Modeling Our Understanding, Understanding Our Models - The Case of Inheritance in FRBR

Allen H. Renear (corresponding author)

Graduate School of Library and Information Science University of Illinois at Urbana-Champaign, Champaign, IL renear@uiuc.edu

Yunseon Choi

Graduate School of Library and Information Science University of Illinois at Urbana-Champaign, Champaign, IL ychoi10@uiuc.edu

IFLA's Functional Requirements for Bibliographic Records (FRBR) presents a compelling and influential model of the "bibliographic universe." However there are interesting variations between FRBR's formal model and the narrative expositions of FRBR's authors and explicators - that is, between the formal model and the framework as more broadly understood by the FRBR community. In this paper we argue that despite a widespread belief to the contrary, attribute inheritance down the "hierarchy" of Group 1 entities is inconsistent both with the formal model and with the general spirit of the project. We believe these observations reveal an ongoing uncertainty about the nature of bibliographic entities as well as difficulties in maintaining a clear and exact understanding of the models we are using to represent those entities - even when those models are our own creation.

Introduction

The *Functional Requirements for Bibliographic Records* (FRBR) is a conceptual model of the "bibliographic universe" (works, texts, editions, documents and the like) that was developed by the International Federation of Library Associations and Institutions (IFLA 1998). It is intended to guide the development of systems for creating and managing bibliographic records. Although primarily aimed at the library cataloguing community, FRBR has been recognized more generally as a compelling empirically grounded framework for intellectual material and is increasingly influential in other domains that are using ontologies for content management.

Although some alternative models have been debated, there is little published analysis of the internal consistency of FRBR, and, specifically, of the relationship of the formal model to the assertions made in various expository accounts. We show that there are interesting variations between FRBR's formal model and the narrative expositions of FRBR's authors and explicators - that is, between the formal model and the framework as more broadly understood by the community. We believe these variations raise fundamental questions about both (i) the nature of bibliographic entities and (ii) our understanding of the models we use to represent those entities.

FRBR identifies four "Group 1" entity types: *work, expression, manifestation*, and *item*, defines relationships between them (items exemplify manifestations, manifestations embody expressions, expressions realize works), and assigns characteristic attributes to each entity - for instance, works have form, expressions may be in a particular language, manifestations may have a typeface, and items may have a provenance. It is not uncommon for writers explaining or commenting on FRBR to go on to say that attributes such as these are "inherited" down the "hierarchy" of Group 1 entity sets - from work to expression to manifestation, to item (Antelman, 2004; Bennett, 2003; Oliver, 2003; Taniguchi, 2002; Taniguchi, 2003; Tillett, 2002; Tillett, 2003; Vellucci, 1997). Presumably this would allow us to say of a particular item that it has a typeface, and of a particular manifestation that it has a language, even though the typeface attribute is assigned explicitly only to manifestations, and the language attribute is assigned explicitly only to expressions.

We argue that despite the initial plausibility of this widespread view such inheritance is (i) not part of the FRBR formal model; (ii) incompatible with the general FRBR modeling strategy, regardless of any particular formal implementation; and (iii) a dubious notion in any case.

We begin by observing that attribute inheritance is not part of the FRBR entity-relationship model and is not entailed by any statement in the the FRBR document. We then

demonstrate that general inheritance of all properties is clearly incompatible with the FRBR model (or any other similar entity/relationship based framework of bibliographic entities). Next we consider the whethr some form of limited inheritance could be defined that would be consistent with the general FRBR approach. We explore one strategy for this, but with inconclusive results. We then argue that the motivation for limited inheritance appears to be without foundation in any event, at least for the attributes in question. We develop three separate arguments against accepting limited inheritance within the FRBR general framework. The last argument includes an account of why the notion that inheritance exists in FRBR (and similar modeling frameworks) is so strong, and shows how the advantages supposedly secured by inheritance are already available within the framework as it exists, making inheritance unnecessary in any case.

Although we do not appear to be the first to suspect a difficulty along these lines we know of just one other brief direct mention of it (Manzanos 2003). There is however an extremely insightful and well developed treatment of many of these issues in Taniguchi (2003).

We think that our analysis has important lessons for understanding the creation and use of conceptual models. But we also believe there may be deeper issues at stake here as well. It seems to us that these problems with understanding of FRBR reveal the presence of unsettled fundamental questions about the basic nature of bibliographic entities; questions need to be better understood, although not necessarily resolved, if FRBR is to achieve its promise.

