The Exhibition Problem. A Real-life Example with a Suggested Solution

Øyvind Eide
Unit for Digital Documentation, University of Oslo

Abstract

As the number of scholarly encoded digital texts is increasing, creating models of these kinds of texts with the help of digital tools is becoming more and more interesting. In connection with this type of work, it is important to have a clear understanding of what these particular models are based on. They will clearly be based on certain readings of the source texts, but we need to keep track of the relationships between the texts, readings of the texts and the models based on such readings.

In this article, a problem of potentially great significance for this kind of modelling is discussed. The problem is called the exhibition problem and is based on the difference in ordinary linguistic communication between asserting a fact, e.g. that a certain person has a certain name, and exhibiting the same fact. In many cases, the latter is modelled as if it was the former. As a solution to this problem, an event-oriented modelling method is proposed.

Correspondence:
Øyvind Eide,
Unit for Digital Documentation at the Faulty of Arts, University of Oslo, P.O. Box 1123 – Blindern, N-0317 Oslo, Norway.
E-mail:
oyvind.eide@edd.uio.no

1 Background

1.1 Text encoding

Text encoding has been an integrated part of the work of many scholars for a number of years. Encoded texts, especially texts coded on the basis of the TEI guidelines (TEI, 2002), are being created and used in a number of ways. Web presentations and printed editions are common ways of using the products of text encoding.

Work related to the understanding of encoding as such has followed text encoding for a long time. An example of this is the discussion concerning overlapping hierarchies in texts (e.g. Renear *et al.*, 1996), which is a discussion that to a great extent emerged as a result of text encoding. More recently, the meaning of markup itself is a topic for discussions, as seen in e.g. Sperberg-McQueen *et al.* (2002). One of the reasons why these discussions are important is that they will influence the methods used for building databases based on facts found in XML texts. These facts can be expressed in the text as such, in the XML tags or through various combinations of the two.

Databases based on facts expressed in XML documents differ from databases built in order to represent the text as such. An example of the latter is the letters published by the Documentation Project in Norway (Eide, 2003). The text of each letter is stored in a database, but the database is merely used as a tool to organize the blocks of SGML fragments as well as the meta-data connected to them.

1.2 Modelling

Databases based on facts expressed in XML documents are built in order to reflect not only a text as such, but also the state of the world as it is described in the text, or, more specifically, the subset of such a

Literary and Linguistic Computing, Vol. 23, No. 1, 2008.

doi:10.1093/llc/fqm040 Advance Access Published on 24 November 2007

state revealed by a specific reading. A modelling of readings of texts of this kind is described by Willard McCarty as one of the important parts of humanities computing (McCarty, 2005).

A short example sentence will illustrate this point:

[1] 'A man of old age spoke.'

We may consider a document containing only sentence [1] above. Based on the text string in [1], such a database will not only store the string in connection with other strings in the text, with reference to the printed source, etc. It will also store facts found through reading and understanding the string, such as:

- There exists a person x.
- x is old.
- x is a male.
- x spoke at time *t*.
- $t <$ present.

The actual facts stored in the database will vary with the model used. But they will usually, if not always, take the extension of expressions in the texts into consideration. Consequently, not only a name such as 'Herman Melville', but also aspects of the person or persons referred to by the name, as well as that particular person's relationships to other entities around him, are included in this kind of model.

2 The Exhibition Problem

In a paper presented at ACH/ALLC 2005, Allen H. Renear *et al.* describe a problem of potentially great significance for the modelling of readings of texts (Renear *et al.*, 2005). They argue that:

> In ordinary linguistic communication we often use a name to refer to something in order to then go on to attribute some property to that thing. However when we do this we do not naturally construe our linguistic behavior as being at the same time an assertion that the thing in question has that name. (Renear *et al.*, 2005, p. 176)

Further, they claim that this distinction is overlooked when conceptual models based on encoded texts are developed. This observation gave rise to some consideration of our own methods at the Unit for Digital Documentation at the University of Oslo, resulting in this article.

