

In Grove, Andrew, Ed. Proceedings of the 70th Annual Meeting of the American Society for Information Science and Technology (ASIST), Milwaukee, WI (US), 2007.

Three of the Four FRBR Group 1 Entity Types are Roles, not Types

**Allen H. Renear (corresponding author), David Dubin
Graduate School of Library and Information Science,
University of Illinois at Urbana-Champaign,
501 E. Daniel St.; Champaign, IL 61820; Tel: 217-333-3280.
renear@uiuc.edu**

We examine the conceptual model of the “bibliographic universe” presented in IFLA’s *Functional Requirements for Bibliographic Records* (FRBR) and argue, applying ontology design recommendations proposed by N. Guarino and C. Welty, that three of the four Group 1 entity *types* might be more accurately conceptualized as *roles*. We show how this approach may generalize the solution to a previously identified puzzle regarding the FRBR entity type of XML documents and speculate as to the sorts of entities that might take on these roles. This view of bibliographic entities, that they are roles that other things have in particular social contexts is consistent with John Searle’s notion of a cascade of social facts established through collective intentionality. We allow that even if our analysis is correct the current FRBR approach may be preferable as there are good reasons for “denormalized ontologies” that treat roles as types, particularly when the objective is not a general ontology, but a practical conceptual model.

Introduction

We examine the conceptual model of the “bibliographic universe” presented in IFLA’s *Functional Requirements for Bibliographic Records* (FRBR) and argue, applying ontology design recommendations proposed by N. Guarino and C. Welty, that three of the four Group 1 entity *types* might be more accurately conceptualized as *roles*. We show how this approach may generalize the solution to a previously identified puzzle regarding the FRBR entity type of XML documents and speculate as to the sorts of entities that might take on these roles. This view of bibliographic entities, that they are roles that other things have in particular social contexts is consistent with John Searle’s notion of a cascade of social facts established through collective intentionality. We allow that even if our analysis is correct the current FRBR approach may be preferable as there are good reasons for “denormalized ontologies” that treat roles as types, particularly when the objective is not a general ontology, but a practical conceptual model. This is a preliminary exploration; our intention at this point is more to convene discussion than to confidently defend a particular analysis.

The FRBR Group 1 Entities

The *Functional Requirements for Bibliographic Records* (FRBR) is “a conceptual model of the bibliographic universe” developed by the International Federation of Library Associations (IFLA 1998). Although primarily aimed at the library cataloguing systems community FRBR is a compelling empirically grounded conceptual framework for cultural objects that is increasingly influential in other domains as well.

The FRBR “Group 1” entity types are *works, expressions, manifestations, and items*. A work is defined as “a distinct intellectual or artistic creation”, an expression is “the intellectual or artistic realization of a work in the form of alphanumeric, musical, or choreographic notation, sound, image, object, movement...”, a manifestation is “the physical embodiment of an expression of a work”, and an item is “a single exemplar

of a manifestation”. Using printed books as an example (which we will do throughout) these concepts correspond roughly to the common notions of work, text, edition, and physical copy, respectively. Each entity type is associated with a characteristic attributes — for instance, works have *form* (novel, play, poem, etc.), expressions may be in a particular *language*, manifestations may have a *typeface*, and items may have a *condition*. A particular work may be realized in any number of expressions (such as different translations or textual variants); an expression may be embodied in any number of different manifestations (such as different editions with different page design or carrier); and a particular manifestation may have any number of individual physical instances. Works, expressions, and manifestations are abstract objects, and items are concrete physical objects.

The following entity relationship diagram represents the relationships between the Group 1 entity types.

