Feasibility of Automated Foundational Ontology Interchangeability

Zubeida Casmod Khan[1,2] and C. Maria Keet[1]

[1] Department of Computer Science, University of Cape Town, South Africa
mkeet@cs.uct.ac.za
[2] Council for Scientific and Industrial Research, Pretoria, South Africa
zkhan@csir.co.za

Abstract. While a foundational ontology can solve interoperability issues among the domain ontologies aligned to it, multiple foundational ontologies have been developed. Thus, there are still interoperability issues among domain ontologies aligned to different foundational ontologies. Questions arise about the feasibility of linking one's ontology to multiple foundational ontologies to increase its potential for uptake. To answer this, we have developed the tool SUGOI, Software Used to Gain Ontology Interchangeability, which allows a user to interchange automatically a domain ontology among the DOLCE, BFO and GFO foundational ontologies. The success of swapping based on equivalence varies by source ontology, ranging from 2 to 82% and averaging at 36% for the ontologies included in the evaluation. This is due to differences in coverage, notably DOLCE's qualities and BFO and GFO's roles, and amount of mappings. SUGOI therefore also uses subsumption mappings so that every domain ontology can be interchanged, preserves the structure of the ontology, and increases its potential for usability.

1 Introduction

The growth in the amount of Semantic Web applications and ontology-mediated interoperability of complex software applications pushes demands for infrastructure to facilitate with semantic interoperability. Already from the early days of the Semantic Web, foundational ontologies have been proposed as a component to facilitate such interoperability, for they provide common high-level categories so that domain ontologies linked to them are also interoperable [7]. Over the past 15 years, multiple foundational ontologies have been developed, such as DOLCE, BFO [7], GFO [1], SUMO [9], and YAMATO [8]. This introduced the issue of semantic conflicts for domain ontologies that are linked to different foundational ontologies, if those foundational ontologies are indeed really different, and new questions for ontology engineers, chiefly:

1. Which foundational ontology should one choose to link one's domain ontology O_A to?
2. If O_A is linked to foundational ontology O_X, then is it still interoperable with domain ontology O_B that is linked to foundational ontology O_Y?

3. Is it feasible to automatically generate links between O_A and O_Y (which one may not know in sufficient detail), given O_A is linked to O_X?
4. If there are issues with the former, what is causing it? Or: in praxis, which entities of O_X are typically used for mappings with domain ontologies that may not be present, or present in an incompatible way, in O_Y?

The first question has been answered with ONSET [5], which considered the requirements but did not take into account what has been used in praxis. That is, ONSET cannot be used to answer the fourth question, which, may indicate actual modelling motivations and content of O_A to choose O_X over O_Y, assuming O_X has those things being represented in O_A that O_Y does not have. If O_Y has all those things as well, then there must be another reason why O_X was chosen.

The aim of this paper is to answer questions 3 and 4, above. To achieve this, we created SUGOI, a *Software Used to Gain Ontology Interchangeability*, which automatically interchanges the foundational ontology a domain ontology is linked to, as, to the best of our knowledge, no such tool exists yet. The current version can interchange between DOLCE, BFO, and GFO, for their mappings have been studied in detail [3, 6], but the system is designed for extensibility so as to handle any 'swap' (only new mapping files have to be provided). SUGOI and its supplementary material are available from the foundational ontology library ROMULUS at `http://www.thezfiles.co.za/ROMULUS/ontologyInterchange.html`. We conducted an evaluation with 16 ontologies, using both quantitative and qualitative means. The overall 'raw interchangeability' based solely on the equivalence mappings among the foundational ontologies, is 36.18% on average, ranging from 2.04% to 81.81%, depending mainly on the source ontology. If one permits subsumption mappings for the interchange, then one's foundational ontology can be fully swapped for another. The main reasons for any 'low' interchange is a combination of limited foundational ontology mappings that maintain consistency of the resultant ontology, and coverage of the foundational ontology, in particular when that is used in the domain ontology to foundational ontology alignment.

The remainder of the paper is structured as follows. The design of SUGOI is described in Section 2, which is followed by the experimental evaluation in Section 3. A discussion is presented in Section 4 and conclusions in Section 5.

2 SUGOI ontology interchangeability tool

To answer the questions concerning interchanging domain ontologies among foundational ontologies, we have developed SUGOI, Software Used to Gain Ontology Interchangeability. We describe the input files and algorithm of SUGOI.

