*Formal Ontology in Information Systems*
*P. Garbacz and O. Kutz (Eds.)*
*IOS Press, 2014*


*doi:10.3233/978-1-61499-438-1-227*

# Applying the Realism-Based Ontology-Versioning Method for Tracking Changes in the Basic Formal Ontology

Selja SEPPÄLÄ[a,1], Barry SMITH[a] and Werner CEUSTERS[b]
[a] *Department of Philosophy, University at Buffalo*
[b] *Center of Excellence in Bioinformatics and Life Sciences, Institute for Healthcare Informatics*, and *Department of Psychiatry, University at Buffalo*

**Abstract.** Changes in an upper level ontology have obvious consequences for the domain ontologies that use it at lower levels. It is therefore crucial to document the changes made between successive versions of ontologies of this kind. We describe and apply a method for tracking, explaining and measuring changes between successive versions of upper level ontologies such as the *Basic Formal Ontology* (BFO). The proposed change-tracking method extends earlier work on *Realism-Based Ontology Versioning* (RBOV) and *Evolutionary Terminology Auditing* (ETA). We describe here the application of this evaluation method to changes between BFO 1.0, BFO 1.1, and BFO 2.0. We discuss the issues raised by this application and describe the extensions which we added to the original evaluation schema in order to account for changes in this type of ontology. The results of our study show that BFO has undergone eight types of changes that can be systematically explained by the extended evaluation schema. Finally, we discuss problematic cases, possible pitfalls and certain limits of our study that we propose to address in future work.

**Keywords.** Ontology versioning, change tracking, quality assessment, evaluation metrics, upper-level ontologies, realist ontology

## Introduction

Numerous domain ontologies use the *Basic Formal Ontology* (BFO) as an upper level reference ontology. BFO is a realist, formal and domain-neutral upper level ontology that is designed to represent at a very high level of generality the types of entities that exist in the world and the relations that hold between them [1-3]. BFO is intentionally very small, since it is intended to provide only the most basic building blocks for the construction of domain-specific ontologies at lower levels. Briefly, it provides a starting point for logical descriptions (formulated through the statement of necessary

[1] Corresponding Author. Department of Philosophy, University at Buffalo, 135 Park Hall, 14260-4150 Amherst, NY, USA. Email: seljamar@buffalo.edu.

and jointly sufficient conditions) of the types of entities in a specific domain. Because of this common starting point, the domain ontologies using BFO appropriately are to a degree interoperable.

Ontologies are continuously changing and so is BFO [3]. As more detailed theories are proposed to aid the understanding of the types of entities in a given domain, the representation of these types in the corresponding domain ontology needs to evolve also, and this can create new demands on the overarching upper ontology. Changes in ontologies arise in some cases because there are changes in reality itself (for example when new pathogens evolve, or new types of information artifacts are created), but they can also because of changes in scientists' capabilities and preoccupations. In recent years, BFO has in fact undergone changes of the latter kind: the widely adopted BFO 1.1 has been updated to a newer version, BFO 2.0.

To remain compatible with future ontologies based on BFO 2.0, existing ontologies based on earlier versions of BFO must be updated accordingly. Information about the changes undergone by BFO is provided in [3], and in the *BFOConvert* tool which has been developed by the He Group and which provides a mapping between previous successive BFO versions.[2] These resources are however too limited to allow for a full understanding of the impact these changes will have on the domain ontologies they will affect. It is, for example, not sufficient to know that some entity types have been deleted or some others added; it is also, for reasons discussed in detail in [4, 5], necessary to know *why* these changes were made. As noted in [5], little information is provided as to the reasons for such changes which are typically made in reflection of problems caused with the older version and known only to those closely involved in the development of the ontology.

In this article, we present an improved method for tracking, explaining and measuring changes between subsequent versions of realist upper-level ontologies and related artifacts. While the method has been developed and tested in relation to realist ontologies, it can be applied also to concept-based ontologies, when the latter are conceived along the lines outlined in [1, section 7]. The method extends the realism-based ontology versioning strategy originally proposed in [5] and subsequently named *Evolutionary Terminology Auditing* (ETA) as applied to the evaluation of the *Gene Ontology* (GO) and of *SNOMED CT* [6, 7]. We focus here on the application of ETA to the *Basic Formal Ontology*[3] as a vehicle for explaining the changes undergone by BFO in a form which will be useful to the authors of domain ontologies developed using BFO 1.0 or BFO 1.1 who want to make them conformant to BFO 2.0. This study will also show some limits of ETA as originally conceived, as well as illuminate the problems users of new versions of BFO encounter because BFO itself does not yet use the ETA method. We discuss the problems that arise and propose solutions.

The following abbreviations used throughout the article are listed here for quick reference: RE=representational element; RU=representational unit; RC=representational configuration; POR=portion of reality.

---

[2] BFOConvert: http://bfoconvert.hegroup.org, last viewed on February 5, 2014.

[3] Contribution of the authors: the first author, external to the development of BFO, realized the main work and the adaptation of the ETA schema to the evaluation of BFO; the second author was the main BFO-expert who helped clarifying the changes of BFO over time; the last author contributed to the application and extension of the ETA method.