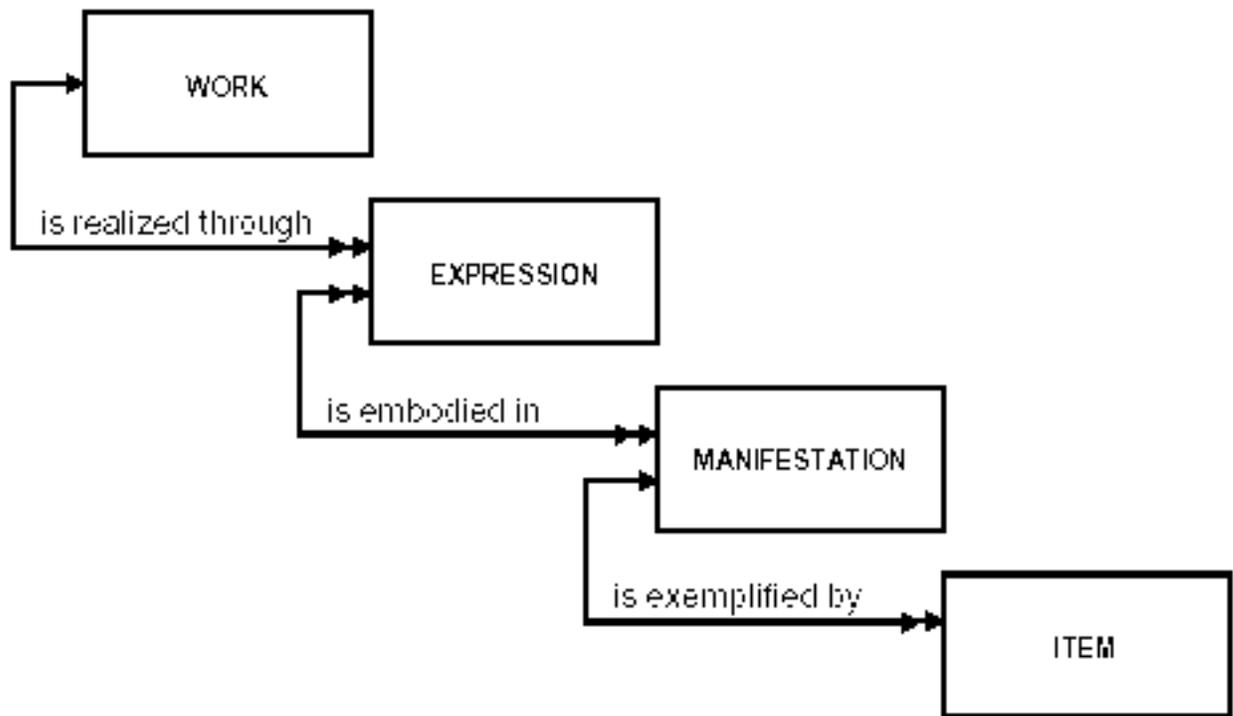


Figure 1: ER Diagram of FRBR Group 1 Entities and Primary Relationships (diagram from IFLA, 1998)

For background on FRBR see Tillett (2004) and IFLA (1998). For discussions of FRBR modeling issues relevant to this paper see Doerr, Hunter, & Lagoze (2003), Renear & Choi (2006), and Doerr and Le Boeuf (2007).

An Entity Type Assignment Puzzle

In 2003 we noticed a case where the assignment of a particular bibliographic entity to the appropriate FRBR entity type class seemed problematic. In this section we rehearse that difficulty and the preliminary solution we favored in 2003. The remainder of the paper describes an approach to FRBR that generalizes this solution.

With the widespread use of XML vocabularies (such as HTML, TEI, DITA, the NLM Journal Publishing Tag Set and many others) *XML documents* are now one of the most common objects found in the contemporary digital “bibliographic universe”. Informally an XML document is a combination of text and descriptive markup from an

XML tag set (Goldfarb 1981; Coombs et al., 1987). Formally an XML document is defined as a string that matches a particular context-free grammar and conforms to several additional non-syntactic constraints (W3C, 1999). It is natural when anticipating the application of FRBR to the digital portion of the bibliographic universe to ask “to which FRBR Group 1 entity class does an XML document belong”?

An XML document cannot be an item as it is not a concrete physical object. Nor would it seem to be a work, an intellectual creation independent of any particular symbolic realization. That leaves two remaining candidate entity classes: expression and manifestation. As something defined by a formal grammar an XML document may at first appear to be a notational entity and therefore, an expression. Moreover the XML markup itself appears to be making assertions about the text (Sperberg-McQueen et al., 2000; Renear et al., 2002; Dubin et al., 2003). Consistent with this assignment is the fact that manifestation-specific features such as typeface and carrier material are not properties of the XML document per se, even though XML documents may of course be rendered in a particular typeface and on a particular carrier.

On the other hand we also noted that from a different perspective an XML document could indeed appear to be a manifestation, with its XML markup

... functioning exactly like [manifestation-level] rendering events (font shifts, type size changes, vertical and horizontal whitespace, etc.) to effect, expedite, and disambiguate the recognition (whether by humans or computer) of the underlying textual objects. (Renear et al., 2003)

And if XML documents seem unlike manifestations in not having typeface, carrier, and such perhaps that is only because they are partial or incomplete manifestations (from a remark by Michael Sperberg-McQueen c. 2002).

We now have conflicting considerations, with one set supporting identifying an XML document as an expression and the other set supporting identification as a manifestation. In response to this puzzle we conjectured that the concept of an XML document does not by itself univocally identify something which is an appropriate candidate for identification as a FRBR entity. Rather it appears that

... a particular XML document can only be assigned to one of the two candidate FRBR entity classes — manifestation or expression — with respect to certain features of its context or intended use. Apart from such a context, XML documents have a “dual aspect”, understood from one perspective an XML document is a manifestation ... understood from another it is itself an expression (Renear et al., 2003).

For instance, in the context of original authorship an XML document is a manifestation:

... a work about, say, whales, consists in a body of abstract intentional acts: assertions about whales, questions about whales, warnings about whales, and so on, arranged in a rhetorical structure. When an author writes a book about whales, she creates an expression the meaning of which is this particular structure of intentional acts. She accomplishes this by creating an exemplar of a manifestation ... that in a reader will effect the apprehension of the expression that in turn has as its meaning the body of intentional objects which is the work about whales. Generally, such a manifestation will have, as well as the marks indicating linguistic characters, specific features that make the apprehension of the embodied expression reliable and efficient. In familiar printed books, these features are, most importantly, graphic devices such as changes in horizontal and vertical spacing, font shifts, color, and so on.