Overview of the FRBR Group 1 Entities

As described in the introduction FRBR recognizes four sets of Group 1 entities, works, expressions, manifestations, and items. Works are defined as “A distinct intellectual or artistic creation”, expressions as “the intellectual or artistic realization of a work in the form of alphanumeric, musical, or choreographic notation, sound, image, object, movement, etc., or any combination of such forms”, manifestations as “The physical embodiment of an expression of a work” and items as “a single exemplar of a manifestation”. (IFLA, 1998)

Each entity type is assigned a set of attributes. Works have such attributes such as title and form; expressions have a language attribute (translations of the same work are different expressions); manifestations have attributes like typeface; and items have attributes such as condition and location. The named relationships mentioned in the definitions hold between adjacent entities: a single work may be realized by a number of

expressions, each of which may be embodied in a number of manifestations, and each of those may be exemplified by many items. Below is an entity-relationship diagram representing these relationships, the arrowheads indicating cardinality.

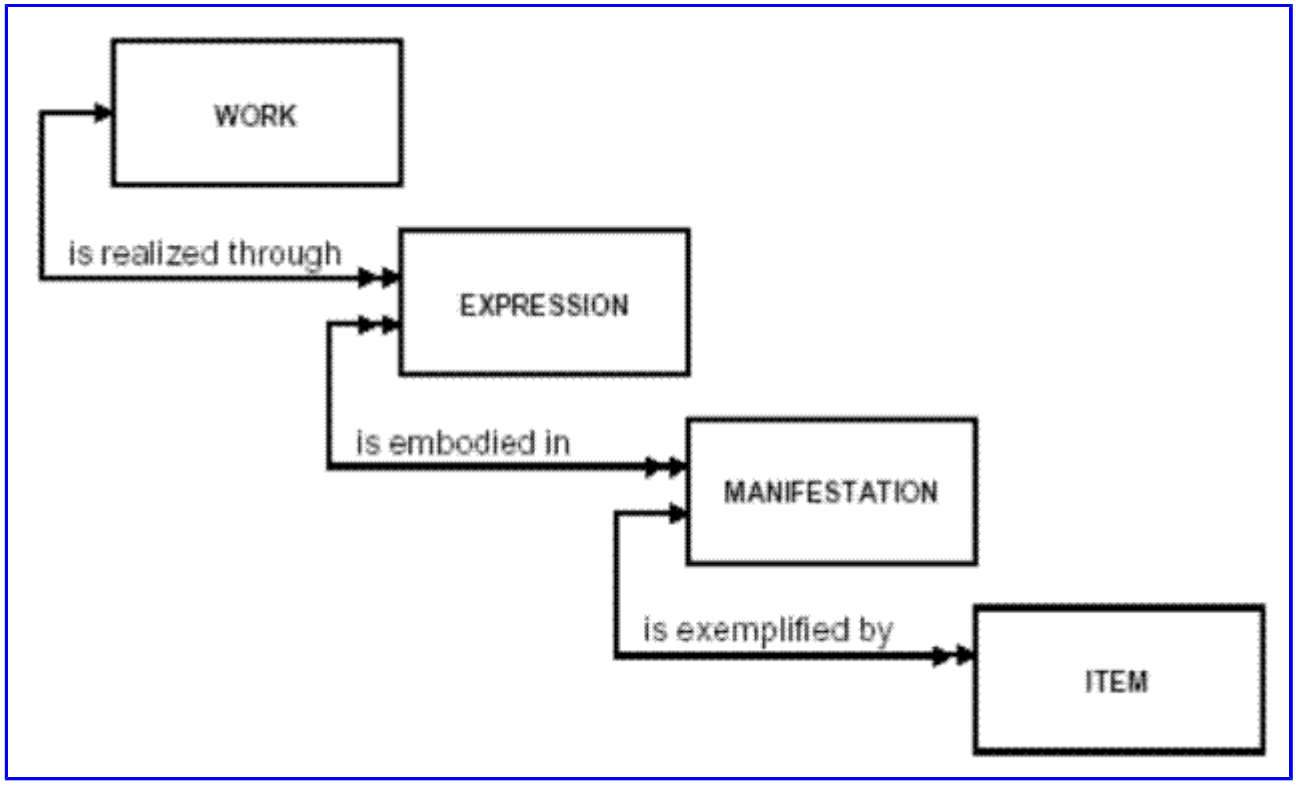

Figure 1: ER Diagram of FRBR Group 1 Entities and Primary Relationships (diagram from IFLA, 1998)

Although this diagram appears straightforward and intuitive, our analysis requires that we confirm some key points. The discussion below follows Chen's original description of entity-relationship modeling (Chen 1976), but is consistent with recent authoritative treatments as well (Elmasri and Navathe 2003, Silberschatz, Korth, and Sudarshan 2006).

In the FRBR entity-relationship diagram the rectangles represent sets of entities and arrows represent sets of relationships, where in both cases "set" is intended in the ordinary mathematical sense and a "relationship" is an association (ordered pair) of individual entities from the joined pair of entity sets. The arrowheads indicate cardinality constraints of one-to-many for realization, many-to-many for embodiment (to accommodate anthologies for instance), and one-to-many for exemplification.