## 1. Realism-Based Ontology Versioning

Ontology versioning methods allow tracking changes between versions of an ontology. Differences are generally computed through automatic mapping thereby generating a log quantifying the number of entity types that were “added, deleted, split or merged” [4]. Standardly, however, the log does not give any explanation as to why the corresponding changes were made and to what extent they impact the overall quality of the ontology. To overcome this limitation, the qualitative ETA versioning method provides explanations of the nature of such changes in such a way that quantitative quality measures can be associated with the changes in question.

### *1.1. Principles*

According to the ETA method, changes in successive versions of an ontology or similar resource can be characterized as involving five types of errors [4]: (1) **assertion errors:** the previous version erroneously asserted the existence of some portion of reality (POR); (2) **relevance errors:** the previous version erroneously considered some POR to be objectively relevant to the purposes of the ontology; (3) **omission errors:** a relevant POR failed to be represented; (4) **encoding errors:** some term in the previous version failed to refer to the intended POR due to encoding errors, such as spelling mistakes; (5) **redundancy errors:** two or more distinct terms in a previous version referred to the same POR.

### *1.2. Original ETA Coding Schema*

The representational elements (REs) in an ontology are either representational units (RUs), for example single terms, or representational configurations (RCs), for example asserted relationships involving multiple terms. To keep track of changes between two successive versions, ETA describes how to tag each RE in the earlier version as a *match* or *mismatch*[4] with the corresponding POR (in the simplest case) by taking the later version as gold standard. The ETA coding schema recognizes seventeen configurations based on the five types of errors distinguished above. The errors depend primarily on whether an RE denoting a POR is *present in* (P) or *absent from* (A) the ontology. Such presence and absence can further be either *justified* (cases P+ and A+), or *unjustified* (P– and A–).

Each configuration depends on two factors and a number of sub-factors. The schema considers the following distinctions: (1) the level of *reality*: whether a POR exists objectively (OE, 2)[5], independently of our perception or understanding thereof, and whether an existing POR is objectively relevant to the purpose of the ontology (OR, 3), and (2) the level of *representation*, which distinguishes between:

- The *ontology authors' belief* in the *existence* (BE, 4) and in the *relevance* (BR, 5) of the represented POR.
- The *encoding* itself i.e. the RE as it appears in the ontology, which can be *intended* or not (IE, 6), for example when it shows an unintended typographic

---

[4] For a detailed description of the encoding schema, see for example Ceusters, 2011, 22–23. For its application, see Ceusters, 2011, 24–27.

[5] The parentheses specify the abbreviation used in Table 1 together with its associated column number.

error (value ‘N’ in column IE (6) in Table 1). The encoding is attributed a type of reference (TR, 7): either correct (R+) or incorrect, either because (a) the encoding does not refer (¬R) or (b) it does refer, but to a POR other than the one which was intended (R–) [4]. Finally, an RE may denote an intended and objectively existing POR that, however, is already denoted by another RE in the terminology (R++). [4]

The 17 configurations consist of one corresponding to a justified presence (P+1), two to a justified absence (A+<1-2>), ten to an unjustified presence (P–<1-10>), and four to an unjustified absence (A–<1-4>). Table 1 shows the details for each configuration.

**Table 1.** Original ETA configurations (from [4]), OE: objective existence; OR: objective relevance; BE: belief in existence; BR: belief in relevance; IE: intended encoding; TR: type of reference; ME: magnitude of error; Y: yes; N: no; na: not applicable; R+: correctly refers; R++: correctly refers, but redundantly; ¬R: does not refer; R–: refers to another POR.

| configuration | reality | | representation | | | | ME |
|---|---|---|---|---|---|---|---|
| | | | authors' belief | | encoding | | |
| | OE | OR | BE | BR | IE | TR | |
| (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| P+1 | Y | Y | Y | Y | Y | R+ | 0 |
| A+1 | N | na | N | na | na | na | 0 |
| A+2 | Y | N | Y | N | na | na | 0 |
| P-1 | N | na | Y | Y | Y | ¬R | 3 |
| P-2 | N | na | Y | Y | N | ¬R | 4 |
| P-3 | N | na | Y | Y | N | R– | 5 |
| P-4 | Y | Y | Y | Y | N | ¬R | 1 |
| P-5 | Y | Y | Y | Y | N | R- | 2 |
| P-6 | Y | N | Y | Y | Y | R+ | 1 |
| P-7 | Y | N | Y | Y | N | ¬R | 2 |
| P-8 | Y | N | Y | Y | N | R– | 3 |
| P-9 | Y | Y | Y | Y | Y | R++ | 1 |
| P-10 | Y | N | Y | Y | Y | R++ | 2 |
| A-1 | Y | Y | Y | N | na | na | 1 |
| A-2 | Y | Y | N | na | na | na | 1 |
| A-3 | N | na | Y | N | na | na | 1 |
| A-4 | Y | N | N | na | na | na | 1 |