XML markup may also perform the same function as these graphic devices, making the apprehension and cognitive navigation of an expression efficient and unambiguous. Consider ... an authorial manuscript about whales, prepared in an XML element set, such as TEI Lite . The author uses markup to indicate the identities and boundaries of textual components, anticipating that subsequent processing by XML/TEI software will be able to follow these cues and process the (digital) manuscript appropriately. ... In this case, XML markup for, say, a paragraph, seems to be functioning ... like extra vertical leading or an initial em space.

What is important to note in this scenario is the obvious fact that the author is creating a book about *whales*, and not about paragraphs, titles, lists, and such. That is, the author is not using markup to describe an expression, and in particular not using markup to describe the expression that realizes a work about whales. If she were she would be creating another work, on another subject — a work not about whales, but about an *expression*. (Renear et al., 2003)

And if someone were to read the work on whales from unformatted XML source then in that situation as well the XML tags would appear to be functioning just like the graphic devices of manifestation-level “presentational markup” (Coombs et al., 1987), cuing the underlying constitutive textual components (DeRose et al., 1990).

But in another context, for instance the preparation of a scholarly edition, the same XML document appears to be an expression:

... a scholarly textual editor is preparing, using the TEI XML vocabulary, a new critical text of an important cultural work. ... The text [of the work being edited]... may well be about whales, but our scholarly editor is probably not a whaling specialist, and, in any case, is not herself creating or modifying a work about whales. The editor is rather creating a work about an expression that realizes a work about whales. To do this, the editor creates an expression that realizes the work that is about the expression that realizes the work about whales. The editor’s own creative achievement, we note, is not zoological; it is philological.

Looking at the phenomenology of the situation more closely, we imagine that the editor or transcriber puzzles over the pages of a physical document and, drawing on her erudition, her knowledge of the relevant history, graphic vocabulary, the content and context of the specific textual item, and so on, comes to a conclusion: “<p>”. Here “<p>” itself is a bit of notation that in context is clearly being used to express an assertion about a expression. The XML document in this context is functioning as an expression, but an expression that realizes a work *about an expression*.

To appreciate how the subtlety of this ambiguity has impeded its identification note that the expression created by the scholarly editor is... a second-order expression: an expression that realizes a work that is itself about an expression — and, specifically, about the very expression the XML document would embody if interpreted as a manifestation. (Renear et al., 2003).

In the authorial context, where a book about whales is being authored, the markup is being used paralinguistically to performatively realize a particular structure of linguistic objects, and is therefore a manifestation level feature of the work about whales. In the editorial context the markup is being used not to assist in making assertions about whales, but to make assertions about the text of a work about whales. An author uses “<p>” to implement paragraphing (much as we say “I promise” in order to promise), but a textual editor uses “<p>” to assert the existence of a paragraph (cf “he promised”). [For the distinction between the performative and descriptive use of identical markup and application of speech act theory to markup

analysis see Renear (2000); for more on intention in markup practices see Dubin (2003).]

This analysis of the problem still seems correct as far as it goes: an XML document can be a manifestation in some circumstances and an expression in others. But we missed drawing the more general and broadly explanatory conclusion: expressions and manifestations are not types of things, but rather roles that things, such as symbol sequences, take on in particular social contexts. Reaching that more general understanding required re-examining FRBR through the lens of recently developed ontology evaluation criteria.

Rigidity, Types, and Roles in Ontology Evaluation

Guarino and Welty have identified a number of “meta-properties” which can be used to evaluate modeling decisions and ontologies (Guarino & Welty, 2000, 2002, 2004). We are concerned here with just one of these, which they call *rigidity*, and define, using the formalism of symbolic modal logic, this way:

D1: A rigid property is essential to its instances, i.e. $\forall x \phi(x) \rightarrow \Box \phi(x)$

A prose version of this definition would be:

D2: ϕ is a *rigid* property =df for all things x , if x has ϕ then x necessarily has ϕ .

The basic idea here is that while some properties seem to be such that the things that have them might not have had them, others are such that anything that has them could not have failed to have them. The latter sort are identified as *rigid*. The notion is a refinement of the traditional concept of an *essential* or *necessary* property.

Guarino and Welty use the properties of *being a person* and *being a student* as examples of rigid and non-rigid properties respectively. Someone who is a student might, if circumstances had been different, not have been a student; for instance they might have accepted a job offer instead of enrolling in school. But it does not seem possible, at least on the ordinary understanding of the concept of person, that something which has the property of being a person might have not had the property of being a person, or could lose the property of being a person (and continue to exist). Guarino and Welty write:

...we normally think of PERSON as rigid; if x is an instance of PERSON it must be an instance of PERSON in every possible world. The STUDENT property on the other hand, is normally not rigid; we can easily imagine an entity moving in and out of the STUDENT property while being the same individual... (Guarino and Welty, 2000).