Entity-relationship modeling has, appropriately, a liberal conception of what can be an entity. Entities are defined as "'a thing' which can be distinctly identified" (Chen 1976 p. 10), "a thing in the real world with an independent existence" (Elmasri and Navathe 2003), and "a 'thing' or 'object' in the real world that is distinguishable from all other objects"

(Silberschatz, Korth, and Sudarshan 2006). At the level of the modeling formalism further restrictions are typically avoided - explanations of entity-relationship modeling almost always say that entities may be “abstract” or “concrete” (Silberschatz, Korth, and Sudarshan 2006), or, in alternative terminology, “physical” or “conceptual” (Elmasri and Navathe 2003). It is natural then to ask with respect to each of the FRBR Group 1 entities whether it is abstract or concrete. FRBR tells us explicitly that works are abstract and items concrete, but does not categorize the other two entity types, expressions or manifestations, in these terms. In light of the absence of evidence to the contrary it is natural to suppose that expressions are abstract - as, whenever abstract objects are allowed at all linguistic objects such as sentence types, which may be multiply instantiated, are almost always classified as abstract.

The manifestation entity type presents a more complicated case. One problem is that manifestations are described by FRBR as “physical”, and “physical embodiments”. On the natural assumption that physical things or “physical embodiments” are concrete, this suggests that manifestations are concrete. But the inference trades on an ambiguity. The Harper and Brothers 1851 edition of *Moby Dick* is, in a sense, a physical object and my copy of that edition is certainly a physical object; but the phrase “physical object” clearly has two different, if systematically related, meanings in the two predications. The term “concrete” has a similar ambiguity: the Harper and Brothers edition of *Moby Dick* is not concrete in the same sense in which my copy of it is. Finally, being physical in the first sense implies being concrete in the first sense, but it does not imply being concrete in the second sense.

The problem is compounded by characterizations such as “the set of copies produced in each case constitutes a manifestation”. Understood literally this seems to make the members of the manifestation entity set themselves sets (sets of physical items), which is not what would be expected here and is not parallel with the other FRBR entities. Moreover, although sets themselves are abstract entities regardless of the status of their members, sets of physical entities are routinely confused with their corresponding aggregates or fusions - and aggregates and fusions of physical objects, unlike the sets they correspond to, are concrete, and in the strictest sense: the sense in which my copy of the Harper and Brothers 1851 edition of *Moby Dick* is concrete.

These problems and others have already been identified (Doerr, Hunter, Lagoze, 2003). We will not press the points further here but will say that in our view, and this is broadly consistent with current thinking, the manifestation entity set must be understood as a set of abstract objects none of which is itself a set or an item and each of these abstract objects will be associated with the set of items, these items being the manifestations exemplars (items). Manifestations are not concrete physical objects, strictly speaking (i.e.

in the sense in which items are concrete and physical), but only in this derivative sense: all of their exemplars are concrete and physical.

This complex clarification was necessary because some of the confusion about whether it is appropriate to speak of inheritance in FRBR seems connected with the possibility of treating bibliographic entities such as manifestations, expressions, and works, alternately as either abstract objects that are instantiated by concrete objects or as the corresponding sets of (those) concrete instantiating objects. But consistently and correctly understood FRBR takes the former approach, not the latter.

An Easily Accommodated Unmodeled Feature: Disjointness

Before going on to take up the problem of inheritance we give a simple example of how feature of the FRBR general approach which, like inheritance, seems to be taken for granted by the FRBR community, and which, like inheritance, is not explicitly part of the FRBR formal model, can nevertheless be accommodated within the FRBR model quite easily - unlike inheritance. The accommodation does require an extension to the entity-relationship modeling language, but the extension is simple and well within the general spirit of the FRBR.

In any discussion of FRBR it is always assumed that the Group 1 entity classes are disjoint, that there is no entity that is a member of more than one entity class. But this disjointness is not explicitly modeled in the FRBR entity-relationship diagram and not explicitly mentioned in the FRBR documentation. The remedy of course is simply to assert disjointness in the text of FRBR, and to add the appropriate symbols to the ER diagram to the same effect.

While it is actually not possible to represent disjointness with the standard constructs of basic entity-relationship modeling, we can draw on the constructs of “enhanced entity-relationship modeling” (EER). Using these constructs the diagram Figure 2 asserts that work, expression, manifestation, and item are disjoint subsets of the class of Group 1 entities.