We have been using SGML- and XML-encoded material as sources for several of our databases (Holmen and Uleberg, 1996, Holmen *et al.*, 2007). The way this is done is by marking up texts both descriptively and interpretatively, followed by the use of software to extract information which is included in the databases. When we build up a database based on information found in written material, it is important to document in the database as to who is responsible for the facts in the database. Which facts are asserted by the text, and which are exhibited in or presupposed by the text? If Renear's argument is correct, we may infer that the databases include assertions that are based on information in the source texts that is, strictly speaking, not grounded in these texts. For example, we could be using a text as the source of a naming in the database whereas the naming is merely exhibited, and not asserted, in the text. When modelling the reading of texts, we have to take this problem into consideration.

But why is it actually a problem? An example of the problem can be found in the difference between sentences [4] and [5] below:

> [4] «Moby Dick was written by Herman Melville.»
> [5] «The name of the author of Moby Dick is Herman Melville.»

The important fact is that sentence [5] asserts a different state of affairs compared to the one asserted by sentence [4]. According to Renear:

> We refer to this relationship as exhibition. We say that the brief document/utterance ''Moby Dick was written by Herman Melville'' exhibits the state of affairs that ''the name of the author of Moby Dick is 'Herman Melville' '', but it does not assert that state of affairs. What it does assert is that Melville is the author of Moby Dick. Although naming is our prototypical example of exhibition

> in this paper, we believe that exhibition is a widespread and diverse phenomenon. (Renear *et al.*, 2005, p. 176)

The CIDOC-CRM is an ontology made in the museum community in order to model cultural heritage information. The ontology will be described further below. In this article, the following thesis will be discussed: the use of CIDOC-CRM as a modelling tool can assist us in solving the problem of modelling the difference between what is asserted and what is exhibited in utterances such as [4] and [5] above. To look into this, let us first consider an example text for which I claim to be able to solve the problem by using CIDOC-CRM. Next, I will return to the Herman Melville problem. I will try the same method to sentence [4] above and explore whether that problem is also solvable. But before we enter into the modelling, I will briefly discuss three resolutions to the exhibition problem described by Renear.

2.1 Renear's three false resolutions

Renear proposes three possible resolutions to the problem, but also states that all of them are false. These resolutions are the following:

(1) TEI encoding represents features of the text only.
(2) The use of two arcs, i.e. 'The Semantic Web community solution', which will be discussed subsequently.
(3) Exhibition is a special case of presupposition.

Based on the description of our work above, it should be obvious that resolution no. (1) is not an alternative for us. Semantic modelling of the real world on the basis of descriptions in texts is part of our work. The relationships between textual descriptions and features outside of the text have to be examined.

I find it difficult to understand how resolution no. (3) may represent a possible solution. Whether exhibition is a type of presupposition or not does not change the basic problem, i.e. in our case, the use of a text as the source of a naming which is merely exhibited in the text. The problem remains the same if the naming is also presupposed in the text, as long as it is not asserted. Although this is an interesting question, it will not in itself solve the problem. To quote Renear:

> Currently we are undecided on this issue but we note that even if exhibition does turn out to be a form of presupposition that would remove neither the difficulty exhibition creates for conceptual modelling, nor its intellectual significance. In fact it would be a rather substantial finding to determine the matter one way or the other. (Renear *et al.*, 2005, p. 178)

My claim is that resolution no. (2) is not false after all, and below I will discuss how the Conceptual Reference Model (CIDOC-CRM, 2005) will solve a similar problem in my example text. Consequently, the major part of this article is concerned with resolution no. (2).

3 Text and Model—An Observation

When modelling a reading of a text, some of the information placed in the model is, semantically speaking, based on the text itself to the extent that the truth conditions of the text are equal to the truth conditions of the modelled information. Our sentence [1] above would normally be considered to be false if the speaker was a woman or a child, both as part of text and in a model based on the text.

Other types of language usage represents different truth conditions for the text compared to the ones for the model, e.g. sentence [2] below:

[2] He said he was from Paris

This sentence would normally be true even if the man saying it was from another place, as long as he did indeed say he was from Paris. In a model, on the other hand, there will often be good reasons to model the fact [3] based on sentence [2].

[3] x is from Paris.

This fact represents different truth conditions compared to sentence [2]. The model then requires additional information in order to convey this. Below, I will describe the way in which this

additional information is entered into the model, as this will be a part of my suggested solution to Renear's problem, and as there are important similarities between sentences such as [2] and the sentences using names, which I will discuss below. At this point, I will only remark that these are the same types of truth conditions as the ones we find in e.g. 'that-sentences', which are well known from semantics and also discussed by Frege as 'abstrakten Nennsätzen' (Frege, 1892, p. 37 ff.)