The ideal configurations (those having zero errors) are: P+1, A+1, and A+2. All other configurations are assigned a score (> 0) corresponding to the magnitude of the error (ME, 8) it involves [4]. This score is calculated by considering the number of values in columns (4) to (7) that differ from the ideal configurations P+1, A+1, and A+2. The pertinent ideal configuration for each P– and A– configuration depends on the values in columns (2) and (3). Thus, for example, configurations P–4, P–5, P–9, A–1, and A–2 are scored with respect to P+1 [7]. If the value ‘na’ (not applicable) appears in the P– and A– rows this counts as zero, and the value R– in column (7) counts double [7]. Finally, the overall quality score of the different versions of the ontology is calculated by applying the mathematical formula in Eq. (1), where “$e_i$ stands for the magnitude of the error (if any) for a given corresponding RU, $n$ for the number of RUs [evaluated] and $m$ for the number of RUs unjustifiably absent.” [7]:

$$\frac{\sum_{i=1}^{n}(5-e_i)}{5n+4m} \quad (1)$$

### *1.3. Application of the Method*

A first step is to list and compare all the REs in the successive versions of the evaluated ontology. REs include (i) RUs, which denote an entity type and (ii) RCs, which denote representational configurations via 'entity type+relation+relatum' triples, as in 'process is_a occurrent', and which are used in the characterization of an entity type. A second step is to mark these REs with a configuration label (P or A) from the ETA schema, and to score them accordingly (see Table 5). The last step is to use the formula in Eq. (1) to generate the quality score. Creating the list of REs involves carrying out an analysis of the RCs for each entity type as represented in the versions being compared. This can be performed partially automatically, for example by extracting the RCs from the successive OWL files corresponding to each version. The extracted data then needs to be supplemented by manual examination of the natural-language specifications of the upper-level ontology and — in the case of BFO — of the definitions formulated in first-order logic.

### *1.4. Measuring the Changes*

ETA originally [4] proposed two ways of scoring the overall quality of ontologies. The first one consists in using reality as benchmark which allows assessing, for example, of how well given ontologies conform to the reality which they claim to represent [5]. The second one, which we will focus on here, consists in using the successive versions of the same ontology to measure its improvement in time. The latter amounts to treating the latest version as a correct representation of reality (gold standard) against which the previous versions must be evaluated. When assessing different versions of the same ontology at different times, the scores must be recalculated at each time *t* with respect to whatever is at *t* the latest version (see Table 5). For instance, if we have three versions of an ontology, the score for the first version (*v1*) is calculated at *t1* (the time when it was the only existing version), at *t2* with respect to the second version, and at *t3* with respect to the third version. Accordingly, the second version (*v2*) is evaluated at *t2* with respect to what PORs are acknowledged by the authors and their beliefs about those PORs at *t2*, and at *t3* with respect to the third version. The third version (*v3*) is only evaluated at *t3*. Thus, *v1* at *t1*, *v2* at *t2*, and *v3* at *t3* all have the maximal score of 1 according to the formula (Eq. (1)) above.

## 2. Applying the ETA Method to BFO

To apply the method to the BFO case would involve extracting all representational elements (REs) from the BFO 1.0, BFO 1.1, and BFO 2.0 OWL files. We focused however only on the representational units, and on a subset of the asserted and inferred RCs consisting of those involving *is_a*. A complication was that BFO does not separate its ontological component (the taxonomy backbone) from the terminological component (the terms in the taxonomy). Terminological changes have indeed occurred for different reasons and this caused problems for the evaluation since it was not possible to assume that the absence of a *term* in a later version as compared to some previous version reflected any change in the ontology, since it may have been that another name might have been used to represent the same entity type. This required a disambiguation effort consisting in identifying those cases for which the same terms do or do not denote the same entities over different versions. To ensure that we did not erroneously identify some terminological change as an ontological one, we marked each RE with a unique identifier (ID) that allows ignoring any change at the level of linguistic encoding. We marked each BFO category with an ID starting with 'RU' for 'representational unit', followed by a number, and each RC with an ID starting with 'RC' followed by the same number as the category of which it is an RC, a dot, and a unique number for each RC type. We checked the ID attribution by referring to the *BFOConvert* mapping and to the BFO specifications ([8] for BFO 1.0, and [3] for BFO 2.0), and by consulting with the authors of BFO. Table 2 shows the disambiguated RUs in BFO (the disambiguation also applies to the RCs that include them, as illustrated in RC7.1 and RC7.2[6]). The first row, for example, shows that what was earlier called 'processual entity' has been renamed 'process' in BFO 2.0.

**Table 2.** Assigning an ID to the REs. RU: representational unit; RC: representational configuration

| | **BFO 1.0/BFO 1.1** | **BFO 2.0** |
|---|---|---|
| RU7 | processual entity | BFO2-process/BFO1-processual entity |
| RC7.1 | processual entity is_a entity | BFO2-process/BFO1-processual entity is_a entity |
| RC7.2 | processual entity is_a occurrent | BFO2-process/BFO1-processual entity is_a occurrent |
| RU16 | temporal instant | zero-dimensional temporal region |
| RU17 | temporal interval | one-dimensional temporal region |
| RU31 | zero dimensional region | zero-dimensional spatial region |
| RU32 | one dimensional region | one-dimensional spatial region |
| RU33 | two dimensional region | two-dimensional spatial region |
| RU34 | three dimensional region | three-dimensional spatial region |

Some results of the RE extraction and disambiguation step are shown in columns (0) and (1) of Table 3, which lists the first three representational elements encountered when computing the transitive closure of the backbone hierarchy, as well as other REs discussed in this paper. The terms in the 'RE' column are given after the disambiguation and are the ones used in the most recent version in which the corresponding entity is denoted.