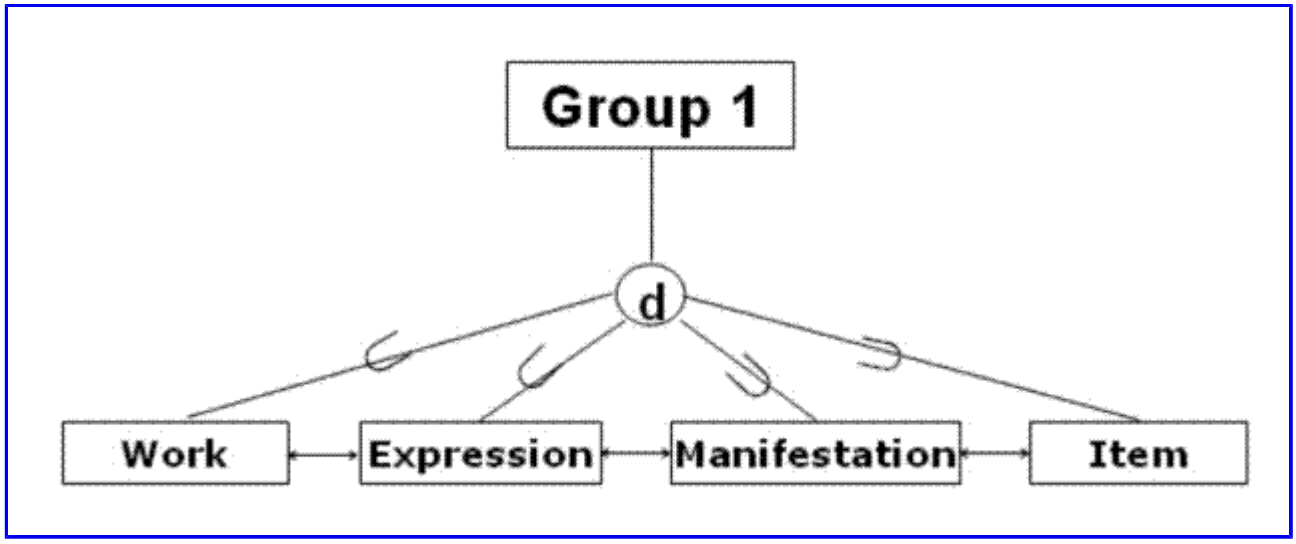

Figure 2: Modeling disjoint relationships with Enhanced Entity-Relationship Modeling (EER)

Of course we should never assume that the addition of a modeling construct is without wider significance. Additional expressiveness in a modeling language can have consequences for computational efficiency and interoperability. Nevertheless from the point of view of the conceptual integrity of the FRBR framework the modification is natural and straightforward and brings the entity-relationship model closer to our general understanding of the FRBR view of the bibliographic universe. We will see that things are quite different for any attempt to capture the unmodeled but widely accepted feature of attribute inheritance.

Inheritance and FRBR

In presenting the FRBR model researchers and practitioners both often use expressions such as hierarchy and inherit, and describe entities at "lower levels" of the Group 1 entity "hierarchy" as "inheriting" the characteristics of their corresponding entities at higher levels. Some examples of sentences along these lines are:

- Vellucci: "hierarchical levels ... each of which incorporates the attributes of its superordinate entity" (Vellucci, 1997)
- Tillett: "a characteristic of a work is carried to all the entities below it in the hierarchy through a transitive relationship" (Tillett, 2002)
- Oliver: "the characteristics of the work belong to all expressions, all manifestations, and all items[...] characteristics (or attributes) are inherited by all lower levels of the hierarchy" (Oliver, 2003).
- Bennett: "FRBR's primary benefits extend from its hierarchical structure, permitting the placement of bibliographic information at its appropriate level of abstraction

and facilitating its inheritance" and also "The structure of the FRBR model implies the existence of certain information that applies at the highest level of the bibliographic hierarchy - the work - and therefore also applies to, or is inherited by, bibliographic entities comprising the lower levels of the hierarchy" (Bennett, 2003)

- Taniguchi: "The inheritance of attribute values from an upper-level to a lower-level [FRBR] entity can easily be understood..." (Taniguchi, 2002)
- Taniguchi: "A typical example of the adoption of this method ['hierarchical way'] is the FRBR model ... the hierarchical way presupposes in principle the inheritance of attributes from an entity to its subordinates... (Taniguchi, 2003)
- Antelman: "Attributes at the appropriate level are associated with the highest possible entity and are inherited-not repeated-by lower level 1, which is an XML-like hierarchical representation of a work record" (Antelman, 2004)

These are subtle issues and without further examination of the context we should not conclude with respect to any particular case above that the writer truly intends to assert general attribute inheritance, but nevertheless the collective effect of such sentences tend in that direction. In at least some cases then the idea seems to be that an attribute such as language is assigned to the entity type understood to be "at the appropriate level of abstraction" (Bennett 2003) or the "highest possible entity" (Antelman 2004), which is to say, in this case, the expression entity type. Then in virtue of the assumed inheritance such an attribute is also understood to be inherited by the "lower" entity types. So if an expression is in a particular language then all manifestations embodying that expression are in that language and all items exemplifying those manifestations are also in that language.