When models like these are used to create databases, one of the uses of the database will be to return information based on questions (queries) from a user. Often, the answers from the database will be presented and/or interpreted as answers from the texts on which the database is based. Regarding the examples above, typical questions could be: is the person young or old? Is the person male or female? What is the birthplace of the person? The answers to the first two questions will be based on the truth values of text [1], whereas the third question could not be answered based on truth values in text [2] only.

Because the results returned from such a query system is not only based on the statements asserted by the source texts on which the system is based, but also on statements exhibited by them, it is useful to have a layer above the question–answer mechanism. In this top layer, general rules of what data to base searches on may be defined. This will be seen by a user as similar to options such as 'Give me answers from text A only' or 'Give me answers from all texts in this collection', but will not be based on text choices. It will be based on the way in which the facts in the database are based on the source text. One option may be: 'Base answers only on assertions made by the text' whereas another may be: 'Include assertions made by voices (characters making assertions) in the text exhibited by the text.' It may even be interesting to have a choice between which voices in the text to base answers on, as well as to calculate possible contradictions between the assertions in the system. These may be contradictions between assertions based on the same voice as well as contradictions based on assertions made by different voices.

4 My Example Text

The text used in my example is based on the work of Major Peter Schnitler. In the 1740s, Major Schnitler was appointed by the Danish government to explore the border area between the northern parts of Norway and Sweden/Finland. Significant parts of the text in the manuscript that he handed over to the Danish government consist of transcripts of local court interviews that were carried out by Schnitler in order to gather information about the local population as well as what they had to say about the border areas. The material includes information directly relevant to the border question, as well as general information of the areas in question, which corresponds to similar material collected through work carried out in Europe at the time (Burke, 2000, p. 128 ff.).

The text fragments below are taken from the very first meeting described in the text (English translation from Danish by me):

> [6] Of the Witnesses, supposed to be the most Cunning on the border issue, Were and stood up in the court 1: Ole Larsen *Riise*, [. . .]
> [7] For these the Kingly order was read out loud [. . .] and they gave their Bodily Oath [. . .]
> [8] Question: 1: What his name is? *Answer: Ole Larsen Riise* (Schnitler, 1962, p. 1)

4.1 Readings of the example sentences

In these quotes, we find that several facts are asserted by the text. Sentence [6] claims the existence of a witness. We will call this witness x. Being a witness implies being a person. Thus, x is a person. We may also note that x is referred to by using the name 'Ole Larsen *Riise*.', abbreviated OLR subsequently.

Sentence [7] claims that a group of people, among them the person we call x, made an oath to speak the truth. Sentence [8] conveys that person x claims that his name is OLR. The source of the naming is person *x*, as spoken out loud at a specific place at a specific date in 1742 (the specification of the date is not included in the text fragment above). In this way, the text puts forward an assertion by person x that he is named OLR.

My semantic model of these facts will include the four assertions in Table 1. The assertions are not the only ones that may be read from the text, but they are the most important ones for our modelling purpose. Note the source of each assertion, which is the entity responsible for the assertion. We may say that the source of each assertion is the one guaranteeing that the assertion is true.

4.2 CIDOC-CRM models

CIDOC-CRM is a modelling language, an ontology in the computer science sense of the word, which can be used to describe the implicit and explicit concepts and relationships found in cultural heritage documentation (CIDOC-CRM, 2005). In this case, a small subset of the concepts (entities, called Es) and relationships (called Ps) are used in order to model the reading of the text described above.

Table 1 Assertions from sentences 6–8

Assertion	Source
(1) There is an x who is a witness	The text
(2) x is a person	The meaning of the word 'witness' and 'person' in this context
(3) x gave an oath	The text
(4) OLR is the name of person x	x. Note: This is based on sentence [8]. The reference to his name in sentence [6] is not included in the models in order to keep them as simple as possible.

The CIDOC-CRM is not developed as a tool for the modelling of readings of texts as it is used in this article, but I find it to be a good tool to apply to exercises such as ours.

The model in Fig. 1 contains the basic facts from the first three assertions in Table 1. The facts are focused on the event of making an oath. This event was carried out by a person. Through this event, the oath as a conceptual object was created, and it was documented in Schnitler's text.

