experimenter removed the UFO alignment themselves before commencing with the BFO alignment process with both versions of the classifier in sequence, also resulting in two new OWL files. During the evaluation, the ontologists were asked to take notes of their perception of the process.

Regarding the BFO Classifier tool versions, the only difference is the decision diagram. Both diagrams were converted to an XML file and the code compiled twice, once for each version of the diagram. Since a few questions with the examples are slightly longer in version 2, resizing of the window to read the full text is required, but this does not affect the diagram itself, which is deemed the key component under investigation.

4.2. Results

The results of the alignments showed a high level of consistency between the original mappings and the results of the classifier tool. The perceptions from the ontologists matched our expectations, while pointing directions for future improvements.

4.2.1. Alignments of AWO

A summary of the tool-based and manual alignments of AWO to BFO is shown in Table 2. Alignment to the AWO was relatively straightforward, except for, mainly, Distribution that both evaluators questioned because it was not clear what it referred to precisely. One evaluator considered that it can be seen as an event unfolding in time (e.g., monkeys spreading around the Atlantic forest throughout the years) or as a fact (e.g., elephants are found in a certain region of Africa) and therewith a quality, and the other would have preferred to align it to immaterial entity pending further investigation into the meaning of it rather than the site that the questioning led them to. It is a similar challenge for habitat, with the addition that the original alignment was intentionally to spatial region only, because habitat was understood as an *n*-dimensional hypercube that is not present in BFO.

Table 2. Alignments from AWO to BFO using the first and second version of the BFO Classifier tool.

AWO entity	Original Alignment (BFO v1.1)	Version 1 Alignment (evaluator 1)	Version 2 Alignment (evaluator 2)
Animal	object	object	object
Distribution	two-dimensional region	site	quality
Eating	process	process-profile	process
Eating disposition	disposition	disposition	function
Habitat	spatial region	role	three-dimensional spatial region
Plant	object	fiat object part	object
Plant parts	object	object aggregate	fiat object part

The plant $\sqsubseteq$ fiat object part alignment for version 1 is due to a perceived unclear question phrasing in version 1, with "Is it a proper part of an object, rather than a member of?", where neither a 'yes' nor a 'no' were seen as ideal, but a 'yes' chosen tentatively. This is disambiguated in version 2, with the examples of 'upper portion of the cake' for 'yes' and 'a plate' as example for 'no', and worked as intended. In the same corner of the decision diagram, plant parts was mis-aligned due to lack of annotation in the ontology and not consulting its subclasses, therewith leaving the meaning unclear to the evaluator. The evaluator also noted that, for safety, they would rather have had it aligned to material entity for the time being and only refine the alignment after consultation, but they felt compelled to continue with the questions.

Last, both evaluators commented on eating disposition. For version 1, "Does this entity exist (by evolution or design) to realise specific processes?", the evaluator intentionally did not answer (hence, the alignment remained at bfo:disposition), noting they could not see a case where the answer is 'no'. The example in version 2—that the function of a pen to is to write—led the other evaluator to answer 'yes'.

4.2.2. Alignments of GORO

The results of the alignments of GORO, together with its original alignment to UFO, are summarised in Table 3. Comparing the output of the two versions of the tool, it can be identified that alignments in the second version are grounded on classes deeper in the BFO hierarchy; e.g., the goro:Task concept (aligned from bfo:Process in version 1 to bfo:History in version 2). Comparing the two versions' output with the alignments to UFO, it can be noticed that the second version's alignment are also more consistent with the UFO theory. For instance, the change from bfo:Occurrent to bfo:Relational Quality on the goro:Conflict element makes it more consistent with the definition of ufo:Relator that represents a relationship between two entities, depending on their existence.

4.2.3. Perceptions from using the BFO classifier tool

The participants shared some similar perceptions about the usability of the classifier tool. Both ontologists agree that the tool facilitate the alignment process, and that the rephrasing of questions and the examples included in version 2 made the questions easier to understand. They suggested adding examples to the "None of the above" option to improve clarity. One of the participants mentioned that using the tool (especially version

Table 3. Summary of the alignments from GORO to BFO using the first and second version of the classifier tool. Note that it just shows the top level elements of GORO (i.e., the ones that inherit directly from UFO).