### *2.1. Determining for each Representational Element the Applicable ETA Configuration*

To partially automate this step, we defined the following principles to be used as rules of thumb, motivated by the nature of BFO as a realist upper level ontology and presupposing the disambiguation carried out in the previous step:

**Principle of Consistency with Established Science (PCES):** *the latest version of BFO is (aiming to be) the version that is most faithful to reality*. [1]

[6] For reasons of space, we do not report all the disambiguated RCs falling under each representational unit.

**Reference Ontology Principle (ROP):** *a reference ontology should cover the terminological content of the settled portions of a given scientific discipline. It should include only general terms, which are assumed to denote corresponding universals in reality and assertions of certain relations between the instances thereof.* [1]

**Principle of Obsoletion (PO):** *if a term in an ontology fails in designation, then it must be obsoleted at the first opportunity*. [1]

**Principle of Inertia of Existence (PIE):** *entities represented by an RE in the latest version of BFO can be assumed to have always existed, to exist now, and to exist always in the future*, *as parts of reality*.

**Principle of Inertia of Relevance (PIR):** *entities marked as objectively relevant in the latest version of BFO can be assumed to have been relevant throughout their entire existence*.

Some of these principles can be applied individually; others must be used in combination. According to the first principle (**PCES**) above, BFO 2.0 is to be considered the reference version against which older versions are to be compared. This principle, combined with **PIR**, allows us to set the values of the OE and OR columns (for 'objective existence' and 'objective relevance') in all versions to the value they have in BFO 2.0. Moreover, whenever an RE is present in the latest version, **PCES** allows us to set the values in columns OE and OR to 'Y', and assign the configuration type 'P+1' to that RE in the last version. Whenever an RE is present in the last version and in some previous version of BFO, **PIE** allows us to set the value by default to 'Y' in the 'belief in existence' (BE) column of that previous version and the type of reference as correct ('R+' in the TR column), and **PIR** allows us to set the value by default to 'Y' in the 'belief in relevance' (BR) column of that previous version. Thus, **PCES**, **PIE**, and **PIR** jointly allow us to determine all 'P+1' configurations in all versions of BFO whenever the RE is present in the last version and some previous one. There is no automatic way to determine the other configurations. However, **PCES** and **PO** jointly allow us to predict that deleted REs will be marked as unjustified presences (P–) in previous versions and as justified absences (A+) in later versions. Moreover, **PCES**, **PIE**, and **PIR** together allow us to predict that REs added in a subsequent version will be marked as unjustified absences (A–) in previous versions and justified presences (P+) in the subsequent ones. One exception to these predictions in BFO concerns two cases that present an 'A+ P– A+' configuration pattern. We discuss this case in section 4.4.

### *2.2. Extended ETA Method*

To determine which particular type of A+, A–, and P– configuration applies in the current and previous versions of BFO, we examined all REs individually and assigned them the values in columns (2) to (7) that best explain the respective changes according to the explanations given in the specifications and by the authors of BFO.

**Table 3.** Sample of the versioning of BFO at time *t3*. ID: unique identifier; RE: representational element; OE: objective existence; OR: objective relevance; BE: belief in existence; BR: belief in relevance; IE: intended encoding; TR: type of reference; ME: magnitude of error; CvX: configuration that applies to the version X; Y: yes; N: no; NC: not considered; na: not applicable; R+: correctly refers; ¬R: does not refer; R–: refers to another POR; Ra: refers ambiguously; P+: justified presence; A+: justified absence; P–: unjustified presence; A–: unjustified absence.