According to Guarino and Welty this difference reflects the fact that in our conceptual scheme *person* is a fundamental ontological type, whereas *student* is role that individuals things (of some type of other) have in certain contingent circumstances. They claim that the “ideal structure of a clean taxonomy” will have types (and other rigid properties) in the “backbone” and roles (and other non-rigid properties) “hanging off” the backbone.

We determine whether or not a property is a necessary property by reflecting on our intuitions with respect to counterfactual cases: Is it possible that someone who is a student might not have been a student? Yes. Is it possible that someone who is a person might not have been a person? No. Presumably this sort of introspection is a reasonable method for determining modal relationships because we are posing questions about our own understanding of a concept. (However there is no particular assumption of infallibility, nor are more empirical methods excluded as possibly providing more useful or accurate results.)

A common device for supporting these thought experiments is the notion of (logically) *possible worlds* as used in Guarino and Welty's quotation above. Consider the complete and entire history of our universe, past and future. That is the *actual* world. But as things might have gone differently then they did in fact go the actual world is just one of many possible worlds. Some of the alternative possible worlds are only slightly different than the actual one (varying in the placement of a speck of dust say), others are quite different (with different biological organisms perhaps, or no organisms at all). Although this informal notion of a possible world can be made more precise (the formal semantics of modal logic re-defines possible worlds in purely mathematical terms (Hintikka 1962, Kripke 1963)), the intuitive notion will suffice for our purposes here. For a comprehensive overview of contemporary modal logic see Hughes and Cresswell (2001).

Rigidity may be re-expressed in the terminology of possible worlds like this:

D3: ϕ is a *rigid* property =df for all things x , if x has ϕ in some possible world, then x has ϕ in all possible worlds in which x exists.

Are FRBR Group 1 Entities Really Types?

Given their position in the FRBR ER model and the nature of the narrative treatment in the FRBR document it seems plausible that the four FRBR Group 1 Entities are being conceptualized as fundamental kinds, or *types*, of things, comparable to PERSON in the Guarino and Welty example. As noted, on the Guarino-Welty account one of the expectations for types is rigidity. If the FRBR entity types are not rigid then they are not types in the Guarino and Welty sense. (Nor could they be any of the other "backbone" entities that Guarino and Welty identify, such as "categories" or "quasi-types", as these also require rigidity).

The remainder of this paper focuses on whether this meta-property, rigidity, does indeed apply to the properties of being an item, being a manifestation, and being an expression. We argue first that being an expression is not a rigid property and then reason from that result to the further claims that being a manifestation is not a rigid property and being an item is not a rigid property. We conclude that these FRBR entities are not types, and that they are better thought of as *roles*. We then speculate as to what entities are take on these roles and in what circumstances. Before beginning there are three important methodological comments.

First, this analysis is an exercise in what might be called *rationalized descriptive* ontology. The emerging ontology is *descriptive* (rather than revisionary) in that we are describing the world as it seems to be to practitioners and not developing radical alternatives to these conceptualizations. Given this as a goal it is appropriate to adopt a method that takes the beliefs and modal intuitions of practitioners at face value initially, as the prima facie evidence for ontological distinctions. However we do allow revisions based on the need to resolve conflicts or accommodate new considerations. Such adjustments are justified internally by reference to other beliefs and perceptions, and are designed to retain as much of the original conceptualization as possible while improving logical consistency and generality — it is in that sense that the ontology is *rationalized*.

Second, we do not assume that these conceptualization are anything other than the products of the evolution of human neurology, cognition, and social practices. Other ways of framing the discussion are possible, but nothing beyond this sort of naturalistic contingent constructivism is necessary. (cf Guarino & Welty, 2000 p. 99)

Finally, we are aware that there may seem to be a mismatch between the nature of the actual FRBR project and the kind of analysis we conduct below. FRBR has a practical agenda and does not present itself with general ontological ambitions. In particular FRBR does not claim that its entity types are intended as *types* in any

sense with specific ontological significance, such as fundamental *kinds*. Consequently our labored arguments may seem willfully obtuse. However we are taking a deliberately naïve approach to FRBR, treating it as if it were, contrary to fact, intended as a general ontology of cultural objects, and not just a sound conceptual model for guiding practice. This is largely because we find FRBR, despite its modest stated objectives, to be extraordinarily promising and compelling as a general ontology. At the end of the paper we will revisit this issue.

Expressions are not types

We now argue that the property of being an expression is not a rigid property, and therefore expressions are not types. This section bears the burden of making plausible a fairly abstract modal claim. To that end we will present slightly different three (but structurally similar) arguments to the same conclusion. These arguments are abstract, complex, and involve concepts (such as meaning, necessity, languages, and works) that are notoriously elusive and regarding which there is little common agreement, but if we succeed in making plausible the claim that being an expression is not rigid we will then be able to use that assumption as the basis for the much simpler and more direct, arguments that manifestations and items are not types.