2.1 SUGOI's input files

SUGOI has been designed to interchange domain ontologies between DOLCE, BFO, and GFO by using OWL mapping files. For the foundational ontology mediation, we will use the results obtained by [3, 6]: its equivalence and subsumption mappings between entities in the three ontologies have been investigated in detail, are logically consistent, and are available as machine-processable OWL files from the ontology repository ROMULUS [4]. Because several ontology files are used in the interchangeability, we describe here the terms used for each one:

- The *Source Ontology* (${}^{s}\mathcal{O}$) that the user wants to interchange, which comprises the *Source Domain Ontology* (${}^{s}\mathcal{O}_d$), with the domain knowledge component of the source ontology, the *Source Foundational Ontology* (${}^{s}\mathcal{O}_f$) that is the foundational ontology component of the source ontology that is to be interchanged, and any equivalence or subsumption mappings between entities in ${}^{s}\mathcal{O}_d$ and ${}^{s}\mathcal{O}_f$.
- The *Target Ontology* (${}^{t}\mathcal{O}$) which has been interchanged, which comprises the *Target Domain Ontology* (${}^{t}\mathcal{O}_d$), with the domain knowledge component of the target ontology, and the *Target Foundational Ontology* (${}^{t}\mathcal{O}_f$) that is the foundational ontology that the user has selected to interchange to, and any equivalence or subsumption mappings between entities in ${}^{t}\mathcal{O}_d$ and ${}^{t}\mathcal{O}_f$.
- *Mapping ontology*: the mapping ontology between the ${}^{s}\mathcal{O}_f$ and the ${}^{t}\mathcal{O}_f$.
- *Domain entity*: an entity from ${}^{s}\mathcal{O}_d$ or ${}^{t}\mathcal{O}_d$.

The algorithm is described in the next subsection.

2.2 Foundational Ontology Interchangeability Algorithm

The general idea of the algorithm behind SUGOI is that it accepts a ${}^{s}\mathcal{O}$ consisting of a ${}^{s}\mathcal{O}_d$ linked to a ${}^{s}\mathcal{O}_f$ (either DOLCE, BFO or GFO) and converts it to a ${}^{t}\mathcal{O}$ with a different ${}^{t}\mathcal{O}_f$. For this, SUGOI must have access to all the foundational ontologies and the mapping ontologies. The ${}^{s}\mathcal{O}$ is provided by the user. It does not matter whether the ${}^{s}\mathcal{O}_d$ is linked to a foundational ontology by an import or a merge. SUGOI accesses the remainder of the ontologies either by loading the ontology from the online URI, or by loading it from an offline file, depending on the version in use. Since the algorithm refers to independent ontology files, any changes in the foundational ontologies and mappings will not affect either the algorithm fundamentally or the software. Also, any implementation can be extended easily, as other foundational ontologies and mappings are developed by including the new ontology file paths or URIs.

Twenty mapping files are pre-loaded into SUGOI, allowing the user to interchange between DOLCE, BFO and GFO modules bi-directionally. These mappings do not result in an inconsistency, because any alignment that did that has been removed [3]. After the interchange process, all the domain entities from the ${}^{s}\mathcal{O}_d$ are present in the ${}^{t}\mathcal{O}_d$. SUGOI links domain entities from the ${}^{s}\mathcal{O}_d$ to the ${}^{t}\mathcal{O}_f$ as follows. SUGOI maps a domain entity's superentity in the ${}^{s}\mathcal{O}_f$ to its corresponding superentity in the ${}^{t}\mathcal{O}_f$ using the mapping ontology. This is illustrated in Fig. 1 for the entity dmop:DataType from the DMOP ontology [2], changing the link from DOLCE to one in GFO, and this resulting axiom is called a *GT, good target linking axiom.* If the domain entity's superentity does not have a corresponding mapping entity, SUGOI then treats that superentity as a domain entity and looks for a corresponding mapping entity at a higher level up in the taxonomy. Thus, eventually, the domain entity from the ${}^{s}\mathcal{O}_d$ is mapped with on-the-fly subsumption. This is displayed for interchanging the entity dmop:Strategy in Fig. 1, and this resulting axiom is called a *BT, bad target linking axiom.*

The main steps of the algorithm are thus as follows:

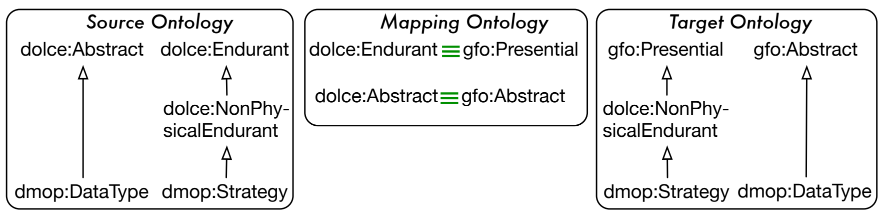

Fig. 1. Examples of interchanging domain entities dmop:DataType and dmop:Strategy from ${}^{s}\mathcal{O}_f$ DOLCE to ${}^{t}\mathcal{O}_f$ GFO with SUGOI, using equivalence and subsumption mappings, respectively.