There are several difficulties in these descriptions. We mention them now and then take them up as needed in our discussion in later sections. First it is not clear that terms like "properties", "characteristics", "attributes", and "attribute values" are always being used in the same way, or, in the case of "attributes" and "attribute values" being used consistently with their technical senses in conceptual modeling. Second it is not clear why FRBR expositors consider FRBR a "hierarchy" with "upper" and "lower" levels and what that actually means in this context. Finally, it is not clear what single relationship [type] exists in FRBR that would "transitively" support inheritance.

We now begin our analysis by noting that there is nothing in the FRBR document, or in the FRBR entity-relationship diagram that expresses the assertion that attributes, attribute values, attribute value pairs, properties, or characteristics are inherited. We have found no word cognate or synonymous with "inherit" in the document and an inspection of the entity-relationship diagram suffices to show immediately that it does not entail any inheritance either.

In fact, it is easy to show that general unlimited inheritance down the FRBR "hierarchy" is certainly impossible. The argument is simple: FRBR describes works as abstract and items as concrete. If all properties of "higher" entities are inherited by "lower" entities then items inherit the property of being abstract, and therefore items will be both abstract and concrete. But nothing is both abstract and concrete - therefore there is no unlimited general property inheritance in FRBR.

It may be objected that this argument is sound but irrelevant as the only properties ever at issue were the attributes (or attribute values, or attribute/value pairs) explicitly specified in FRBR. So for the work entity the only relevant properties are ones such as title, form, context, and so on; but not the property of being abstract. On this account only specified attributes (and/or attribute values) are inherited and the argument given does not apply as "abstract" and "concrete" are neither attributes nor attribute pairs. We believe that even this limited version of inheritance is misconceived, but before presenting our arguments against it we explore it further.

Modeling Limited Inheritance

In entity-relationship modeling there is no general expectation that inheritance will occur across an arbitrary relationship, so if inheritance is desired it must be explicitly specified. Given the general nature of entity-relationship modeling (distinct entity types, each with an appropriate attribute set) any inheritance will be limited to certain attributes and not others. If FRBR does wish to incorporate inheritance into its framework it must articulate rules for limited inheritance. In what follows we explore a little how this might be done, but end inconclusively - it is a difficult project. In the end however the difficulty in defining inheritance turns out not to be a problem for us, because, as we show in the next section, where bibliographic attributes are concerned the expectation of inheritance rests on a mistake.

Inheritance may be described in general as one entity receiving "properties or characteristics of another, normally as a result of some special relationship between the giver and receiver" (Danforth and Tomlinson 1988, cited by Taivalsaari 1996). A first cut at an inheritance rule might be:

D1: F is inherited with respect to R =df $(\forall x)(\forall y)[Rxy \supset (Fx \supset Fy)]$

This definition says that a particular property F is inherited with respect to a relation R whenever it is the case that if something x bears R to something y, then if x is F then y is also. Perhaps then is blue (F) inherits with respect to has as a part (R) because if

something is part of a blue thing then it is itself a blue thing. A bibliographic example might be: being in English is inherited with respect to the relationship embodies.

One may also more generally, and more usefully, define inheritance not for individual determinate properties, but for general determinable attributes combined with their values:

D2: Amz is inherited with respect to R =df ($\forall x$)($\forall y$)($\forall z$)[(R*xy* $\supset$ (A*xz* $\supset$ A*yz*)]

Here A is a two place predicate representing a determinable attribute such as has the color and z ranges over possible colors, such as blue, understood as values, not predicates. The application to the color example would be: the color of a thing is inherited by its parts. The bibliographic example here might be: the language of a manifestation is inherited from the expression the manifestation embodies.

These are reasonable beginnings and certainly necessary conditions for relationship-based inheritance, but hardly unproblematic: among other things all of the material conditionals in **D1** and **D2** may be trivially satisfied, producing counterexamples against the sufficiency of the analysis, or at least a problem for our intuitive notion of inheritance. Possible repairs involving modal notions suggest themselves of course, but whether in the end a sound intuitive definition is possible remains to be seen.

Apart from such technical problems with defining inheritance logically how close are we to formalizing the related claims of FRBR expositors? We can suppose that eventual versions of **D1** and **D2** be used to support individual inheritance across each Group 1 relationship, but Tillett refers to "a transitive relation" that supports inheritance of all properties down the "hierarchy" [Tillett 2003]. What relation would that be exactly? Although there is no single named FRBR relationship that fits the bill here one may be defined in terms of given FRBR relationships:

D3: carries(*x,y*) =df
[realizes(*x,y*) v embodies(*x,y*) v exemplifies(*x,y*)]
v
($\exists$ z)[carries(*x,y*) & carries(*x,y*)]

Carries (we take the term from Tillett's remark, mentioned above) is transitive and may be used for R in **D1** and **D2** to ensure that all inheritable properties are inherited down the hierarchy.