GORO Concept	Original Alignment (UFO)	Version 1 Alignment	Version 2 Alignment
Assumption	Belief	Disposition	Generically Dependent Continuant
Conflict	Relator	Occurrent	Relational Quality
Contribution Type	Quality	Generically Dependent Continuant	Quality
Requirement	Goal	Generically Dependent Continuant	Generically Dependent Continuant
Resource	Object	Material Entity	Object
Risk Event	Situation	Process	History
Stakeholder	Agent	Object	Object
Stakeholder Belief	Belief	Quality	Quality
Stakeholder Intention	Intention	Quality	Quality
Task	Plan	Process	History
Threatening Proposition	Proposition	Specifically Dependent Continuant	Specifically Dependent Continuant
Threatening Situation	Situation	Process	History

2) increased his confidence on the alignment of the ontology as compared to manual alignment. Another participant raised the need to include the ability to explore multiple paths of the decision tree in the tool without the need to commit to each one, until the end. For instance, when aligning awo:Habitat, one may have arguments of different strengths for considering the answers yes or no for the question "Is this entity a property of another entity or depends on at least one other entity?". Following the two alternative paths after considering yes and no, a user's answers may either lead them to either align awo:Habitat to bfo:role or bfo:three-dimensional spatial region. This then forces the user to interrogate their interpretation of the class under consideration. Technically, the consideration of multiple paths is feasible to implement. However, it requires the user to tediously re-answer questions after choosing what they believe to be the appropriate path in the decision tree.

5. Discussion

From the initial inspection of domain ontologies that were linked to BFO, it was discovered that several BFO entities were not used at all. It is unclear whether this non-use was due to terms that were not explained properly, or that the classes were simply not relevant for the domains of the selected ontologies. Conversely, perceived 'gaps' may not only be gleaned from the direct alignments, but also in domain ontologies. An example is the quantitative quality as subtype of bfo:quality and sibling of the lone bfo:relational quality in the COVID-19 epidemiology and monitoring ontology (CEMO) [31]. Additional indications on perceived gaps is the 'none of the above' option, which leaves some users with some consternation that they have reached the end unexpectedly. When unexpected, it suggests another answer option, i.e., sub-entity, may be needed.

An observation obvious in hindsight is that the questions themselves force a modeller to examine, perhaps again, the ontological nature of the entity they want to align. If one is uncertain or not knowledgeable enough about it, the alignment process is likely to falter either entirely, a repeated alignment may end up being aligned to a different entity, or when asking several people to do so it is aligned to a different entity. Any decision diagram cannot solve that of itself, but, rather, assist in pinpointing the issue to facilitate the ontological investigation thanks to having to answer the questions and therewith laying bare any uncertainties there may be about that domain entity, as was also observed with the awo:Distribution alignment. The flip side is that if one knows one's entity, the decision diagram, according to one of the evaluators, served as a confidence-building measure that the alignment was done correctly.

The analysis of a selection of BFO alignments may not have uncovered all types of alignment issues and the examples are for BFO alignments only. The approach, however, can be reused for other FOs to better guide the FO alignment process without the need for continuous requests to its developers. Further, since it is non-trivial to gather sufficient domain ontologists for a quantitative evaluation of FO alignment methods and tools, events like a hackathon dedicated to the task may yield additional data to assist enhancing FO alignment guidance.

Finally, the BFO Classifier as tool had mixed perceptions among the authors on whether it assists or hinders alignment. Looking ahead to explore consequences of a decision is easier with the diagram than the tool and it provides an overview that the tool does not show at present. The interface can be changed to include that view of the decision diagram. Other avenues for improvement concern ease of use, such as autocomplete of the entity to align, automatically resetting it once an axiom is inserted, and saving a question-and-answer trace for reuse for a similar entity. Integration with an editor may simplify keeping track of which entity is already aligned where as it was a back-and-forth with Protégé in the evaluation. The main reason why the BFO Classifier is a standalone tool, however, is brittleness of ontology editors, as is the problem with the D3 plugin [5] to an editor that is now defunct. Alternatively, one may create a web-based version of it. Either way, the current tool architecture should remain, in that it it needs only a modification to the XML file to update the decision diagram and a path declaration to the foundational ontology; hence, the tool design is generalisable to other ontologies.

6. Conclusions

This paper presents an improved decision diagram for aligning domain ontologies to BFO, BFO Classifier version 2. The improvements for the decision diagram were crafted following an assessment of existing domain ontologies that were linked to BFO. The alignments in the ontologies were analysed, revealing several alignment errors that were then used to create an improved decision diagram. The results of the evaluation indicated that the use of version 2 resulted in more consistent alignments as compared to prior literature. The improved questions and the addition of answer examples heightened users confidence in the classifier outputs.