In this figure, as well as in the following figures subsequently, the boxes containing names starting with an E represent entities, whereas the ovals with names starting with a P represent the properties linking them together.

The modelling of the naming event from assertion (4) in Table 1 is shown in Fig. 2. The event through which the attribute was assigned (the naming event, a speech act) is an *E13 Attribute assignment* which states that x carried out this particular speech act, as seen in Fig. 2. The part of the model where the responsibility for the naming is shown is the *P14 Carried out by*, which indicates that person x was the one carrying out the naming.

Note also that all of the four assertions can easily be combined in a single CIDOC-CRM model, as seen in Fig. 3. Here we also see clearly in what way both the events are documented in Schnitler's book.

4.3 The assertions in the example sentences revisited

When looking at these two model figures, it is striking to what extent the modelling of the giving

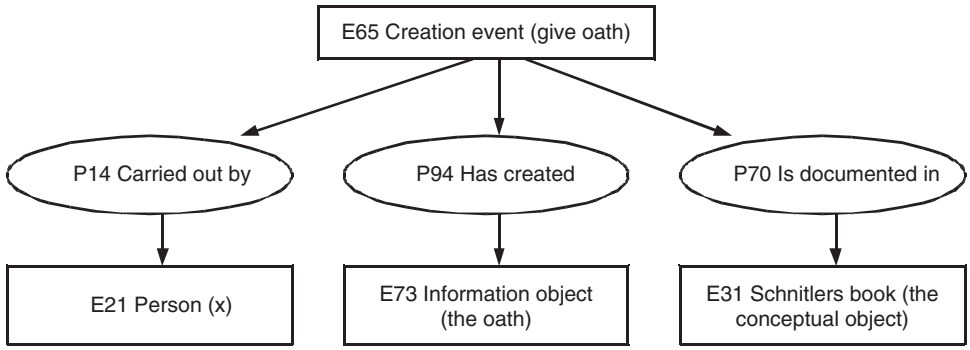

Fig. 1 CIDOC-CRM model based on assertions 1–3 from Table 1

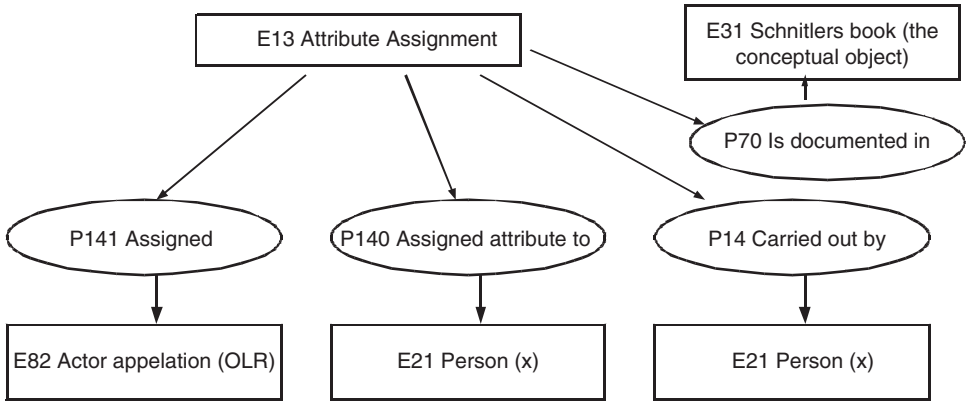

Fig. 2 CIDOC-CRM model based on assertion 4 from Table 1

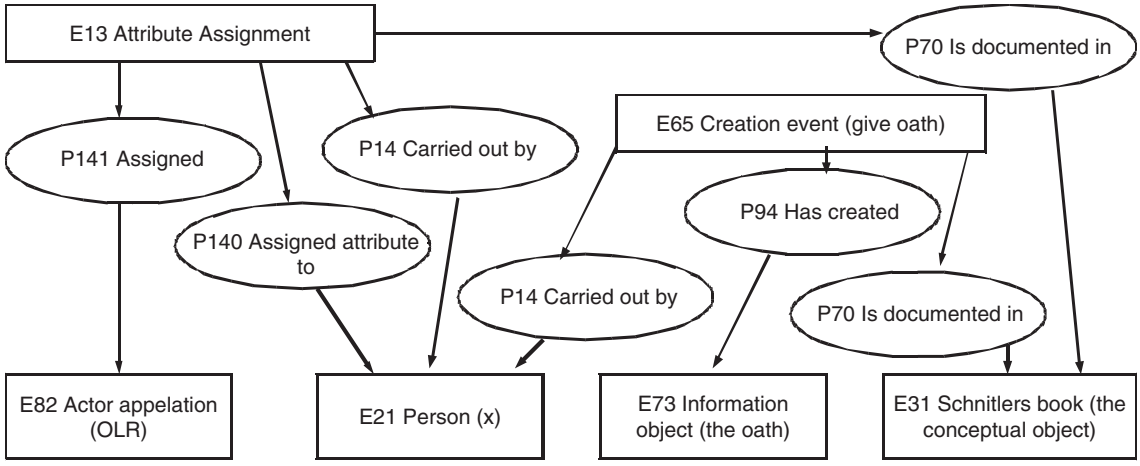

Fig. 3 The CIDOC-CRM models from Figures 1 and 2 combined

of the oath in Fig. 1 compares to the naming of x in Fig. 2. The explanation is that they are similar situations. Our traditional way of reading them made us structure them differently in Table 1, whereas represented in the CIDOC-CRM structure they came out the same in Figs 1 and 2. In order to present explicitly in what way they correspond, note that line 4 from Table 1 above could be rewritten as seen in Table 2.

This is a good example of the way modelling may help us understand a text better. What we have done is to rethink the difference between an event (x gave an oath) and a fact (ORL is the name of x).

Table 2 Assertions from Table 1 with number 4 rewritten

Assertion	Source
(1) There is an x who is a witness	The text
(2) x is a person	The meaning of the word 'witness' and 'person' in this context
(3) x gave an oath	The text
(4) x named himself OLR	The text

In order to model the fact correctly, i.e. to show that it was exhibited rather than asserted in the text, we had to consider it a naming event. Considering

it an event is more feasible in that an event typically has actors who are responsible for the outcome. Further, this makes more sense in that both expressions are speech acts. When it is considered a speech act, the naming event is the same kind of event as the making of an oath.