The Set Theoretic Alternative to FRBR - and Some Ambiguities

FRBR has been analyzed critically by a number of authors (Weinstein, 1998; Abrams, 2002; Taniguchi, 2002, Carlyle 2004). Some have proposed alternatives with more or fewer entities, or different assignments of attributes to entities, or a radically different general framework altogether. A significant alternative of the latter sort is a set-theoretic approach (e.g., Svenonius (2001). Both set-theoretic and abstract entity models use sets to model domains, but the set-theoretic strategy is distinctive in using sets of similar concrete objects as a strategy for eliminating the need for abstract entities.

In the set theoretic model we identify the set of concrete objects that would be colloquially described as belonging to the same "edition". We perform this identification by specifying a common characteristic, a particular similarity, or a relationship to some other object - anything other than instantiation of an abstract object. Now we can explain the meaning of a casual expression that seems to refer to an abstract object by reference to this set. Abstract entities (work, expression, manifestation) are replaced with classes of physical entities exhibiting common characteristics. For instance items belonging to the same edition are represented as a set of all items having a certain common publishing history, rather than as related in a particular way to an abstract object (a manifestation). That a number of items have the same text is represented by identifying a set of textually similar items, or items faithfully produced (directly or indirectly) from an original. These sets form a taxonomic class hierarchy: all items with the same publishing history are also textually similar, all textually similar items also have the same intellectual content. This alternative is a standard strategy in philosophy for avoiding ontological commitment to (presumably dubious) abstract entities.

Both taxonomic class hierarchies and entity-relationship diagrams are often represented as boxes and arrows, But to see inheritance as (generally) implied by the arrows between entity classes in an ER diagram would be, of course, a misunderstanding. The ER boxes do represent classes of entities, and the relationships are sometimes, in some sense, hierarchical. But the relationships are not typically class/subclass relationships - and that is what is needed to conclude that properties are inherited on the basis of the boxes and arrows alone. FRBR expositors are of course not confused about this, but the ambiguity is part of what makes any discussion confusing.

In traditional taxonomic class hierarchies the boxes represent classes of things and an arrow indicates that all members of a class (box) pointed to are members of the class pointed from. This relationship can be formalized either in terms of sets or in terms of properties related by implication. The latter might be described as a concept hierarchy rather than a class hierarchy, but the two formulations, while different, are obviously functionaly equivalent. (All whales are mammals if and only if the set of whales is a subset

of the set of mammals.). The consistent direction of the arrows representing the subclass relationship has the result that any property possessed by all members of a class S is also possessed by all members of any subclass: if all dogs are mammals, and all mammals have kidneys, then all dogs have kidneys. The subclass relation being transitive this reasoning may be iteratively applied along the chain of arrows, to establish inheritance down the tree ("transitive closure"). However, in ER diagrams the boxes continue to represent sets of entities, but the arrows no longer represent subclass relationships, instead they represent arbitrary substantive relationships that may hold between members of the entity sets. In an ER diagram, unlike the taxonomic class diagram, that two boxes are linked by an arrow does not imply a subclass relationship, or, in itself, inheritance.

Inheritance is most often conceptualized as sub/super class inheritance: a member of a subclass inherits all the properties of its superclass in virtue of the fact that all members of a subclass are members of the superclass- the relationship expressed by a taxonomic class hierarchy. However, the FRBR ER diagram is not a taxonomic class hierarchy and in fact basic EER does support sub/super class inheritance. EER however does and so could generate a parallel taxonomic hierarchy similar to the principal hierarchy of a set-theoretic approach. That is, for every cascade of work, expression, manifestation, and items a taxonomic class hierarchy could be generated that would support class-based inheritance. However this would not seem to do justice to the intuition had by the FRBR explicators that properties are inherited by entities in virtue of their inherent bibliographic relationships to other entities, rather than through a subclass/superclass relationship.

Graphical modeling systems like ER diagrams and taxonomic class hierarchies have advantages over linear symbolic systems like first order logic. Not least their visual intuitive nature allows casual use, without reference to explicit formalization. But experience within the AI community in the 1970s-80s revealed the disadvantage: it is easy to overload the graphics and misunderstand the relationships. (Woods 1975, Brachman 1983).

Against Inheritance in FRBR

The project of defining a notion of limited inheritance in order to accommodate the inheritance seen in FRBR is interesting, but we believe it to be both unnecessary and mistaken in its assumptions. We would argue that in fact there should be no inheritance of specified attributes at all, in any case, and so there is no need to accommodate limited inheritance formally.

We have already given a simple argument against general unlimited property inheritance. But noted that it depended on properties that are not explicit named attributes or attribute values. Now we present three arguments against accommodating even limited inheritance of explicitly specified attributes. These latter considerations are not decisive, unlike the argument against general property inheritance, but we believe they are very strong nonetheless.