Since the OBO foundry amounts to an intricate fabric of aligned ontologies and copies of alignments across domain ontologies, future work would look into how to possibly re-examine them to improve alignments. More practical engineering steps could also be taken, such as improving the tool's usability of the current functionality as well as extending it with, e.g., options to annotate why one answer was selected over an other, one's confidence in the answer, and perceived 'gap' where a BFO entity was expected but not present. Another avenue for research is to investigate whether the non-use of BFO terms is due to misunderstanding or irrelevance in domain ontologies. Finally, an additional research path is to study how the BFO classifier can be used to inform and constrain automatic alignment tools.

Acknowledgments The authors would like to thank Toky Raboanary for recompiling the BFO Classifier and verifying that mere swapping XML files of the trees indeed works.

Supplementary Materials Statement: The dataset of the evaluated ontologies, the data file, the generated alignment files, and the two BFO Classifier versions (source code and jar file) are temporarily available from `https://doi.org/10.5281/zenodo.7885614`.

References

[1] Keet CM. The use of foundational ontologies in ontology development: an empirical assessment. In: Antoniou G, et al., editors. 8th Extended Semantic Web Conference (ESWC'11). vol. 6643 of LNCS. Springer; 2011. p. 321-35. Heraklion, Crete, Greece, 29 May-2 June, 2011.

[2] Stevens R, Lord P, Malone J, Matentzoglu N. Measuring expert performance at manually classifying domain entities under upper ontology classes. Journal of Web Semantics. 2019;57:100469.

[3] Bernabé CH, Queralt-Rosinach N, Souza VES, da Silva Santos LOB, Jacobsen A, Mons B, et al.. The Use of Foundational Ontologies in Bioinformatics; 2022. Available from: `https://doi.org/10.21203/rs.3.rs-1929507/v1`.

[4] Emeruem C, Keet CM, Khan ZC, Wang S. BFO Classifier: Aligning domain ontologies to BFO. In: FOUST-VI: 6th Workshop on Foundational Ontology, part of JOWO'22. vol. 3249 of CEUR-WS; 2022. 15-19 August 2022, Jönköping, Sweden.

[5] Keet CM, Khan MT, Ghidini C. Ontology Authoring with FORZA. In: Proceedings of the 22nd ACM international conference on Conference on Information & Knowledge Management (CIKM'13). ACM proceedings; 2013. p. 569-78. Oct. 27 - Nov. 1, 2013, San Francisco, USA.

[6] Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A. Ontology Library; 2003. Http://wonderweb.semanticweb.org. WonderWeb Deliverable D18 (ver. 1.0, 31-12-2003).

[7] Smith B. Classifying processes: an essay in applied ontology. Ratio. 2012;25(4):463-88.

[8] Khan ZC, Keet CM. ROMULUS: a Repository of Ontologies for MULtiple USes populated with foundational ontologies. Journal on Data Semantics. 2016;5(1):19-36.

[9] Khan ZC, Keet CM. Automatically Changing Modules in Modular Ontology Development and Management. In: Proceedings of the South African Institute of Computer Scientists and Information Technologists (SAICSIT'17). SAICSIT. New York, NY, USA: ACM; 2017. p. 19:1-19:10. Thaba 'Nchu, South Africa. September 26-28, 2017.

[10] Junior AGL, Carbonera JL, Schmidt D, Abel M. Predicting the top-level ontological concepts of domain entities using word embeddings, informal definitions, and deep learning. Expert Systems with Applications. 2022;203:117291.

[11] Garcia LF, Abel M, Perrin M, dos Santos Alvarenga R. The GeoCore ontology: A core ontology for general use in Geology. Computers & Geosciences. 2020;135:104387.

[12] Pietra C, De Lotto R, Bahshwan R. Approaching Healthy City Ontology: First-Level Classes Definition Using BFO. Sustainability. 2021;13(24).

[13] Beisswanger E, Schulz S, Stenzhorn H, Hahn U. BioTop: An Upper Domain Ontology for the Life Sciences: A Description of Its Current Structure, Contents and Interfaces to OBO Ontologies. Applied Ontologies. 2008 dec;3(4):205–212.

[14] Obrst L, Chase P, Markeloff R. Developing an Ontology of the Cyber Security Domain. In: da Costa PCG, Laskey KB, editors. Proceedings of the Seventh International Conference on Semantic Technolo-

gies for Intelligence, Defense, and Security, Fairfax, VA, USA, October 23-26, 2012. vol. 966 of CEUR Workshop Proceedings. CEUR-WS.org; 2012. p. 49-56.