5 Why Renear's Solution No. 2 is not False After All

In order to be able to see the problem with Renear's solution no. 2, or to realize that the problem is not really there, we have to read his text *in extensio*:

> Another approach, this one anticipated from the Semantic Web community, is simply to insist on an unambiguous corrected conceptual representation: one arc for being named ''Herman Melville'', one for authoring *Moby Dick*. But this resolution fails for the reasons presented in the preceding section. Although this model would be in some sense an accurate representation of ''how the world is'' according to the document, it would not represent what is asserted by the document. The authorship arc in the corrected RDF graph model will correspond to relationships of exhibition, not assertion; and there is no accommodation for this distinction in the modelling language. (Renear *et al*., 2005, p. 178)

In the first couple of sentences of this paragraph, the resolution of using an 'unambiguous corrected conceptual representation' is said to have failed. In the next couple of sentences Renear weakens his statement by saying that only RDF does not accommodate this; 'there is no accommodation for this distinction in *the* modelling language' (my emphasis). There are no arguments to support why a different modelling language could not solve the problem. In fact, the CIDOC-CRM does solve this, by giving the modeller an opportunity to state explicitly who is the source of an assertion, as demonstrated in Fig. 2.

By showing that the problem can be solved in CIDOC-CRM, we have also shown that the problem can be solved in an RDF system. This is because a document from ICS-FORTH shows that CIDOC-CRM can be expressed in RDF Schema (CIDOC-CRM, 2006), which implies that the solution described above can be expressed in RDF.

On the other hand, one has to agree that the distinction between assertion and exhibition is not often spelled out in RDF–in fact, I have never seen it. Indeed, even models in which authors of books seem to be strings of characters rather than human beings are not uncommon. An example of this can be found in Miller's introductory text, where he models the text 'John Smith is the author of Document 1' in a way which gives the impression that the author is the string 'John Smith' and not a human being referred to by this string (Miller, 1998).

What remains is that if my argument in this article is correct, i.e. if the CIDOC-CRM can handle the exhibition problem, then a solution can be made in RDF as well because any CIDOC-CRM model can be expressed in RDF.