1. Create a ${}^{t}\mathcal{O}$.
2. Copy axioms from the ${}^{t}\mathcal{O}_f$ to the ${}^{t}\mathcal{O}$.
3. Copy the ${}^{s}\mathcal{O}_d$ (domain axioms) to the ${}^{t}\mathcal{O}$.
4. Map the ${}^{s}\mathcal{O}_d$ domain entities to the ${}^{t}\mathcal{O}_f$ using the mapping ontology.
5. Perform on-the-fly subsumption if a domain entity from previous step is not linked to a ${}^{t}\mathcal{O}_f$.
6. Delete ${}^{s}\mathcal{O}_f$ entities that are not referenced by the domain entities in the ${}^{t}\mathcal{O}$.

The full algorithm can be accessed at the ROMULUS repository at the aforementioned URL, and it is illustrated in the next example.

Example 1. The BFO-aligned Subcellular Anatomy Ontology (SAO) is linked to DOLCE by SUGOI as follows.

1. Create a new ontology file, a ${}^{t}\mathcal{O}$: `sao-dolce.owl`.
2. Copy the entire ${}^{t}\mathcal{O}_f$ to the ${}^{t}\mathcal{O}$: copy the OWLized DOLCE ontology into `sao-dolce.owl`.
3. Copy the axioms from the ${}^{s}\mathcal{O}_d$ to the ${}^{t}\mathcal{O}$: e.g., the axiom sao:Membrane Surface $\sqsubseteq$ bfo:Object_boundary exists in the ${}^{s}\mathcal{O}$ SAO, which is added to the `sao-dolce.owl` ${}^{t}\mathcal{O}$ and is referred to as a 'new' axiom.
4. Change the 'new' axioms to reference ${}^{t}\mathcal{O}_f$ entities, if mappings exist: for the example in the previous step, no mapping exists for bfo:Object_boundary between BFO and DOLCE, so it proceeds to the next step.
5. If a mapping does not exist, perform on-the-fly subsumption: continuing with the example, bfo:Object_boundary has a superclass bfo:Independent_Continuant and the mapping ontology has bfo:Independent_Continuant $\equiv$ dolce:endurant, so bfo:Object_boundary $\sqsubseteq$ dolce:endurant is added to `sao-dolce.owl`.
6. Delete entities that exist in the ${}^{t}\mathcal{O}$ that are from the ${}^{s}\mathcal{O}_f$ but do not appear in an axiom with entities from the ${}^{t}\mathcal{O}_d$, resulting in the final ${}^{t}\mathcal{O}$, `sao-dolce.owl`.

3 Experimental Evaluation

The first purpose of the quantitative evaluation is to assess whether SUGOI successfully interchanges a ${}^{s}\mathcal{O}$ to a ${}^{t}\mathcal{O}$ and to determine the amount of the ontology that will be effectively interchanged, which refers to those entities within the ${}^{t}\mathcal{O}$ that have been mapped with equivalence relations, thereby not required to use parts of the ${}^{s}\mathcal{O}_f$ in the ${}^{t}\mathcal{O}$. Second, to carry out a qualitative assessment of the entities and axioms to uncover what contributes to (un)successful interchangeability. Finally, we consider two domain ontologies linked to foundational ontologies to cross-check whether there are any major differences between the manual and automated mappings.

3.1 Materials and methods

Materials The SUGOI desktop application online version was used for the automated interchangeability, Protégé v4.3 for data on ontology changes, the foundational ontologies in OWL, the mapping ontologies available in ROMULUS, and a set of real ontologies to evaluate. Based on the ${}^{s}\mathcal{O}_f$ of the ontologies, SUGOI loads five mapping files, for interchanging between DOLCE↔BFO, BFO↔GFO, GFO↔DOLCE, DOLCE↔GFOBasic, and GFOBasic↔BFO. The sample size was 16 source ontologies covering various subject domains, such as data mining, animals, dermatology, and spatial scenes, of which 5 have DOLCE as ${}^{s}\mathcal{O}_f$, 6 with BFO, and 3 with GFO, and 2 ontologies for the comparison between manual alignments to BFO and DOLCE and the automated alignments. All test files can be downloaded from the SUGOI page on ROMULUS.

Methodology The procedure for the experiment is as follows. *Preprocess domain ontologies* to checking whether each ${}^{s}\mathcal{O}$ uses the latest version of its respective ${}^{s}\mathcal{O}_f$ (DOLCE-Lite v397, BFO v1.1, and GFO v1.0) and whether the ontology import URIs are correct, and fix where necessary. *Perform interchangeability* by running SUGOI twice for each ${}^{s}\mathcal{O}$ to acquire its respective target ontologies; e.g., if the ${}^{s}\mathcal{O}_f$ is GFO, we generate two versions for the ${}^{t}\mathcal{O}$, one with DOLCE and another with BFO. *Evaluate interchangeability* as follows:

1. Compare the metrics of the domain entities of the target ontologies to those of the ${}^{s}\mathcal{O}$.
2. Compare the ${}^{t}\mathcal{O}$ to the ${}^{s}\mathcal{O}$ using the compare feature in Protégé. Protégé generates a list of entities that have been added, deleted and modified. Its 'modified entities' refers to the axioms that are used to define the entities, and whether they have been changed. In each ${}^{t}\mathcal{O}$, the entities that are modified are those that reference the mappable classes from their respective foundational ontologies.
3. Running the reasoner for the ${}^{t}\mathcal{O}$ to detect if there are unsatisfiable entities.
4. Analyse the metrics of the ${}^{s}\mathcal{O}_f$ entities that exist in the target ontologies.
5. Analyse the *raw interchangeability* of each ${}^{t}\mathcal{O}$, i.e., the amount of the ${}^{t}\mathcal{O}$ that has been correctly interchanged using equivalence mappings thereby not referring to the ${}^{s}\mathcal{O}_f$ entities. This is calculated from the ${}^{t}\mathcal{O}$ as follows: Let

GT, good target linking axioms, represent the sum of axioms that link domain ontology entities and ${}^t\mathcal{O}_f$ entities in the ${}^t\mathcal{O}$. Let *BT, bad target linking axioms*, represent the sum of axioms that link domain ontology entities and ${}^s\mathcal{O}_f$ entities in the ${}^t\mathcal{O}$; the raw interchangeability is calculated as follows:

$$\textit{Raw interchangeability} = \frac{|GT|}{|GT\ +\ BT|} \times 100 \qquad (1)$$

For instance, recall Fig. 1: the subsumption with dolce:NonPhysicalEndurant in the ${}^t\mathcal{O}$ counts toward the bad target linking axioms, whereas dmop:DataType $\sqsubseteq$ gfo:abstract counts as a good target linking axiom. SUGOI generates the raw interchangeability and ontology metrics in its log file for each ${}^t\mathcal{O}$.

6. Analyse and compare the DOLCE- and BFO-linked BioTop and Stuff ontologies with SUGOI's interchangeability. We interchange in both directions and compare the output with the original ontologies.

Because there are not many domain ontologies linked to a foundational ontology, there will be insufficient data to conduct a full statistical analysis to compare the results for different interchanges.

3.2 Results and discussion

We describe the analysis of the interchanged ontologies, a more detailed entity-level analysis, and then the comparison with the manual mappings.

Analysis of the interchanged ontologies After minor preprocessing of DMOP and SAO, all ontologies were successfully interchanged. Table 1 displays the domain entities in the ${}^s\mathcal{O}_d$ of the ${}^s\mathcal{O}$ and the domain entities of the ${}^t\mathcal{O}_d$-component of the ${}^t\mathcal{O}$ for the source ontologies. The 'modified entities' are those where a subsumption changed in the process; e.g., in the original DOLCE version of the DMOP ontology, the dmop:DecisionBoundary class is a subclass of dolce:abstract, which is modified in the GFO version, where dmop:DecisionBoundary is a subclass of gfo:Abstract. Such changes were collected from Protégé, which is illustrated for the example in Fig. 2.

Comparing the metrics of the domain entities of the ${}^s\mathcal{O}_d$ and ${}^t\mathcal{O}_d$ shows there are one or more extra domain entities in the ${}^t\mathcal{O}_d$, indicating the amount of ${}^s\mathcal{O}_f$ entities that have been added to the ${}^t\mathcal{O}_d$ (recall Example 1 and Fig. 1). The number of domain ontology entities increases from ${}^s\mathcal{O}_d$ to ${}^t\mathcal{O}_d$: e.g., $\delta = (749 - 739) = 7$ when interchanging DMOP to a GFO-aligned version (see Table 1), which is due to some absent mappings between the ${}^s\mathcal{O}_f$ and ${}^t\mathcal{O}_f$.

The same occurs with the set of BFO-aligned and GFO-aligned ${}^s\mathcal{O}$, and they differ for each case. Among others, bfo:Object_boundary is added to the ${}^t\mathcal{O}$ when the SAO ontology is interchanged from BFO to DOLCE, because there is no mapping from bfo:Object_boundary to a DOLCE entity, whereas in the interchange to GFO, bfo:object_boundary maps to gfo:Material_boundary, so sao:Membrane Surface becomes a subclass of gfo:Material_boundary.

The number of entities in each ${}^t\mathcal{O}$ that are from the ${}^s\mathcal{O}_f$ follows from the extra domain entities in Table 1. In terms of these metrics, ontologies with a small

Table 1. Comparison of the ${}^{s}\mathcal{O}_d$ entities to the ${}^{t}\mathcal{O}_d$ entities and raw interchangeability; interchang. = raw interchangeability and avg = average.

