ISSN 1744-1986

Technical Report No 2011/ 01

Coherence relations in ontologies

Richard Power

19 January, 2011

Department of Computing
Faculty of Mathematics, Computing and Technology
The Open University

Walton Hall, Milton Keynes, MK7 6AA
United Kingdom

http://computing.open.ac.uk

Coherence relations in ontologies

Richard Power*

1 Introduction

One of our aims in the SWAT project[1] is to produce texts that verbalise ontologies *coherently*. As a step towards achieving this, it is useful to consider the following sub-problem: given an ontology $\mathcal{O}$ and two statements S_1 and S_2 from $\mathcal{O}$, (a) are these statements rhetorically related, and (b) if so, what is their relationship?

(1) Pixie is a dog. Pixie lives in a kennel.
(2) Pixie lives in a kennel. Every dog lives in a kennel.
(3) Pixie is a dog. Every rabbit lives in a burrow.

Consider for example the pairs of statements (1)-(3), taken (we will assume) from the same ontology, but shown here in English rather than in OWL. For pair (1), a plausible judgement is that the sentences are indeed related by EVIDENCE or perhaps CONSEQUENCE, so that we could combine them by saying 'Pixie is a dog and therefore lives in a kennel'. For pair (2), the sentences are also related, the second generalising the first, so that we could say 'Pixie, and more generally every dog, lives in a kennel'. The sentences in pair (3) are unrelated so we would not wish to express them together.

These examples reveal two sources of evidence from which rhetorical relationships can be inferred. First, and most obviously, we can look for repeated references within the statements themselves: thus the sentences in (1) are related because they share the same subject, while the sentences in (2) share the same predicate. Secondly, we can look *elsewhere* in the ontology for statements containing the same terms. Thus for pair (1) it is

*Department of Computing, Open University.

[1]SWAT (Semantic Web Authoring Tool) is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grants G033579/1 (Open University) and G032459/1 (University of Manchester).

OWL statement	**Example**
ClassAssertion(A,I)	Pixie is a corgi
ObjectPropertyAssertion(P,I,J)	Mary owns Pixie
ClassAssertion(ObjectSomeValuesFrom(P,A),I)	Mary owns a corgi
SubClassOf(A,B)	Every corgi is a dog
SubClassOf(A,ObjectHasValue(P,I))	Every corgi likes Mary
SubClassOf(A,ObjectSomeValuesFrom(P,B))	Every corgi likes a person

Table 1: Coverage

relevant that every dog lives in a kennel, while for pair (2) it is relevant that Pixie is a dog. The purpose of this paper is to formalise these two intuitions, apply them systematically to an OWL fragment, and classify the rhetorical relationships that result.

2 Basic coherence model

We begin by constructing a simple model which covers OWL statements based on three axiom functors: *ClassAssertion*, *ObjectPropertyAssertion*, and *SubClassOf*. Elsewhere (Power and Third, 2010) we have shown that these functors cover around 80% of all axioms[2]. The commonest patterns are presented in table 1, along with sample English realisations conforming to ACE and most other verbalisers (Schwitter et al., 2008).

For the axioms considered, we can give a simple uniform semantics in which each statement links two sets, one denoted by the subject, the other by the predicate; the meaning of the statement is that the predicate set contains the subject set. To accommodate individuals within this scheme we can replace them by enumerated classes with only one member (in OWL these can be constructed using the functor *OneOf*). Thus 'Pixie is a corgi' means that the set containing only Pixie is a subset of the set of corgis; 'Mary owns Pixie' means that the set containing only Mary is a subset of the set of things that own Pixie, and so forth. Both statements in a pair can then be reduced to a pair of sets SP, where S is the subject set and P is the predicate set, the structure of the pair being $S_1P_1 - S_2P_2$.[3] With four

[2] These results were obtained through a study of over 200 ontologies and 600,000 axioms.

[3] Note that this semantics is derived from the underlying OWL formulas, and would not be applicable to some sentences in English (e.g., ones expressing existential statements,

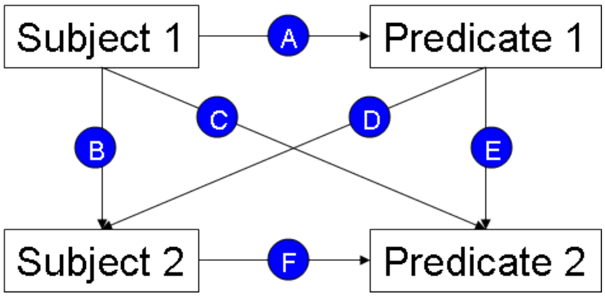

Figure 1: Subject-Predicate relations for two statements

sets we now have six potential relationships to consider: S_1P_1, S_1S_2, S_1P_2, P_1S_2, P_1P_2, and S_2P_2. Two of these (S_1P_1, S_2P_2) correspond to the original statements; the other four may be addressed elsewhere in the ontology, thus providing additional information on whether and how the statements are rhetorically related. The six relationships are shown diagrammatically in figure 1.

The next question is how these relationships among sets should be classified. Among various possibilities, an obvious method is shown in figure 2: given two sets A and B, either A will be narrower than B, or wider, or equal, or distinct, or overlapping. These relations are represented in OWL as follows: (1) narrower by *SubClassOf(A,B)*; (2) wider by *SubClassOf(B,A)*; (3) equal by *EquivalentClasses(A,B)*; (4) distinct by *disjointClasses(A,B)*; and (5) overlapping, implicitly, by absence of the above.