6 Generalization

The example described above is quite unusual in that it includes an explicit naming. But it can be argued that all personal names, at least in eighteenth-century Scandinavia, are based on naming events, as people are normally baptized. One may even argue that this is always true, as a name has to be attributed in one way or another in order to be used. As long as we believe that this is the case, we can include into the model an explicit attribute assignment event as the one in Fig. 2 for each name found in the text.

With respect to this naming event, we do not know who carried it out or when it took place. That is not necessarily a problem. In CIDOC-CRM, the modelling of entities we presume to exist without knowing who or what they are is quite possible. And there will always be things we do not know when we read historical texts. What may be a problem is that the naming event we model in this way is an event that is not documented in the text we are basing the model upon. Whether this is acceptable or not is a decision one has to make when building up these kinds of models.

Let us return to the following sentence discussed above:

[4] Moby Dick was written by Herman Melville.

Based on my reading of the sentence, we can extract the assertions shown in Table 3. All the assertions are in some way based on the text, but in some cases more world knowledge is required than in other cases. This is indicated in the source column in Table 3.

Note that the events e1 and e2 are not described in the text; they are included because they are required in order for something to have a name, and because I need them in order to reach a solution to the exhibition problem.

6.1 CIDOC-CRM models

We do not have any direct documentation for event e1 in the text [4], which is the *E13 Attribute Assignment* in Fig. 4. The event is assumed to have taken place because the name is used. The person responsible for the naming, v, could be any legal or physical person. The naming event could be a formal naming, such as a baptism. There may be many events which we could use as e1. The reason for stating that such an event exists is that for a name to be usable in communication, it has to exist when it is used, and in order to exist, it must have come into existence at some point.

V is here not related to any other person, unlike x in the earlier naming event (self-naming). The fact that the person is not related to any other person in the model, and the fact that the naming event is not documented in the text, implies that the naming is exhibited (or presupposed, or both) in the text.

In Fig. 5, a similar *E13 Attribute Assignment* is the event e2 in Table 3. Note that I assume that the person naming the novel Moby Dick is y, the person who wrote the book and who is known by the name Herman Melville. This is a fact I choose to include in the model because I believe it to be true, but I have no source for this, and I could very well be wrong.

And, to complete the modelling, the assertion 'y wrote z' in Table 3 is modelled in Fig. 6. Note that the creation event is documented in sentence [4]. It would now be easy to combine these three models as demonstrated in the combination of the first two models in Fig. 3, but this is not done here.

By including the naming of HM in the model of Fig. 4, and by not stating in the model that the

Table 3 Assertions from sentence 4

Assertion	Source
There exist a person y	[4] and world knowledge
y is named HM	naming event e1
There exist a work z	[4] and world knowledge
z is named MD	naming event e2
y wrote z	[4]
There exist a person v	Extra-textual world knowledge
v named y HM in e1	Extra-textual world knowledge
y named z MD in e2	Extra-textual world knowledge (or just a guess?)