The Argument from Category Error

This argument is simple, although turns almost immediately on fundamental ontological intuitions. If attributes are inherited then a concrete physical item might have the properties of, say, physical condition and language. It might be, for instance, both torn and in German. This of course is the very point of inheritance. But can such ontologically disparate properties really be had by the very same thing? We think not. An item may be torn, but it is not the kind of thing that can be in German. No physical object is "in German". Only linguistic entities such as sentences or texts can be in German - and yet these things, being abstract and not physical, cannot be torn. Of course we do commonly say that an item is both torn and in German. But this should not be taken literally. The meaning of such a sentence would be more accurately expressed this way: an item exemplifies a manifestation that embodies an expression that is in German. Much too cumbersome for ordinary conversation of course, but a more accurate account nevertheless.

The Argument from Modeling Strategy

The use of entity-relationship modeling involving abstract entities and concrete entities is very common and may be seen as a kind of practical but avowedly ontologically agnostic exploitation of the view just presented. It is a picture of the world "as if" the platonist ontology of abstract and concrete objects were correct. Modelers are not necessarily platonists, but they often find it convenient to talk as if they were; such a perspective facilitates the assignment of attributes to the "appropriate level of abstraction" and is conducive to creating normalized relational models. So regardless of one's ontological beliefs, to add attribute inheritance to FRBR is still inconsistent with the general approach: using entity-relationship model of related abstract and concrete objects.

The Argument from Redundancy

Here we argue that the practical motivations for inheritance are mistaken as the supposed advantages of inheritance can be more naturally secured in other ways. In particular we show that inheritance is not needed for collocation or for database denormalization.

It might be thought that supporting the collocating function of bibliographic systems depends on inheritance. But it does not. Suppose we wish to do something we might informally describe as "identify all items that have for the class number attribute the value P95". We might think that this requires, conceptually, that the items being collocated really do have a class number attribute, and that they also have same value, P95, for that attribute. And if so then apparently they must have inherited that attribute/value combination from their corresponding works. Not so. The items in question are collocated because they all have this property: they exemplify a manifestation that embodies an expression that realizes a work that has P95 as a value for the attribute class number. And that is a property those items already have - no inheritance is necessary.

Another concern is achieving efficiency in relational databases by denormalization, reducing order to the number of joins needed to process certain queries. The idea here is that in a relational database developed from the FRBR conceptual model we may wish to, for instance, add a class number column to the item table and use it to hold the number of each corresponding work's class number, reducing the number of table lookups for queries that ask with respect to an item what the class number of its corresponding work is. This is all fine, but it has nothing to do with attribute inheritance. To see why notice that the significance of the column values are different in the work table and the item table. In the work table a class number in a record is the class number for the entity associated with that record, but in an item table a class number in the new column is not the class number of that item record's entity. Of course the headers of both columns may very well each say "Class Number", but if so those headers are being used with different meanings and are not indicating the same attribute. This is consistent with the technical definition of attribute as the role, or name of the role, played by a domain of values, and not the domain values itself.

Another way of understanding the assertion of inheritance is that what is being inherited are not attributes at all, but attribute values. Although this may appear more promising that appearance is an illusion. On this account the value of an item's Class-Number-of-Corresponding-Work attribute is inherited from the Class Number attribute of its corresponding work. And that of course is similar to what is actually going on in denormalization. But why should this be called inheritance? Ordinarily inheritance refers to an attribute or an attribute value pair, or, more generally a property or characteristic, not an attribute value. The very notion of an attribute value being inherited

is obscure. The association of an attribute value, by itself, with an object tells us nothing about the object: The association of an attribute value (e.g. "red") with an object does not distinguish "is colored red" from "likes red"; there is no implicit reference to the attribute whose domain the value "comes from" and that value may in fact be in the domains of many attributes. So neither bibliographic attributes nor attribute values are inherited. Using "inheritance" to refer to denormalization is not innocent - there may be cases of real inheritance in the bibliographic universe and we should not confuse the two.

Speaking Strictly vs. Speaking Loosely

Of course it goes without saying that the logically perspicuous locutions are cumbersome, and that humans have no problem navigating idioms, ambiguities, and overloaded expressions. Indeed we cannot imagine practical work without these shortcuts, so we certainly do not suggest that we stop, in ordinary conversation, saying things like "that book is torn … and it is in English." And indeed we may allow that there is a derivative sense of "being in English", defined as being an item that exemplifies a manifestation that embodies an expression that is in English. Nevertheless we do recommend avoiding claiming, when talking about the bibliographic universe, that attributes, or attributes values, or other bibliographic properties are, in general, inherited.