Argument A: Because meaning is assigned to sentences as a matter of contingent social/linguistic circumstances the property of being an expression is not rigid.

Are all expressions necessarily expression? As an example of a FRBR expression consider a particular text which realizes the work *Moby Dick*, for instance the text embodied in the 1851 Harper and Row edition. It indeed seems difficult to conceive of this text as possibly not realizing the work that it does in fact realize. Or, in the terminology of possible worlds, as not realizing in some other possible world the work that it in fact realizes in the actual world.

Difficult at first, but then when we allow our counterfactual imaginations full scope it is perhaps not so hard. After all, what meaning is attached to sentences (the constituents of texts) is a contingent matter, based on social circumstances in general and linguistic conventions in particular. So apparently those same sentences could, if social circumstances and linguistic conventions were different, mean something other than they do in fact mean, in which case they would realize a different work. The 1851 text of *Moby Dick* might, in different social/linguistic circumstances, have a meaning other than the one it has, and therefore not be an expression of the work *Moby Dick*, but, rather an expression of some other work. If this is right then texts do not realize, *necessarily*, the work that they do realize as a matter of fact. Or, in the terminology of possible worlds, a text does not realize in every possible world in which it exists the work that it realizes in this possible world. Generalizing further: a text can realize different works in different possible worlds.

It may be objected that the sentences in question would not be *those* sentences if they had some other meaning than the one they in fact have. They might be composed of the same symbols perhaps, or might have the same visual appearance, but they would not be the same sentences because the identity of sentences (the argument goes) is intrinsically tied to the syntactic and semantic rules, implicit in linguistic practice, that give those sentences meaning. While this response has some plausibility, it seems to us not quite decisive. Reflection on, for instance, how natural it is to say that the “same word” or “same sentence” can be used to mean different things seems consistent with the initial intuition that sentences (and therefore texts) have the meanings they have as a matter of contingent linguistic convention, and as those conventions could have been other than they are the meanings of those *same sentences* could have been otherwise as well.

However even if an expression could have realized some work other than the one it in fact realizes that alone does not establish that being an expression is not rigid, but only that the property of realizing some particular work is not rigid. Perhaps while it is true that expressions need not realize the work they in fact realize they must nevertheless realize some work or other. That is, we have not shown that an expression need not be an expression, but only that it need not be the expression that it is. However the very same considerations which suggest that sentences can have different meanings also suggest that they can have no meaning at all, and therefore not realize any work at all. Given that only contingent social conventions confer meaning it seems plausible for syntactically well-formed sequences of symbols, in a particular language, to have no meaning at all. And, again, we note that it is not uncommon to speak of syntactically well-formed sentences as “meaningless”.

Argument B: Because symbol sequences are sentences only as a matter of contingent social/linguistic circumstances the property of being an expression is not rigid.

While we believe that Argument A is sound there is a refinement which does not depend on the controversial possibility of meaningless sentences. This refinement replaces the claim that meaningless sentences are possible with the claim that the things which are sentences are not sentences necessarily, but only contingently. In other words, we argue to the conclusion that the property of being an expression is not rigid from the assertion that the property of being a sentence is not rigid.

The non-rigidity of the property of being a sentence will follow immediately if the things which are sentences in some worlds are not sentences in other worlds. A plausible candidate for an entity which would be a (meaningful) sentence in some worlds and not a sentence in others, is *a symbol sequence*. In worlds in which the appropriate social/linguistic circumstances obtain a particular symbol sequence will in fact be a sentence, but in other worlds, worlds without the appropriate social/linguistic circumstances, the same symbol sequence will not be a sentence.

Let us say then that sentences are symbol sequences in particular contexts of collective linguistic intention (cf Searle 1995). We can then allow that although meaningless sentences are impossible, the symbol sequences which are in fact sentences in this possible world (or in some other) are not necessarily sentences (i.e. are not sentences in all possible worlds), and so those symbol sequences are not necessarily meaningful and do not necessarily realize a work. And therefore being an expression is not a rigid property.

[This argument does require simultaneously denying that “all sentences are necessarily meaningful” while allowing “necessarily all sentences are meaningful”. If that seems puzzling it is probably because of a well-known and rather subtle modal ambiguity: the *de dicto* / *de re* distinction (Plantinga, 1969). Compare “Necessarily all students are students” with “All students are necessarily students.” “Necessarily all students are students” is true because the proposition “All students are students” is true in all possible worlds. “All students are necessarily students” is false because not every person who is a student in the actual world is a student in every possible world in which they exist. The first is a case of attributing necessity *de dicto* and says that a *proposition* is necessarily true, the second is a case of attributing necessity *de re* and says that an *individual thing* has some property in every possible world in which it exists. Rigidity as defined by Guarino and Welty is a matter of necessity *de re*. See Hughes and Cresswell (2001) for a contemporary review of this distinction.]

Argument C: Because being a work is assigned to a meaningful text as a matter of contingent social/linguistic circumstances, being an expression is not a rigid property.