| | | ALL | | BFO 1.0 | | | | | | BFO1.1 | | | | | | BFO2.0 | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| RUID | RE | OE | OR | BE | BR | IE | TR | ME | Cv1 | BE | BR | IE | TR | ME | Cv2 | BE | BR | IE | TR | ME | Cv3 |
| (0) | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) | (13) | (14) | (15) | (16) | (17) | (18) | (19) | (20) | (21) |
| RU1 | entity | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU2 | continuant | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC2.1 | RU2 is_a RU1 | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU3 | occurrent | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC3.1 | RU3 is_a RU1 | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU5 | generically dependent continuant | Y | Y | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC5.2 | RU5 is_a RU36 | N | na | NC | na | na | na | 0 | A+4 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RU6 | specifically dependent continuant | Y | Y | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC6.2 | RU6 is_a RU36 | N | na | NC | na | na | na | 0 | A+4 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RU7 | BFO2-process/BFO1-processual entity | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC8.3 | process boundary is_a RU7 | N | na | Y | Y | Y | ¬R | 3 | P-1 | Y | Y | Y | ¬R | 3 | P-1 | N | na | na | na | 0 | A+1 |
| RU11 | material entity | Y | Y | Y | N | na | na | 1 | A-1 | Y | N | na | na | 1 | A-1 | Y | Y | Y | R+ | 0 | P+1 |
| RU12 | immaterial entity | Y | Y | Y | N | na | na | 1 | A-1 | Y | N | na | na | 1 | A-1 | Y | Y | Y | R+ | 0 | P+1 |
| RU14 | realizable entity | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC14.2 | RU14 is_a RU36 | N | na | Y | Y | Y | Ra | 4 | P-12 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RC14.4 | RU14 is_a RU6 | Y | Y | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU15 | process profile | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RU22 | continuant fiat boundary | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RC22.3 | RU22 is_a RU12 | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RU23 | spatial region | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU26 | disposition | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC26.2 | RU26 is_a RU36 | N | na | Y | Y | Y | Ra | 4 | P-12 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RC26.5 | RU26 is_a RU6 | Y | Y | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU27 | function | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC27.2 | RU27 is_a RU36 | N | na | Y | Y | Y | Ra | 4 | P-12 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RC27.3 | RU27 is_a RU26 | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RC27.5 | RU27 is_a RU14 | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC27.6 | RU27 is_a RU6 | Y | Y | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU31 | zero-dimensional spatial region | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RC31.3 | RU31 is_a RU12 | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RC31.5 | RU31 is_a RU23 | Y | Y | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 | Y | Y | Y | R+ | 0 | P+1 |
| RU35 | history | Y | Y | NC | na | na | na | 1 | A-5 | NC | na | na | na | 1 | A-5 | Y | Y | Y | R+ | 0 | P+1 |
| RU36 | dependent continuant | N | na | Y | Y | Y | Ra | 4 | P-12 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RU38 | scattered spatiotemporal region | Y | N | Y | Y | Y | R+ | 1 | P-6 | Y | Y | Y | R+ | 1 | P-6 | Y | N | na | na | 0 | A+2 |
| RU39 | connected spatiotemporal region | Y | N | Y | Y | Y | R+ | 1 | P-6 | Y | Y | Y | R+ | 1 | P-6 | Y | N | na | na | 0 | A+2 |
| RU42 | processual context | N | na | Y | Y | Y | ¬R | 3 | P-1 | Y | Y | Y | ¬R | 3 | P-1 | N | na | na | na | 0 | A+1 |
| RU43 | process aggregate | N | na | Y | Y | Y | Ra | 4 | P-12 | Y | Y | Y | Ra | 4 | P-12 | N | na | na | na | 0 | A+1 |
| RU47 | BFO1-process | Y | N | Y | Y | Y | R+ | 1 | P-6 | Y | Y | Y | R+ | 1 | P-6 | Y | N | na | na | 0 | A+2 |
| RU48 | object boundary | Y | N | Y | Y | Y | R+ | 1 | P-6 | Y | Y | Y | R+ | 1 | P-6 | Y | N | na | na | 0 | A+2 |

This thorough examination revealed some limits to the original evaluation schema. This led us to introduce new values for two columns: 'not considered' (NC) for column BE and 'refers ambiguously' (Ra) for column TR [7], which yielded five new configurations. We thus extended the existing schema as shown in Table 4. We discuss problematic cases and motivate our solutions in more detail in sections 4.3 to 4.5. REs added to BFO 2.0 make their absence in previous versions unjustified either because of (1) changes in the authors' beliefs about their relevance, or (2) because the authors had not considered the question of the existence of the represented PORs.

[7] For details and examples of ambiguously referring encodings, i.e., referring to two distinct universals, see section 4.5.

**Table 4.** New ETA configurations characterizing justified and unjustified absences and presences in different versions of an ontology. OE: objective existence; OR: objective relevance; BE: belief in existence; BR: belief in relevance; IE: intended encoding; TR: type of reference; ME: magnitude of error; Y: yes; N: no; NC: not considered; na; not applicable; Ra: refers ambiguously; A+: justified absence; P–: unjustified presence; A–: unjustified absence

| configuration | reality | | representation | | | | ME |
|---|---|---|---|---|---|---|---|
| | | | authors' belief | | encoding | | |
| | OE | OR | BE | BR | IE | TR | |
| (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| A+3 | Y | N | NC | na | na | na | 0 |
| A+4 | N | na | NC | na | na | na | 0 |
| P-11 | Y | Y | Y | Y | Y | Ra | 1 |
| P-12 | N | na | Y | Y | Y | Ra | 4 |
| A-5 | Y | Y | NC | na | na | na | 1 |

The first case applies, for instance, to the BFO terms 'material entity' (RU11) and 'immaterial entity' (RU12), since the specifications and BFO-related literature indicate that the authors believed in their existence but did not consider them relevant for inclusion. These cases correspond to A–1 configurations. The second case, i.e. A–5 configurations, applies to the terms 'generically' (RU5) and 'specifically dependent continuant' (RU6) in BFO 1.0, and to 'process profile' (RU15) and 'history' (RU35) in BFO 1.0 and BFO 1.1. It also applies to RCs such as 'function is_a disposition' (RC27.3) that correspond to changes triggered by RUs moved within the BFO-hierarchy thus yielding new *is_a* RCs.