[15] Keet CM, Lawrynowicz A, d'Amato C, Kalousis A, Nguyen P, Palma R, et al. The Data Mining OPtimization ontology. Web Semantics: Science, Services and Agents on the World Wide Web. 2015;32:43-53.

[16] Smith B, Ashburner M, Rosse C, Bard J, Bug W, Ceusters W, et al. The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration. Nature Biotechnology. 2007;25(11):1251-5.

[17] Liu X, Shen B. In: Shen B, editor. Ontologies and Knowledge Graphs for Exercise Medicine. Singapore: Springer Nature; 2022. p. 1-16.

[18] Duncan WD, Thyvalikakath TP, Haendel MA, Torniai C, Hernández P, Song M, et al. Structuring, reuse and analysis of electronic dental data using the Oral Health and Disease Ontology. Journal of Biomedical Semantics. 2020;11.

[19] Özgür A, Hur J, He Y. The Interaction Network Ontology-supported modeling and mining of complex interactions represented with multiple keywords in biomedical literature. BioData Mining. 2016 Dec;9(1):41.

[20] Lin Y, Tarui H, Simons P. From Ontology for Genetic Interval (OGI) to Sequence Assembly - Ontology applying to next generation sequencing. In: Marshall MS, Burger A, Romano P, Paschke A, Splendiani A, editors. Proceedings of the Workshop on Semantic Web Applications and Tools for Life Sciences, Amsterdam, The Netherlands, November 20, 2009. vol. 559 of CEUR-WS; 2009. .

[21] Bandrowski A, Brinkman R, Brochhausen M, Brush MH, Bug B, Chibucos MC, et al. The Ontology for Biomedical Investigations. PLOS ONE. 2016 04;11(4):1-19.

[22] He Y, Liu Y, Zhao B. OGG: a Biological Ontology for Representing Genes and Genomes in Specific Organisms. In: Hogan WR, Arabandi S, Brochhausen M, editors. Proceedings of the 5th International Conference on Biomedical Ontology, ICBO 2014, Houston, Texas, USA, October 8-9, 2014. vol. 1327 of CEUR Workshop Proceedings. CEUR-WS.org; 2014. p. 13-20.

[23] Lin Y, He Y. The ontology of genetic susceptibility factors (OGSF) and its application in modeling genetic susceptibility to vaccine adverse events. Journal on Biomedical Semantics. 2014;5:19.

[24] Young MD, Wakefield MJ, Smyth GK, Oshlack A. Gene ontology analysis for RNA-seq: accounting for selection bias. Genome Biol. 2010 Feb;11(2):R14.

[25] Amdouni E, Gibaud B. Imaging Biomarker Ontology (IBO): A Biomedical Ontology to Annotate and Share Imaging Biomarker Data. Journal on Data Semantics. 2018 Dec;7(4):223-36.

[26] Ferré S, Rudolph S. Advocatus Diaboli – Exploratory enrichment of ontologies with negative constraints. In: ten Teije A, et al., editors. 18th International Conference on Knowledge Engineering and Knowledge Management (EKAW'12). vol. 7603 of LNAI. Springer; 2012. p. 42-56. Oct 8-12, Galway, Ireland.

[27] Bernabé CH, Souza VES, de Almeida Falbo R, Guizzardi RS, Silva C. GORO 2.0: Evolving an Ontology for Goal-Oriented Requirements Engineering. In: International Conference on Conceptual Modeling. Springer; 2019. p. 169-79.

[28] Guizzardi G. Ontological foundations for structural conceptual models. University of Twente; 2005.

[29] Falbo RA. SABiO: Systematic approach for building ontologies. In: Proc. of the 1st Joint Ws. on Ontologies in Conceptual Modeling and Inf. Systems Engineering.. vol. 1201. CEUR; 2014. .

[30] Ruy FB, Falbo RA, Barcellos MP, Costa SD, Guizzardi G. SEON: A Software Engineering Ontology Network. In: Knowledge Engineering and Knowledge Management. LNCS. vol. 10024. Springer; 2016. p. 527-42.

[31] Queralt-Rosinach N, Schofield PN, Hoehndorf R, Weiland C, Schultes EA, Bernabé CH, et al. The COVID-19 epidemiology and monitoring ontology. In: ISCB Bio-Ontologies 2021. Center for Open Science; 2021. .