We begin by distinguishing the literal meaning of a text (an expression) from the work realized by the text. To see why this might be appropriate imagine a possible world much like the actual one except that Melville never writes *Moby Dick* — but an orthographically identical manuscript is produced in 1975 by the American humorist S. J. Perelman. This is the same text as the 1851 Harper and Row edition, and in at least some sense of meaning the sentences of these texts have the same meaning, or close. But the creative *work* realized by the 1975 Perelman text would almost certainly be vastly different than the work realized by the 1851 Melville text: given Perelman's different artistic interests and cultural context. In such a world the text of *Moby Dick* exists, and has the meaning it has in our world, but there that text does not realize the work *Moby Dick* — the work that is realized by the 1851 Harper and Row text in our world. [This is an adaptation of an argument by Levinson (1980, 1990) which distinguishes sound sequences from musical works.]

In fact it might be argued that the English text of *Moby Dick* exists in any world in which the underlying language exists. This may seem very counterintuitive, but is actually consistent with discourse: a linguist may say "I am sure there is a sentence that has these properties, and I will find it", even though the sentence in question has never been uttered. We also appear to recognize the existence of moves in chess that have never been executed and tone sequences that have never been played. Similarly we might supposed that in any world in which the English language exists the text of *Moby Dick* will also exist, and with the same meaning it has in the actual world. But the existence of the *work* that is *Moby Dick* requires not just the existence of the meaningful text in some thin abstract sense, it requires actual authorial social action, which Melville, in our hypothetical world, has not taken. [This is again an adaptation of arguments by Levinson (1980, 1990) which distinguish sound sequences from musical works.]

We conclude by noting that each of the above arguments (A, B, C) proceeds by identifying a contingent relationship where one might have expected a necessary one. For the purposes of establishing the conclusion that being an expression is not rigid it doesn't matter whether we locate the contingency as lying between the sentence and its meaning, the symbol sequence and the sentence, or the meaning and the work, or somewhere else. All we require is that being a FRBR expression is not rigid is that one of those distinctions, or some other appropriately located distinction, is contingent. So while there is undoubtedly more to say for and against the these arguments we think we have at least shifted the burden of proof to anyone who would hold that being an expression is a rigid property.

Manifestations are not types

Making plausible the possibility that the property of being an expression is not rigid involved a rather lengthy appeal to relatively uncertain modal intuitions. However once we allow ourselves the result as assumption, however tentative, we can reason quite a bit more simply and directly to the conclusion that being a manifestation is not rigid. The argument that follows assumes that being an expression is not a rigid property, and more specifically that expressions are symbol sequences that realize works, and that symbol sequences realize different works in different possible worlds, and no works in some possible worlds (as in Argument B above).

- 1) Consider something *X* which in possible world *PW(1)* in a manifestation.
- 2) By the definition of manifestation *X* "embodies" an expression *Y* (in *PW(1)*).
- 3) Assume, as argued above, that an expression is a symbol sequence in certain contingent social circumstances.

- 4) Let S be the symbol sequence in virtue of which X embodies the expression Y
- 5) Assume, as argued above, that being an expression is not a rigid property (that symbol sequences that are expressions are not expressions necessarily, but only in the appropriate contingent social circumstances).
- 6) If so then there must exist a possible world $PW(2)$ such that the symbol sequence S which is expression Y in $PW(1)$ is not an expression at all in $PW(2)$,
- 7) Now suppose X exists in $PW(2)$.
- 8) In that case X will exist in $PW(2)$, but will not be a manifestation in $PW(2)$ (because in $PW(2)$ the social circumstances necessary for X 's symbol sequence, S , to be an expression are not present)
- 9) Therefore something, X , which is a manifestation one world $PW(1)$, is not a manifestation in another world $PW(2)$.
- 10) Therefore being a manifestation is not a rigid property.

At this point one may wonder: exactly what sort of thing, specifically, could be a manifestation in some possible worlds but not in others? Fortunately there is a plausible candidate: an (abstract) *physical kind*, identified by such things as composition, appearance, common causal antecedents. Examples of physical kinds would be species of plants and animals, chemicals, geomorphological structures, and so on, as well as artifacts considered apart from their social significance. An edition is a physical kind which embodies an expression in virtue of an associated symbol sequence. Presumably a physical kind could exist in a possible world even though the symbol sequence associated with it was not given meaning by the social agents in that world, in which case that symbol sequence would not be an expression and the corresponding physical kind would not be an edition. Imagine a freak accident at a printing plant that resulted in nothing but gibberish being inscribed in the books in a third shift print run. In that case instances of the physical kind would be identifiable, but they would not be manifestations (editions) because they would not embody an expression. And yet in some other possible world that gibberish *is* an expression.

We make three comments in defense of this conclusion. First, it would be inappropriate at this point to balk at the modest realism implied by allowing abstract physical kinds as "things" in possible worlds — FRBR straightforwardly interpreted is already committed to abstract objects of some sort or other. Second the argument is agnostic with respect to whether all physical kinds exist in every possible world; some physical kinds may be "initiated types" (Levinson 1980)). Third, we make no assumption as to whether there are physical kinds without instances.