For REs deleted from BFO 2.0, we started by assigning the now appropriate values to the OE and OR columns. Lack of documentation forced us to consult the authors of BFO, who confirmed that deleted terms such as 'dependent continuant' (RU36), 'processual context' (RU42), and 'process aggregate' (RU43) are now understood as having never denoted. The values for the corresponding OE and OR columns for these terms and their RCs are thus 'no' ('N') and 'not applicable' ('na') respectively. These values yielded A+1 configurations in the latest version, and either P–1 or P–12 in previous versions, where the corresponding PORs were believed to exist and to be relevant. P–1 applies to REs deleted for lack of reference, such as 'processual context' (RU42) and 'process boundary is_a BFO2-process/BFO1-processual entity' (RC8.3); P–12 to REs deleted because of ambiguous reference, such as 'process aggregate' (RU43), and REs that include 'dependent continuant' (e.g., RC14.2, RC26.2, RC27.2). The other deletions are explained by the lack of objective relevance of the represented universals, thus yielding A+2 configurations in the latest version and P–6 cases in BFO 1.0 and BFO 1.1.

## 3. Results

The results of the full multi-step evaluation of successive versions of BFO in Table 5 show that its quality has considerably increased since the first two versions. The qualitative improvement from BFO 1.1 to BFO 2.0 was much larger than the improvement observed from BFO 1.0 to BFO 1.1 as (this would have been) evaluated at the time when BFO 1.1 was the reference. Of course, the latest version having the maximal quality score is an artifact of taking that version as gold standard, but the essential part is the evolution over time. It might for instance have been the case that compared from the situation at BFO 2.0, BFO 1.1 turned out to be worse than BFO 1.0.

If such dips do not happen over subsequent versions, there is reason to say that the BFO authors are consistent in their approach.

**Table 5.** Results of the versioning of BFO at three successive times (*t1*, *t2*, *t3*).

| | *t1* | *t2* | | *t3* | | |
|---|---|---|---|---|---|---|
| **C** | **BFO 1.0** | **BFO 1.0** | **BFO 1.1** | **BFO 1.0** | **BFO 1.1** | **BFO 2.0** |
| **P-1** | 0 | 0 | 0 | 5 | 5 | 0 |
| **A-1** | 0 | 0 | 0 | 13 | 13 | 0 |
| **A-5** | 0 | 13 | 0 | 53 | 42 | 0 |
| **P+1** | 138 | 133 | 151 | 85 | 96 | 151 |
| **P-6** | 0 | 0 | 0 | 36 | 36 | 0 |
| **P-11** | 0 | 5 | 0 | 0 | 0 | 0 |
| **A+2** | 0 | 0 | 0 | 0 | 0 | 36 |
| **A+1** | 0 | 0 | 0 | 0 | 0 | 19 |
| **A+4** | 0 | 0 | 0 | 2 | 0 | 0 |
| **P-12** | 0 | 0 | 0 | 12 | 14 | 0 |
| **scores** | **1** | **0.92** | **1** | **0.62** | **0.66** | **1** |

The analysis of the results at time *t3* reveals eight configuration patterns that show which kinds of changes BFO has undergone throughout its three versions (Table 6).

**Table 6.** Configuration patterns characterizing changes in successive versions of BFO at time *t3*. #: number; RE: representational element; POR: portion of reality; P+: justified presence; P–: unjustified presence; A+: justified absence; A–: unjustified absence

| Configuration patterns | # of REs | Explanation. (Examples in Table 3) |
|---|---|---|
| P+1 P+1 P+1 | 81 | No changes. (RU1-RC3.1) |
| A–5 P+1 P+1 | 11 | REs that were not considered at all were introduced in newer versions. (RC22.3) |
| A–5 A–5 P+1 | 46 | REs that were not considered at all were introduced in the last version. (RU5, RC26.5) |
| A–1 A–1 P+1 | 13 | REs that were not believed to be relevant were introduced in the last version. (RU11, RU12) |
| A+4 P–12 A+1 | 2 | REs that were not considered at all were introduced in a newer version and subsequently deleted because the POR did in fact not objectively exist. (RC5.2, RC6.2) |
| P–1 P–1 A+1 | 5 | REs were deleted in the latest version because the POR does in fact not objectively exist and the REs did not refer to anything. (RU42) |
| P–12 P–12 A+1 | 12 | REs that referred ambiguously were deleted in the latest version because the POR does in fact not objectively exist. (RC14.2, RU36) |
| P–6 P–6 A+2 | 40 | REs previously considered objectively relevant are now considered not relevant. (RU47) |

## 4. Discussion

### *4.1. Dealing with Changes in Encoding/Terminology*

Although BFO's lack of separation of its ontological and terminological content does not prevent skilled users from understanding what the ontology represents, this lack is a serious shortcoming for evaluating BFO's evolution: for it means that one cannot take for granted that over different versions the same term denotes the same POR nor that different terms do not denote the some POR. For example, 'process' (in BFO 1.0, 1.1, and 2.0) and 'processual entity' (not in BFO 2.0) were particularly difficult cases. The

*BFOConvert* file mapping from BFO 1.1 to BFO 2.0 suggested that both 'processual entity' and 'process' in BFO 1.1 should correspond to BFO 2.0 'process' (hereafter 'BFO2-process') thus offering three possible explanations: (1) 'processual entity' was deleted from BFO 2.0 leaving only its child category 'process' (hereafter 'BFO1-process'), then 'BFO2-process' corresponds to 'BFO1-process'; (2) 'BFO1-process' was deleted and 'processual entity' renamed to 'BFO2-process'; (3) 'processual entity' and 'BFO1-process' were merged into 'BFO2-process'. BFO-experts report that 'processual entity' was renamed 'process', and that the entity type named 'BFO1-process' is no longer regarded objectively relevant for inclusion.