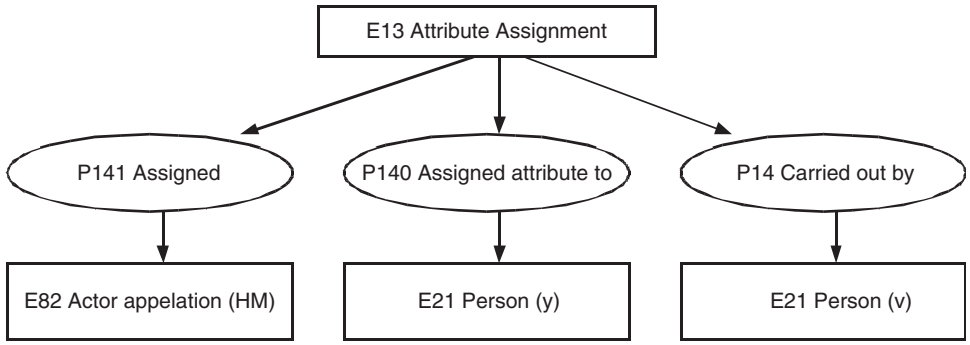

Fig. 4 CIDOC-CRM model of naming event e1, the naming of Melville

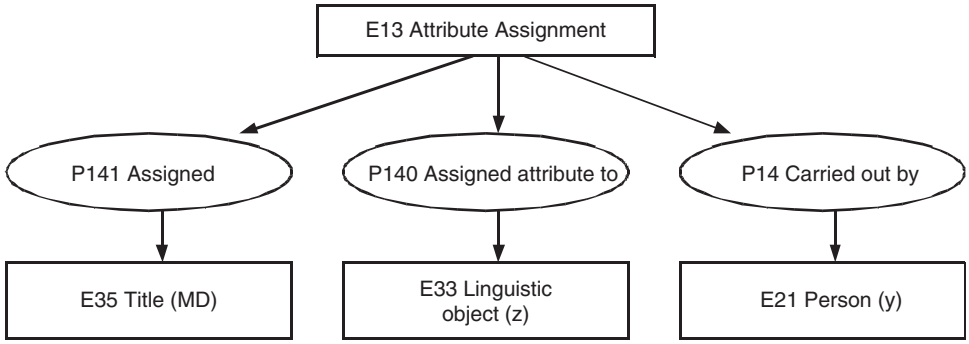

Fig. 5 CIDOC-CRM model of naming event e2, the naming of Moby Dick

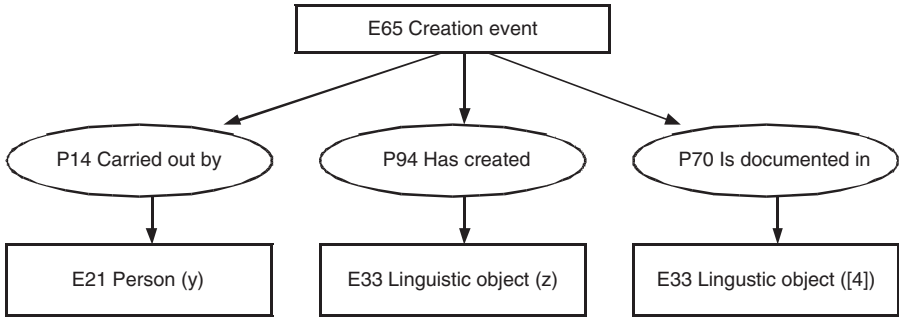

Fig. 6 CIDOC-CRM model of the writing of Moby Dick by Melville

naming is documented in the text, we may infer that:

(1) The naming is known when the text is read, but
(2) The text does not put forward the naming as an assertion.

This is very close to saying that what the model expresses is a presupposition. This then takes us back to Renear's resolution no. (3), i.e. that exhibition is a special case of presupposition.

But even if this is so, i.e. that exhibition is a special case of presupposition, it does not necessarily follow that a model expressing a presupposition also expresses an exhibition, as it does not follow from 'all p's are q' that 'all q's are p'. So the relationship between the model in Fig. 4 and text [4] is, in my opinion, not yet entirely clear. It is not enough to show that we have modelled a presupposition, because it does not necessarily follow from that, that we have modelled an exhibition as well.

With respect to the first part of this article, note also that when a model is built, facts based on assertions in the text will have the same truth value in the text as in the model, whereas when modelling facts that are merely exhibited in the text, a difference in terms of truth values may appear.

This happens when we model facts that are exhibited in the text as assertions in the model. When we do that, it is vital to include the source of the assertion in the model in order to make sure that the text is not considered as evidence for an assertion in the model which is based on something only exhibited in the text. It may be the case that

a similar method could be used in modelling the 'that-sentences' mentioned above.

7 Conclusion

There is reason to believe that the problem described in Renear's paper is an important one. But there is a solution to the problem. I have argued that for one specific type of text, i.e. when somebody explicitly names himself, the problem may be solved by using CIDOC-CRM modelling including explicit statements of the sources of the assertions exhibited in the text. I have also presented this as a possible solution to the more general situation without explicit naming. Further research may disclose whether this solution will apply to other linguistic constructions as well, e.g. 'that-sentences'.

It is worth noting that behind this solution lies the fact that we are using an event-oriented method, since this is an important part of the way CIDOC-CRM works. This is an indication that some of the problems with previous methods is that they do not include the dynamic aspects of reading and understanding a text. It would be interesting to study in which way CIDOC-CRM can be related to semantic theories such as discourse representation theory (Kamp and Reyle, 1993).

In the self-naming modelled in Figs 2 and 3, the CIDOC-CRM model makes it explicit as to who is responsible for the naming, and thereby shows that it is merely reported by the text. Based on this it follows that the naming is exhibited in and/or presupposed by the text.