Source and Target ontology	${}^{s}\mathcal{O}_d$ **to** ${}^{s}\mathcal{O}_f$ **links**	**Domain Classes**	**Domain OP**	**Domain Entities**	**Modified entities**	**Raw interchangeability**
Source domain ontologies linked to DOLCE with target domain ontologies						
dmop	409	739	140	1350		
dmop-bfo		749	154	1377	43	4.88%
dmop-gfo		745	154	1373	54	11.65%
naive_animal	41	416	15	438		
naive_animal-bfo		422	24	453	20	21.95%
naive_animal-gfo		421	24	452	22	25.58%
ontoderm	14	239	28	301		
ontoderm-bfo		244	30	308	47	28.57%
ontoderm-gfo		243	30	307	49	42.85%
scene	18	172	74	246		
scene-bfo		175	78	253	3	50.00%
scene-gfo		174	78	252	4	50.00%
sego	43	75	43	139		
sego-bfo		89	55	165	32	29.54%
sego-gfo		88	53	162	32	36.36%
Correlation ${}^{s}\mathcal{O}_d$-${}^{s}\mathcal{O}_f$ links and interchang.:			*-0.79*	*Avg for DOLCE:*		*30.02%*
Source domain ontologies linked to BFO with target domain ontologies						
bco	26	63	62	146		
bco-dolce		67	62	150	19	65.39%
bco-gfo		66	62	149	20	67.86%
epidemiology	15	169	4	173		
epidemiology-dolce		174	4	178	13	60.00%
epidemiology-gfo		173	4	177	15	68.75%
ero	97	3910	123	4114		
ero-dolce		3918	123	4122	13	32.99%
ero-gfo		3917	123	4121	51	52.85%
ido	77	150	0	150		
ido-dolce		151	0	151	64	81.81%
ido-gfo		151	0	151	65	81.81%
proper_name	13	30	34	64		
proper_name-dolce		33	34	67	31	61.53%
proper_name-gfo		32	34	66	26	76.92%
sao	54	728	36	809		
sao-dolce		732	36	813	29	50.00%
sao-gfo		731	36	812	29	55.17%
Correlation ${}^{s}\mathcal{O}_d$-${}^{s}\mathcal{O}_f$ links and interchang.:			*-0.33*	*Avg for BFO:*		*62.90%*
Source domain ontologies linked to GFO with target domain ontologies						
pid	98	135	2	137		
pid-dolce		138	9	147	22	14.29%
pid-bfo		139	9	148	4	2.04%
gfo-bio	70	90	6	96		
gfo-bio-dolce		103	12	115	19	20.27%
gfo-bio-bfo		103	13	116	19	17.56%
gfo-bio-meta	14	127	6	139		
gfo-bio-meta-dolce		142	13	161	6	20%
gfo-bio-meta-bfo		142	13	161	8	19.69%
Correlation ${}^{s}\mathcal{O}_d$-${}^{s}\mathcal{O}_f$ links and interchang.:			*-0.66*	*Avg for GFO:*		*15.64%*
Correlation all ${}^{s}\mathcal{O}_d$-${}^{s}\mathcal{O}_f$ links and interchang.:			*-0.44*	*Avg for all:*		*36.18%*

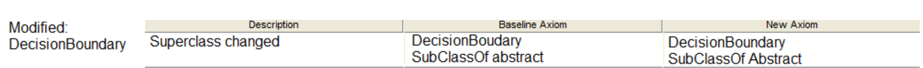

Fig. 2. A change for the dmop:DecisionBoundary class when interchanged to a GFO ${}^{t}\mathcal{O}$.

difference in ${}^s\mathcal{O}_d$ and ${}^t\mathcal{O}_d$ numbers, i.e., having a low number of ${}^s\mathcal{O}_f$ entities, perform best because they only contain few entities that cannot be mapped with equivalence, while ontologies with a high number of ${}^s\mathcal{O}_f$ entities perform worst because they contain many domain entities that cannot be mapped with equivalence. The IDO ontology has the least number of ${}^s\mathcal{O}_f$ entities in its ${}^t\mathcal{O}_d$ (only 1), thus it performs the best, whereas the gfo-bio-meta performs the worst with 15 ${}^s\mathcal{O}_f$ entities in its ${}^t\mathcal{O}_d$. The extra domain entities in a ${}^t\mathcal{O}$ ontology cause an increase in the number of *BT, bad target linking axioms*, which causes a lower raw interchangeability. The raw interchangeability values for the ${}^t\mathcal{O}$ files are shown in the last column of Table 1.

The interchanged ontologies are consistent, except for `ido-dolce.owl` and `ido-gfo.owl`, but `ido.owl` ${}^s\mathcal{O}$ was already unsatisfiable due to conflicting disjointness and subclass axioms among some domain entities.

After reasoning the ontologies, we manually compared the inferences of the domain entities of the ${}^s\mathcal{O}$ ontologies to the ${}^t\mathcal{O}$ ontologies to investigate whether different foundational ontologies influence these domain-specific inferences. For this set of domain ontologies, there was no change in the inferences.