Another difficult case was that of the BFO 1.0 and BFO 1.1 'object boundary' and the BFO 2.0 'continuant fiat boundary'. Here, the introduction to the BFO 2.0 specifications provides some information about the 'new simplified treatment of boundaries and regions'. However, the reader has to be already acquainted with the issue to understand the changes and, more importantly, the explanations do not explicitly state whether we can consider both terms to refer to the same entity type. A closer look at their definitions would suggest that 'object boundary' corresponds to the new 'two-dimensional continuant fiat boundary', but a clarification in the BFO 1.0 specifications suggests that 'object boundary' refers to more than just that entity type: "*In physical entities* boundaries are *often* a closed two-dimensional surface (inner or outer) of a thing." [8] When we asked different authors of BFO about these REs, the responses diverged with respect to their identity, but not with respect to the evolution of the understanding of boundaries in general, leading to our decision to consider 'object boundary' irrelevant since the last version

We found mismatches between the specifications of the ontology and its encoding in the corresponding OWL file. For example, the specifications of BFO 1.0 have the terms 'fiat part of object' and 'boundary of object' that appear in the corresponding OWL file as 'fiat object part' and 'object boundary', respectively. Shifts of this kind should be avoided in future revisions. While we do not here report mismatches between the specifications of the ontology and its encoding in OWL, we do note that the versioning method here described can also be used for checking the consistency between the specifications of the different versions of an ontology and their OWL implementations, as well as the overall internal coherence of such specifications. Such application would allow identifying possible other types of changes that the extended ETA method should account for and, if needed, extending the method further. As we are interested only in the ontological aspects of BFO, we restricted the possible values assigned to the 'intended encoding' (IE) column to either 'Y' or 'na', assuming that the authors of BFO did not make any encoding errors at the ontological level.

### *4.2. Alternative Application of the Evaluation Method*

The evaluation of changes could have been done differently, by marking unjustified presences and absences in previous versions as justified. The rationale for this would be that, because we are dealing with an all-purpose upper level ontology, whenever the authors believe that an RE is relevant, then the types of things it represent should also be seen by them as being objectively relevant since otherwise they would not be included in the ontology. This alternative application of the evaluation method rejects the **PIR** principle and would allow for the values in the OR column to vary for the same RE from version to version. The now deleted REs that were believed by the authors of previous versions to be relevant (since they were included in BFO) would have been assigned a P+1 configuration whenever the OE column has the value 'Y'. And, similarly, the previously absent REs that represent types of things that were not believed to be relevant for inclusion in BFO would have been marked as A+2 configurations whenever OE='Y'. In both cases, the penalties for each change would have been minimal and the quality score for BFO 1.0 and BFO 1.1 relatively high (above 0.9). However, the authors of BFO consider that it is up to users to demonstrate a need for REs representing objectively existing PORs and make a request based on evidence that they are objectively relevant. It seems therefore more accurate to consider, inversely, that the belief of the authors in the relevance of an RE depends on the objective relevance of the represented entity type, as the latter depends on the scientific community's uses and needs. We thus chose to evaluate BFO according to the latter solution, i.e. accepting **PIR**, whereby the values of the BFO 2.0 OE and OR columns apply to all versions of BFO.

### *4.3. Dealing with Objective Relevance and Pragmatic Considerations*

Some of the REs such as 'scattered' (RU38) and 'connected spatiotemporal region' (RU39) lack objective relevance in BFO 2.0 because, according to Barry Smith, they "cause too many problems for people (we postpone them to BFO 3.0)". Therefore, the REs previously under 'spatiotemporal region' are now merged into a single category. These cases show that beliefs in objective relevance can also be influenced by pragmatic considerations.

### *4.4. Dealing with the Authors' Beliefs in Existence of Some Type of Thing*

Whenever an RE is included in BFO, we have to assume that BFO's authors believe in the existence of the types of things represented by that RE. However, this raises some questions. For instance, can we say that the authors of BFO 1.0 *did not believe* in the existence of the entity types represented by an RE that was introduced in BFO 1.1 or BFO 2.0? The issue here is not one of belief, but of whether a given putative type of entity is considered at all as a candidate for existence: we cannot have beliefs about something we haven't even thought about. Therefore, we cannot restrict the possible values in the 'belief in existence' (BE) column of the evaluation schema to 'Y' and 'N'. To explain why an RE was absent from older versions, we thus introduced a new value 'not considered' (NC), which yields two new 'justified absence' configurations (A+3 and A+4) as well as a new 'unjustified absence' configuration (A–5), as shown in Table 4.