In the naming of HM modelled in Fig. 4, the model states that somebody is responsible for the naming, but does not relate that person to the text. Based on this it can be argued that the naming is merely reported by the text, and thus that the naming is exhibited in and/or presupposed by the text. But this argument is weaker than the argument above concerning the self-naming.

It would have been an improvement if I could be able to express more clearly the fact that the naming is merely exhibited in the case of the latter naming, but I cannot see how that could be done in a CIDOC-CRM context. This is a problem that requires further research.

The difference between assertion and exhibition has not been modelled directly in any of the two examples. The difference is modelled indirectly by using attribute assignments with explicit responsibility. This method requires further study and development in order to enable me to establish its fidelity.

References

Burke, P. (2000). *A Social History of Knowledge: From Gutenberg to Diderot*. Cambridge: Polity Press.

CIDOC-CRM (2005). *Definition of the CIDOC Conceptual Reference Model*. Produced by the ICOM/CIDOC Documentation Standards Group, continued by the CIDOC CRM Special Interest Group. Version 4.2, June 2005.

CIDOC-CRM (2006). *CIDOC CRM v4.2 Encoded in RDFS*/Created by ICS-FORTH. http://cidoc.ics.forth.gr/rdfs/cidoc_v4.2.rdfs (accessed 27 September 2006).

Eide, Ø. (2003). Putting the dialogue back together: re-creating structure in letter publishing. *Computers and the Humanities*, **37**(1): 65–75.

Frege, G. (1892). Über Sinn und Bedeutung, *Zeitschrift für Philosophie und philosophische Kritik, NF*, 100: 25–50.

Holmen, J. and Uleberg, E. (1996). Getting the most out of it – SGML-encoding of archaeological texts. Paper at the IAAC'96 Iasi, Romania. http://www.dokpro.uio.no/engelsk/text/getting_most_out_of_it.html (accessed 14 November 2005).

Holmen, J., Jordal, E. K. A., Olsen, S. A., and Ore, C. E. (2007). From XML encoded text to objects and events in a CRM compatible database. A case study. Article based on a paper at Beyond the Artifact, CAA 2004, Computer Applications and Quantitative Methods in Archaeology. http://www.edd.uio.no/artiklar/arkeologi/jordal_caa2004.pdf (accessed 7 November 2007).

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic. Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic.

McCarty, W. (2005). *Humanities Computing*. Basingstoke: Palgrave Macmillan.

Miller, E. (1998) 'An Introduction to the Resource Description Framework'. *D-lib magazine*, May 1998.

Renear, A. H., Mylonas, E., and Durand, D. G. (1996). Refining our notion of what text really is: the problem of overlapping hierarchies. In Hockey, S and Ide, N. (eds), *Research in Humanities Computing 4: Selected Papers from the ALLC/ACH Conference, Christ Church Oxford, April 1992*. Oxford: Oxford University Press, pp. 263–80.

Renear, A. H., Lee, J. H., Choi, Y., and Xiang, X. (2005). Exhibition: a problem for conceptual modeling in the Humanities. In: *ACH/ALLC 2005. Conference Abstracts*. 2nd Edition, Victoria: Humanities Computing and Media Centre, University of Victoria, pp. 176–179.

Schnitler, P. (1962). *Major Peter Schnitlers grenseeksaminasjonsprotokoller 1742-1745. Bind 1 [Major Peter Schnitler's border examination protocols 1742–45.* Volume 1*]*. Oslo: Norsk historisk kjeldeskrift-institutt.

Sperberg-McQueen, C. M., Dubin, D., Huitfeldt, C., and Renear, A. H. (2002). Drawing inferences on the basis of markup. In Usdin, B. T. and Newcomb, S. R. (eds), *Proceedings of Extreme Markup Languages 2002*. Montreal, Canada. http://www.idealliance.org/papers/extreme/Proceedings/xslfo-pdf/2002/CMSMcQ01/EML2002CMSMcQ01.pdf (accessed 29 September 2006).

TEI (2002). *TEI P4: Guidelines for Electronic Text Encoding and Interchange*/Sperberg-McQueen, C. M. and Burnard, L. (eds.) Text Encoding Initiative Consortium. XML Version: Oxford, Providence, Charlottesville, Bergen.