Entity-level analysis Let us now consider those extra domain entities, which are those that are commonly used in domain ontologies, but do not have corresponding equivalence mappings among the foundational ontologies; or: the main 'culprits' for a low interchangeability. Table 2 displays these results. For DOLCE-aligned ${}^s\mathcal{O}$ ontologies, the object property dolce:has-quality has been referenced the most at 308 times, followed by dolce:has-quale 180 times, which are used for relating an endurant (e.g., apple) to a property (e.g., colour) and a value (e.g., red). Hence, domain ontologies linked to DOLCE heavily use DOLCE's features for representing properties and values. While there is some support for representing properties and values in BFO and GFO, they are not represented in the same way. BFO does not have any object properties, so while properties are supported using bfo:quality, there is no object property to link together an entity and its property. GFO does have a gfo:has_value and a gfo:value_of that correspond to those DOLCE entities 'in spirit', but this is not asserted in the corresponding mapping file due to conflicting domain and range axioms that would result in an unsatisfiable ontology. Other DOLCE entities that have been referenced many times include dolce:inherent-in, and dolce:abstract-region. For BFO-interchanged ontologies, the bfo:Role entity has been used the most, at 72 times; perhaps the results could be improved if we consider interchanging these ontologies using the Functional-Participation module of DOLCE that covers roles. Other frequently used BFO entities include bfo:Continuant, and bfo:Site. It might appear that bfo:Continuant could be mapped to the dolce:Endurant and gfo:Presential. This is not the case: bfo:Continuant subsumes bfo:quality, and dolce:quality is disjoint from dolce:Endurant so it would result in an inconsistency in the ${}^t\mathcal{O}$ if we did. It causes other inconsistencies when bfo:Continuant is mapped to gfo:Presential.

Recall that the raw interchangeability measures the amount of the domain entities that have been interchanged using equivalence mappings (see Table 1). Given the set of satisfiable equivalence mappings—7 for DOLCE to BFO, 10 for

Table 2. The number of times (N) a source foundational ontology entity is referenced in target ontologies for the total set of interchanged ontologies.

DOLCE entity	N	BFO entity	N	GFO entity	N
has-quality	308	Role	72	plays_role	80
has-quale	180	Continuant	36	part_of	52
inherent-in	88	Site	30	has_participant	47
abstract-region	60	Function	19	has_part	34
non-physical-endurant	28	ProcessualEntity	18	has_property	34
particular	26	ObjectAggregate	17	on_level	32
non-physical-object	20	FiatObjectPart	8	played_by	30
mediated-relation	18	RealizableEntity	6	Biological_level	28
mediated-relation-i	14	GenericallyDependent Continuant	6	has_role	20
part	14	Disposition	4	instance_of	20
other DOLCE entities (aggregated)	170	other BFO entities (aggregated)	13	other GFO entities (aggregated)	124

BFO to GFO, and 15 for GFO to DOLCE [3]—it is no surprise that the average raw interchangeability for the source ontologies is only 36.18%. The set of BFO ontologies had the highest raw interchangeability (62.90%), followed by DOLCE (30.02%) and lastly GFO (15.64%). BFO has the highest raw interchangeability probably because it is a bare taxonomy with no entity axioms (other than disjointness axioms) and no object properties. The entities of DOLCE and GFO have many axioms that cause dependencies between entities, therefore if a domain entity is related to a foundational ontology entity, other foundational ontology entities are also affected.

In general, the raw interchangeability differs greatly for the target ontologies which is due to two counterweighting factors. First, the number of links between the ${}^s\mathcal{O}_d$ and ${}^s\mathcal{O}_f$ has a moderate negative correlation with the raw interchangeability for DOLCE and GFO; see Table 1. Thus, a larger number of links between ${}^s\mathcal{O}_d$ and ${}^s\mathcal{O}_f$ entities for DOLCE and GFO ontologies can cause a *lower* raw interchangeability values (for the set of BFO ${}^s\mathcal{O}_f$, the correlation is much weaker). Second, the raw interchangeability is slightly *higher* when there are more mappings between source and target foundational ontologies among the interchanged ones: there are more DOLCE to GFO mappings (15) than DOLCE to BFO mappings (7), and the average interchangeability for the test ontologies are 33.29% and 26.99%, respectively. The same pattern exists for BFO to DOLCE vs BFO to GFO (58.62% vs 67.23%) and for GFO to BFO vs GFO to DOLCE (13.10% vs 18.19%). This does not hold for their aggregates, though, where the effect is dampened due to the large variation in raw interchangeability. The 'low' raw interchangeability values reveals that foundational ontology coverage and entity representation differs considerably. In some cases, there is no corresponding entity to interchange to while at other times there are seemingly similar entities to map to (recall property and value treatment in the ontologies), but the entity definition differs such that they cannot be mapped.