A further issue related to beliefs arises whenever an RE is absent from initial version(s), introduced at some later point, and then subsequently deleted. These cases should be considered as having a general A+ P– A+ pattern, since the RE in question is absent in the latest version. Adding the new 'NC' value in the author's 'belief in existence' (BE) column allowed us to provide two distinct explanations for the justified absences of 'generically dependent continuant is_a dependent continuant' (RC5.2) and 'specifically dependent continuant is_a dependent continuant' (RC6.2) in different versions of BFO: in BFO 1.0, these REs are justifiably absent because the authors had not considered the existence of the representational configurations they represent (OE='N' and BE='NC'). Their absence from BFO2 is justified because the authors no longer believe in the existence of the entity types that these REs were believed to represent in BFO 1.1 (OE='N' and BE='N').

### 4.5. Dealing with Ambiguous Reference

In one of the cases encountered in BFO, we were not able to tell just by comparing, for example, BFO 1.0 and BFO 2.0 whether 'dependent continuant' (RU36) was deleted and two more specific categories added. It could also have been that 'dependent continuant' was merged with one of the two subcategories, most plausibly with 'specifically dependent continuant' (RU6), and the other subcategory ('generically dependent continuant', RU5) added as a sibling. If we look at BFO 1.1, which has all three categories, we can suppose that the first explanation is the correct one. However, this kind of intermediate representation might not always be available. Furthermore, the explanation of this change may be absent from the specifications of the latest version of the ontology, which is the case for 'dependent continuant' in the BFO 2.0 specifications. As it turns out, neither of these explanations was correct: the RE 'dependent continuant' was deleted, in accordance with the **ROP** principle, because it does not represent a universal, but a defined class formed by the conjunction of 'generically' and 'specifically dependent continuant'; we marked it therefore to be referring ambiguously to two distinct universals. The other solution would have been to mark it as unjustifiably present since BFO, as a realist ontology, does not allow defined classes to be part of its backbone[8] [1]. However, the first explanation (ambiguity) is more specific. Another ambiguously referring RE is 'process aggregate' (RU43), which was argued '[to be] not precise enough and to which nothing really corresponded in reality'. To account for ambiguous REs, we introduced a new value 'Ra' for 'ambiguous reference' to the possible values in the 'type of reference' (TR) column of the evaluation schema. This yielded two new 'unjustified presence' configurations (P–11 and P–12) that account for the deletion of an ambiguous RE.

## Conclusion

We improved the qualitative Evolutionary Terminology Auditing (ETA) method originally proposed by Ceusters and Smith [5] and applied it to BFO. The results of our study show that BFO has undergone changes which fall in eight distinct patterns that are explained by the ETA configurations that form these change patterns. To do so, we had to introduce five additional configurations to the original 17.

This revealed, once again [6, 7], that identifying the motivations for changes, i.e., assigning the right configuration to each element, is hard to do *a posteriori*. Examining the BFO specifications did not always prove successful. What this study emphasizes, it that for a reliable assessment of the successive versions of an ontology, the method should be applied in collaboration with its authors, and ideally used during the revision process. We therefore strongly encourage ontology authors, including authors of BFO, to systematically use the extended ETA method to make their motivations for changes explicit at the time the changes are made. Once these quality assessment tables are established for each version, the authors can systematically complement the specifications with more detailed explanations.

[8] BFO does not allow double inheritance at the ontological level. However, it can be acceptable at the terminological level, as with the expression 'dependent continuant'.

## Acknowledgements

Work on this paper was supported by the Swiss National Science Foundation (SNSF) and funded in part by grant 1R01DE021917-01A1 from the National Institute of Dental and Craniofacial Research (NIDCR). The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the NIDCR or the National Institutes of Health. Our thanks go in addition to Alan Ruttenberg for his helpful discussions and comments on the evolution of BFO, and to Jonathan Bona for his assistance with OWL-related matters.

## References

[1] Smith B, Ceusters W. Ontological Realism: A Methodology for Coordinated Evolution of Scientific Ontologies. Applied Ontology. 2010;5:139-88.
[2] Smith B. On classifying Material Entities in Basic Formal Ontology. Interdisciplinary Ontology Proceedings of the Third Interdisciplinary Ontology Meeting2012. p. 1-13.
[3] Smith B, Almeida M, Bona J, Brochhausen M, Ceusters W, Courtot M, et al. Basic Formal Ontology 2.0: DRAFT SPECIFICATION AND USER'S GUIDE. 2012.
[4] Ceusters W. Realism-Based Versioning for Biomedical Ontologies. Final report on grant 1R21LM009824-01A1 from the National Library of Medicine; 2011.
[5] Ceusters W, Smith B. A Realism-Based Approach to the Evolution of Biomedical Ontologies. AMIA Annual Symposium Proceedings2006. p. 121.
[6] Ceusters W. Applying Evolutionary Terminology Auditing to SNOMED CT. AMIA Annual Symposium Proceedings2010. p. 96.
[7] Ceusters W. Applying evolutionary terminology auditing to the Gene Ontology. Journal of Biomedical Informatics. 2009;42(3):518--29.
[8] Spear AD. Ontology for the Twenty First Century: An Introduction with Recommendations. Saarbrücken, Germany: Institute for Formal Ontology and Medical Information Science; 2006.