Conclusion

FRBR is a profound and influential formal model for intellectual content. We look forward to both its wide adoption and to the basic research it will inspire. Not surprisingly, given the difficulty and ambition of the agenda, some of the details are not yet quite clear. The required clarification appears certain to immerse us in some of the deepest issues in the modeling and representation of intellectual objects - we look forward to that as well. We believe efforts to make our discourse about bibliographic entities consistent and adequate to its objects will probably bring us back, again, to the fundamental question: what exactly are those objects?

References

References Abrams, S. L. (2002) *Taxonomic views of a digital object* Retrieved June 20 2006 from http://hul.harvard.edu/~stephen/Model.doc

Albertsen, K., & Nuys, C. (2005) Paradigma: FRBR and digital documents *Cataloging & Classification Quarterly* 39(3/4), 125-149

Antelman, K. (2004) Identifying the serial work as a bibliographic entity *Library*

Resources & Technical Services 48(4)

Bennett, R. & Lavoie, B. (2003) The concept of a work in WorldCat: An application of FRBR *Library Collections, Acquisitions & Technical Services* 27(1)

Brachman, R. J. (1983) What IS-A is and isn't: An analysis of taxonomic links in semantic networks *Computer* 16:10, 30-36

Carlyle, A. (2004) FRBR and the bibliographic universe, or how to read FRBR as a Model *2004 ALA Annual Conference, Orlando*

Chen, P. (1976) The entity-relationship model - Toward a unified view of data *ACM Transactions on Database Systems* 1:1

Danforth, D., & Tomlinson, C. (1988) Type theories and object-oriented programming *ACM Computing Surveys* 20:1

Doerr, M. Hunter, J. & Lagoze, J. (2003) Towards a core ontology for information integration *Journal of Digital Information* 4(1)

Elmasri R. & Navathe, S. B. (2003) *Fundamentals of Database Systems* Addison-Wesley

IFLA Study Group on the Functional Requirements for Bibliographic Records, & International Federation of Library Associations and Institutions. (1998) *Functional requirements for bibliographic records : final report* München: K.G. Saur

Jones, E. (2004) Functional requirements for bibliographic records (FRBR): Summit on Serials in the Digital Environment *CONSER: Cooperative Online Serials*

Manzanos, N. (2003) Re: Presention. 22/07/2003 *FRBR Listserv Archive* http://www.ifla.org/VII/s13/wgfrbr/archive/FRBR_Listserv_Archive.pdf

Oliver, C. (2003) Functional requirements for bibliographic records: What is FRBR and why is it important? *Canadian Metadata Forum* http://www.collectionscanada.ca/metaforum/n11-228-e.html

Renear, A. H., Phillippe, C. Lawton, P. & Dubin, D. (2003) An XML document corresponds to which FRBR Group 1 entity? *Extreme Markup Languages 2003 Montreal, CA*

Silberschatz, A. H. F. Korth, & S. Sudarshan. (2006) *Database System Concepts* McGraw-Hill 5th edition

Svenonius, E. (2001) *The Intellectual Foundation of Information Organization* Cambridge, Massachusetts: MIT Press

Taivalsaari, A. (1996) On the notion of inheritance *ACM Computing Surveys* 28:3

Taniguchi, S. (2002) A Conceptual Model Giving Primacy to Expression-level Bibliographic Entity in Cataloging *Journal of Documentation* 58(4), 363-382

Taniguchi, S. (2003) A Conceptual Model Giving Primacy to Text-level Bibliographic Entity in Cataloging: A Detailed Discussion http://www.slis.tsukuba.ac.jp/~taniguch/report200301.pdf

Tillett, Barbara B. (2002) The FRBR model (Functional Requirements for Bibliographic Records) *Workshop on Authority Control among Chinese, Korean and Japanese Languages, March 14-18, 2002, Karuizawa, Tokyo, Kyoto Institutions* http://www.nii.ac.jp/publications/CJK-WS3/cjk3-07a.pdf

Tillett, Barbara B. (2003) The FRBR Model *ALCTS Institute on Metadata and AACR* http://www.loc.gov/catdir/cpso/frbreng.pdf

Tillett, Barbara B. (2003b) *What is FRBR?: A conceptual model for the bibliographic Universe* Library of Congress Catalog Distribution Service, Washington, DC, 2003 http://www.loc.gov/cds/FRBR.html

Vellucci, S. L. (1997) Bibliographic Relationships *International Conference on the Principles and Future Development of AACR* Toronto, Canada, October, 1997

Weinstein, Peter (1998) Ontology-Based metadata: Transforming the MARC legacy *Proceedings of the Third ACM Digital Library Conference*

Vellucci, S. L. (1997) Bibliographic Relationships *International Conference on the Principles and Future Development of AACR. Toronto, Canada, October, 1997*

Woods, W. A. (1975) What's in a link: foundations of semantic networks In D.G. Bobrow and A.M. Collins (eds) *Representation and Understanding: Studies in Cognitive Science* New York: Academic Press