Computing a 'transitive interchangeability' is a moot point, as the raw interchangeability is already substantially less than 100%. Besides the extra domain entities from the base cases, this is exacerbated when the ${}^s\mathcal{O}$ does not import the

Table 3. The BioTop and Stuff ${}^{t}\mathcal{O}$ ontology metrics regarding interchangeability and change in cross comparison.

Target ontology (${}^{t}\mathcal{O}$), with the last component of the name the ${}^{t}\mathcal{O}_f$	**Raw interchangeability**	**New entities**	**Modified entities**	**Additional mappings**
biotop-bfo-ro-dolce.owl	41.18%	25	60	4
biotop-dolce-ro-bfo.owl	27.59%	67	54	6
stuff-bfo-dolce.owl	80.00%	12	17	1
stuff-dolcelite-bfo.owl	45.45%	8	14	3

entire ${}^{s}\mathcal{O}_f$. For instance, DMOP has only a subset of DOLCE, so interchanging to BFO, resulting in `dmop-bfo.owl`, and then back, resulting in `dmop-bfo-dolce.owl`, SUGOI includes the entire DOLCE foundational ontology in the second interchange, causing it to have more axioms than the original DMOP ${}^{s}\mathcal{O}$.

Comparing SUGOI to manual mappings Lastly, we evaluate SUGOI's interchangeability with the BFO and DOLCE versions of BioTop and Stuff to compare the existing manual mappings with the automatically generated ones.

The ontologies were interchanged in both directions (the BFO versions to DOLCE and vv.), and the raw interchangeability measure, and other metrics for the ${}^{t}\mathcal{O}$ ontologies are displayed in Table 3, which are in the same range as with the other ontologies (cf. Table 1). The interchangeability measure for the BioTop ontologies stems from the different coverage in the two foundational ontologies. For instance, in the original DOLCE-aligned version, biotop:physical boundary $\sqsubseteq$ dolce:feature, while in the original BFO-aligned version, it is not directly subsumed by a BFO entity. This also means that the manual versions of BioTop will differ from the interchanged versions. Comparing the interchanged versions of BioTop (e.g., `biotop-bfo-ro-dolce.owl`) to the original manual versions (e.g., `biotop-dolce.owl`), we note that there are some new and modified entities, and additional ${}^{t}\mathcal{O}_d$ to ${}^{t}\mathcal{O}_f$ subsumption axioms identified by SUGOI. One of the additional links in `biotop-dolce-ro-bfo.owl` is, biotop:ImmaterialObject $\sqsubseteq$ bfo:MaterialEntity, which is a consequence of biotop:ImmaterialObject $\sqsubseteq$ dolce-physical-endurant in the original `biotop-dolce-ro.owl`, and there is a new subsumption biotop:ValueRegion $\sqsubseteq$ dolce:endurant in `biotop-bfo-ro-dolce.owl`, which is also due to the ${}^{s}\mathcal{O}$, for biotop:ValueRegion $\sqsubseteq$ bfo:IndependentContinuant was asserted in the original `biotop-bfo-ro.owl`.

For the cross comparison of the Stuff ontologies, there are new and modified entities, and additional mapping axioms in the Stuff ${}^{t}\mathcal{O}$ ontologies. One of the additional links in `stuff-dolcelite-bfo.owl` is, stuff:Endurant $\equiv$ bfo:IndependentContinuant, while in `stuff-bfo-dolce.owl`, there is, stuff:Perdurant $\equiv$ dolce:process (a consequence of stuff:perdurant $\equiv$ bfo:process in the `stuff-bfo.owl`).

Overall, the Stuff ontology performed better in terms of raw interchangeability than BioTop, and compares well to the manual effort. However, the importance of using SUGOI for interchangeability in both ontologies is demonstrated by the fact that there were some missing mappings from the manual ontologies. Thus, it is best to use SUGOI in conjunction with manual interchange to ensure that all the relevant mappings have been implemented.

4 Discussion

Considering the results together, the average raw interchangeability for all the target ontologies is 36.18% (ranging between 2.04% to 81.81%), which means there are typically more links thanks to subsumption rather than equivalence. This is due to the fact that the set of equivalence mappings among the foundational ontologies is limited, and in some cases, those non-mapped entities from the ${}^{s}\mathcal{O}_f$ are heavily used in the alignment of the ${}^{s}\mathcal{O}_d$ to the ${}^{t}\mathcal{O}_f$, as seen by dolce:has-quality (Table 2). Foundational ontology developers may wish to add those entities to broaden the foundational ontology's coverage and therewith increase its interoperability. For the time being, it means that domain ontology developers should choose a foundational ontology carefully.

Interchangeability surely can be performed, and the subsumption mappings added by SUGOI improve the quality of the ${}^{t}\mathcal{O}$ in that extra domain entities are subsumed by the relevant ${}^{t}\mathcal{O}_f$ entities, resulting in a 'clean' taxonomy, i.e., entities that cannot be mapped via equivalence are not by default mapped as subclasses of owl:Thing or topObjectProperty outside the scope of the ${}^{t}\mathcal{O}_f$.