Items are not types

The specific argument to the conclusion that being item is not rigid is exactly parallel to the preceding argument that being a manifestation is not rigid so we only sketch it here. In summary: if being a manifestation is not rigid there are things which are items in some possible worlds but not in others — because the manifestations in virtue of which the things which are items are items (in the possible worlds in which they are items) are not manifestations in every possible world in which the things which are items exist. As always the right social circumstances must obtain.

Once again one must wonder: modal logic-chopping aside exactly what sort of thing could be an item in some possible worlds but not in others? And again we have a good candidate: a *physical object*, individuated by concrete spatio-temporal career and actual concrete causal interactions.

Summary

We argued that being an expression is not a rigid property, and that if being an expression is not rigid, neither is being a manifestation nor being an item. If none of these properties is rigid, then, according to Guarino and Welty, they are not genuine types. Instead they appear to be some variety of *role*. In the course of this analysis we have identified three entities that *do* appear to be genuine types and to provide a foundation for refactoring the original FRBR notions. These new entities are physical objects, physical kinds, and symbol sequences.

The picture that emerges can be summarized this way: (i) expressions are *symbol sequences* in the role of realizing a work, (ii) manifestations are *physical kinds* in the role of embodying expressions, and (iii) items are *physical objects* in the role of exemplifying manifestations. In each case the actual instantiation of the role relationship by an entity is brought about by contingent social circumstances. This characterization is consistent with Searle's notion of natural objects (physical objects, physical kinds, symbol sequences) that count as social objects (items, manifestations, expressions) given appropriate social/linguistic circumstances of "collective intentionality". (Searle 1995).

But are physical objects, physical kinds, and symbol sequences themselves genuine types? They do seem to pass the Guarino and Welty test: it is hard to imagine that there is something which is a physical object in some possible world but not a in others; something which is a physical kind in some possible world but not in others; and something which is a symbol sequence in some possible world but not in others. Nevertheless the arguments given above do not themselves do not make these assumptions. Perhaps physical objects, physical kinds, or symbol sequences are roles as well, and more particularly social roles, also brought into existence by collective intentionality. Such cascades of social facts supervening on other social facts are also described by Searle (1995) — so there may be another turtle or two in the picture, even if it isn't turtles all the way down.

These considerations may not yet be decisive, but we believe we have done enough here to shift the burden of proof to anyone who would defend an alternative account.

The Part Where We Take It Back

Suppose our analysis is sound. One might still wonder just what the significance is for FRBR, and, in particular, whether it was really error for FRBR to treat items, manifestations, and expressions, as entity types.

It is natural to think of conceptual models as representing a relatively stable world of entities that lose and gain attributes and that enter into, and exit from, relationships (roles). We want a simple, expressive, non-redundant representation of these relationships. Guarino and Welty's ontology evaluation criteria are intended to promote robust and efficient models of this sort.

However in the case of actual bibliographic entities encountered in a particular possible world the FRBR relationships in question (realizing, embodying, exemplifying) are ones that, while perhaps not had rigidly are nevertheless such that there is never, or rarely, any practical likelihood of change. Once a symbol sequence (for instance) is recognized as realizing a work we need not, in planning cataloguing or content management systems, accommodate the possibility of that symbol sequence ever ceasing to have the role of being an expression that realizes that work, let alone realizing no work. The property of being an expression seems to be had permanently (within a single possible world) once it is had at all. Similarly, as it is hard to imagine a symbol sequence being recognized as within the domain of an information system for bibliographic control prior to its having the role of an expression which realizes a work, there seems to be little need to

accommodate the possibility of identified symbol sequences *acquiring* the role of realizing a work. Similarly for manifestations and items. The reader is no doubt eager to remind us that in the Borges story “Pierre Menard, Author of *The Quixote*”, the text of *The Quixote* does come to realize a new work. But this story is as striking as it is precisely because we do not expect such a turn of events.

In short, roles are well-suited for representing change, but where change is logically impossible, physically impossible, or even just highly unlikely, the advantage of converting types to roles in a conceptual model may perhaps be slight and add unnecessary complexities to cataloguing practice and system design. A *denormalized* ontology, such as FRBR, reflects this reality. For the problems of bibliographic control it would seem that this world is perhaps world enough.