The interchanged ontologies are usable and SUGOI can be used as an initial tool used to achieve semantic interoperability with regards to foundational ontologies. The best results (higher raw interchangeability) were obtained for DOLCE ontologies when interchanging to GFO, for BFO ontologies when interchanging to GFO, and for GFO ontologies when interchanging to DOLCE.

We now return to the questions posed in the introduction. Regarding question 3: it is indeed feasible to automatically generate links between a domain ontology and a different foundational ontology, although the results based on equivalence-only mappings depend on the source ontology and its amount of links to its ${}^{s}\mathcal{O}_f$. Permitting subsumption, then the whole ontology can be interchanged to another foundational ontology. Regarding question 4: the issues observed are due to a combination of varying foundational ontology coverage (notably quality properties and roles), the amount of mappings between foundational ontologies, and the amount of links between the domain and foundational ontology components of the source ontology. The former problem could be solved with foundational ontology developers extending the coverage of their ontologies. The latter is more complex and requires a deep semantic change and unification about entity representation among ontology developers.

5 Conclusion

We presented the SUGOI tool, which automatically changes a source ontology's foundational ontology to another, maintaining alignments between the domain ontology component and the chosen foundational ontology (either DOLCE, BFO, or GFO). This automation enabled an investigation into the feasibility of aligning automatically one's ontology to another foundational ontology. The success of such a 'swap' based only on equivalence among entities in foundational ontologies differs by source ontology, ranging from 2 to 82% success, and averaging at

36% for the 16 ontologies included in the evaluation. Comparing SUGOI to manual dual mappings, it did outperform manual efforts, in the sense of having found additional alignments, but also missed a few, thus a final manual check is advisable. The large differences in interchangeability success are due mainly to differences in coverage of the foundational ontology (notably: qualities and roles), the number of alignment axioms between the source domain and foundational ontology, and to a lesser extent also the amount of mappings between each pair of foundational ontologies. SUGOI also uses subsumption mappings so that every domain ontology can be interchanged, preserving the structure of the ontology.

For future work, we consider creating mappings between other foundational ontologies and the existing ontologies in SUGOI. The community could also assist with this by submitting mappings in ROMULUS's community page. Given the insights on usage of a foundational ontology's content coverage and domain to foundational ontology mappings, we also plan to extend ONSET [5] with such fine-grained aspects.

References

1. Herre, H.: General Formal Ontology (GFO): A foundational ontology for conceptual modelling. In: Theory and Applications of Ontology: Computer Applications, chap. 14, pp. 297–345. Springer, Heidelberg (2010)
2. Keet, C.M., Lawrynowicz, A., d'Amato, C., Hilario, M.: Modeling Issues, Choices in the Data Mining OPtimization Ontology. In: Proceedings of the 10th International Workshop on OWL: Experiences and Directions (OWLED 2013). CEUR Workshop Proceedings, vol. 1080. CEUR-WS.org (2013), Montpellier, France, May 26-27
3. Khan, Z., Keet, C.M.: Addressing issues in foundational ontology mediation. In: 5th International Conference on Knowledge Engineering and Ontology Development (KEOD'13). pp. 5–16. SCITEPRESS – Science and Technology Publications (2013), Vilamoura, Portugal, 19-22 September
4. Khan, Z., Keet, C.M.: The foundational ontology library ROMULUS. In: 3rd International Conference on Model & Data Engineering (MEDI'13). LNCS, vol. 8216, pp. 200–211. Springer (2013), September 25-27, Amantea, Calabria, Italy
5. Khan, Z., Keet, C.M.: ONSET: Automated foundational ontology selection and explanation. In: ten Teije, A., et al. (eds.) 18th International Conference on Knowledge Engineering and Knowledge Management (EKAW'12). LNAI, vol. 7603, pp. 237–251. Springer (2012), Oct 8-12, Galway, Ireland
6. Khan, Z., Keet, C.M.: Toward semantic interoperability with aligned foundational ontologies in ROMULUS. In: Seventh International Conference on Knowledge Capture (K-CAP'13). ACM proceedings (2013), 23-26 June, Banff, Canada. (poster/demo)
7. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A.: Ontology library. WonderWeb Deliverable D18 (ver. 1.0, 31-12-2003). (2003), http://wonderweb.semanticweb.org
8. Mizoguchi, R.: YAMATO: Yet Another More Advanced Top-level Ontology. In: Proceedings of the Sixth Australasian Ontology Workshop. pp. 1–16. Conferences in Research and Practice in Information (2010), Sydney : ACS
9. Niles, I., Pease, A.: Towards a standard upper ontology. In: Welty, C., Smith, B. (eds.) Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001) (2001), Ogunquit, Maine, October 17-19, 2001