References

- Coombs, J. H., Renear, A. H., DeRose, S. D. (1987). Markup systems and the future of scholarly text processing. *Communications of the ACM*, 30(11), 933-947.
- DeRose, S. D., Durand D. G., Mylonas, E., Renear, A. H. (1990) What is text really? *Journal of Computing in Higher Education* 1(2) 3-26. Reprinted in ACM/SIGDOC **Journal of Computer Documentation* 21(3) (August).
- Doerr, M., Hunter, J., and Lagoze, C. (2003). Towards a core ontology for information integration. *Journal of Digital Information*, 4(1).
- Doerr, M., LeBoeuf, P. (eds). (2007) FRBR: Object-oriented definition and mapping to FRBR-ER. Version 0.8. International Working Group on FRBR and CIDOC CRM Harmonisation, May.
- Dubin, D., Sperberg-McQueen, C.M., Renear, A., Huitfeldt, C. (2003). A logic programming environment for document semantics and inference. *Journal of Literary and Linguistic Computing* 18(1), 39-47.
- Dubin, D. (2003). Object Mapping for Markup Semantics. *Proceedings of Extreme Markup Languages 2003*, Montreal, Quebec, August, 39-47.
- Goldfarb, C. (1981), A generalized approach to document markup. *Proceedings of the ACM SIGPLAN--SIGOA Symposium on Text Manipulation*, New York: ACM 68-73.
- Guarino, N. and Welty, C. A. (2000). A formal ontology of properties. In *Proceedings of the 12th European Workshop on Knowledge Acquisition, Modeling and Management*. Lecture Notes In Computer Science. Springer-Verlag, 97-112.
- Guarino, N. and Welty, C. (2002). Evaluating ontological decisions with OntoClean. *Communications of the ACM*, 45(2), (February) 61-65.
- Guarino, N. and Welty, C. (2004). An overview of OntoClean. In *Handbook on Ontologies*, Springer Verlag.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press.
- Hughes, G. E. and Cresswell, M. J. (2001) *A New Introduction to Modal Logic*, Routledge.
- International Federation of Library Associations (IFLA). (1998). *Functional Requirements for Bibliographic Records: Final Report*. UBCIM Publications-New Series. Vol. 19, München: K.G.Saur, 1998.
- Kripke, S. (1963). Semantical considerations on modal logic, *Acta Philosophica Fennica*, 16, 83-94.
- Levinson, J. (1980). What a musical work is. *The Journal of Philosophy*, 77, 5-28.
- Levinson, J. (1990). What a musical work is, again. In *Music, Art and Metaphysics: Essays in Philosophical Aesthetics*. Ithaca: Cornell, 215-263.
- Plantinga, A. (1969). De dicto et de re. *Nous* 3(3) 235-258.
- Renear, A. H. (2000) The descriptive/procedural distinction is flawed. *Markup Languages: Theory & Practice*, 2(4), 411-420.
- Renear A. H., Dubin, D., Sperberg-McQueen, C. M., and Huitfeldt C. (2002). Towards a semantics for XML markup. In *Proceedings of the 2002 ACM Symposium on Document Engineering*. McClean, Virginia. Association for Computing Machinery, 119-126.

- Renear, A. H., Phillippe, C., Lawton, P., and Dubin, D. (2003). An XML document corresponds to which FRBR Group 1 entity? In *Proceedings of Extreme Markup Languages 2003*. Montreal, Canada, August.
- Renear, A. H., Xiang X., Lee, J. H., Choi, Y. (2005). Exhibition: A problem for conceptual modeling in the humanities. ACH/ALLC 2005, Victoria.
- Renear, A. H., Choi, Y., Lee, J. H., and Schmidt, S. (2006). Axiomatizing FRBR: An exercise in the formal ontology of cultural objects. Digital Humanities 2006 Conference Abstracts, Paris, France, 164-166.
- Renear, A. H. (2006). Is An XML document a FRBR manifestation or a FRBR expression? — Both, because FRBR entities are not types, but roles. *Proceedings of Extreme Markup Languages 2006*.
- Renear, A. H., Choi, Y. (2006). Modeling our understanding, understanding our models —the case of inheritance in FRBR. *Proceedings of the 69th ASIS&T Annual Meeting*, Austin, Texas.
- Searle, J. R. (1995). *The Construction of Social Reality*. New York: The Free Press.
- Sperberg-McQueen M., Huitfeldt, C., Renear, A. H. (2000). Meaning and interpretation of markup. *Markup Languages: Theory & Practice* 2(3), 215-234.
- Tillett, B. B. (2004). *What is FRBR? A Conceptual Model for the Bibliographic Universe*. Washington: Library of Congress, Cataloging Distribution Service. Online: <http://www.loc.gov/cds/downloads/FRBR.PDF>
- W3C. (1999). Extensible Markup Language (XML) 1.0 (second edition) Recommendation. W3C.

Acknowledgements

Early versions of this analysis were developed collaboratively in joint meetings of the Electronic Publishing Research Group, and the Knowledge Representation for Content Management seminar, spring 2006, at the Graduate School of Library and Information Science, University of Illinois, Urbana-Champaign. Participants included Yunseon Choi, Ingbert Floyd, Jin Ha Lee, Karen Medina, Sara Schmidt, Richard Urban, Xin Xiang, and Oksana Zavalina. A short report was given at Extreme Markup Languages 2006, (Renear 2006). We thank Ann Wrightson for first suggesting that the Guarino and Welty ontology evaluation criteria might be usefully applied in our document modeling projects.