With this model, the rhetorical relationship between two statements can be profiled by six integers describing the six relations; if we assume that subjects are always narrower than predicates, two of these relations (A and F in figure 1) will always be 1. This leaves a potential 5^4 or 625 combinations for the other four relations (B to E in figure 1). However, most of these combinations are contradictory; by writing a brief program that applies consistency constraints[4], we have identified what we believe to be the set of consistent profiles. These are shown in table 2, each accompanied by a simple hand-crafted example[5].

or statements with multiple quantification).

[4]The Prolog code is given in Appendix A.

[5]All statements in the examples describe subclass relations between two atomic classes; no individuals or properties are employed. The set relations assigned to the examples

N	Profile	Example
1	111111	Every corgi is a dog; every mammal is an animal
2	111211	Every corgi is a mammal; every dog is an animal
3	111221	Every corgi is an animal; every dog is a mammal
*4	111231	Every corgi is an animal; every dog is an animal
5	111251	Every pembroke is a pet; every corgi is a dog
*6	111311	Every corgi is a dog; every dog is an animal
7	111511	Every corgi is a dog; every pet is an animal
8	111551	Every corgi is a dog; every pet is a possession
9	121211	Every dog is a mammal; every corgi is an animal
10	121221	Every dog is an animal; every corgi is a mammal
*11	121231	Every dog is an animal; every corgi is an animal
12	121251	Every pet is an animal; every corgi is a possession
13	122221	Every mammal is an animal; every corgi is a dog
*14	123221	Every dog is an animal; every corgi is a dog
15	125221	Every dog is an animal; every corgi is a pet
16	125251	Every dog is an animal; every corgi is a possession
*17	131211	Every corgi is a dog; every corgi is an animal
*18	131221	Every corgi is an animal; every corgi is a dog
*19	131231	Every corgi is a dog; every corgi is a dog
*20	131251	Every corgi is a dog; every corgi is a pet
21	141211	Every corgi is a dog; every poodle is an animal
22	141221	Every corgi is an animal; every poodle is a dog
*23	141231	Every corgi is a dog; every poodle is a dog
24	141251	Every corgi is a dog; every poodle is a pet
25	141411	Every corgi is a dog; every cat is an animal
26	141451	Every corgi is a dog; every siamese is a pet
27	141511	Every corgi is a pet; every cat is an animal
28	141551	Every corgi is a possession; every cat is an animal
29	144221	Every cat is an animal; every corgi is a dog
30	144251	Every siamese is a pet; every corgi is a dog
*31	144441	Every corgi is a dog; every siamese is a cat
32	144451	Every corgi is a dog; every siamese is a pet
33	144551	Every corgi is a pet; every mouse is a rodent
34	145221	Every cat is an animal; every corgi is a pet
35	145251	Every cat is a mammal; every corgi is a pet
36	145451	Every mouse is a rodent; every corgi is a pet
37	145551	Every kitten is a youngster; every corgi is a pet
38	151211	Every corgi is a dog; every puppy is an animal
39	151221	Every corgi is an animal; every puppy is a dog
*40	151231	Every corgi is a dog; every puppy is a dog
41	151251	Every corgi is a pet; every petpuppy is a dog
42	151511	Every corgi is a pet; every puppy is an animal
43	151551	Every corgi is a pet; every puppy is a mammal
44	155221	Every puppy is an animal; every corgi is a pet
45	155251	Every puppy is a mammal; every corgi is a pet
46	155551	Every puppy is a youngster; every corgi is a pet

Table 2: All consistent patterns for inclusion statements

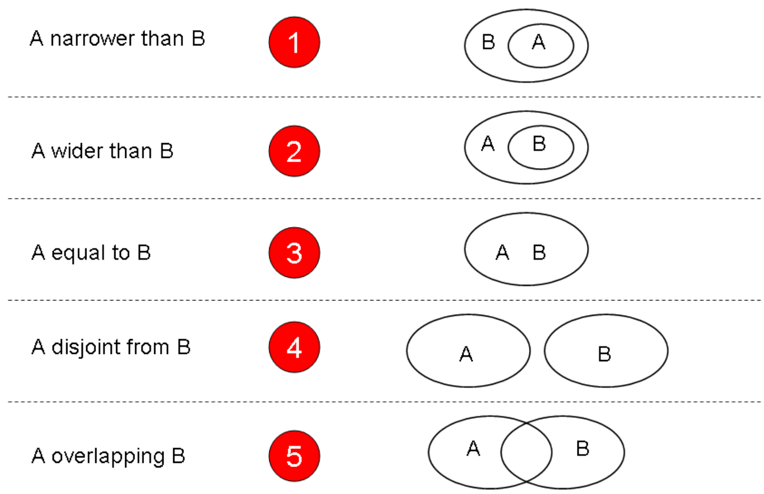

Figure 2: Relations among two sets

As the examples should demonstrate, some of these patterns are more coherent than others. At one extreme, the statements in pattern 46 ('Every puppy is a youngster; every corgi is a pet') have nothing in common apart from a general aura of dogginess; at the other extreme, the statements in 17 ('Every corgi is a dog; every corgi is an animal') are tightly connected and easily associated with familiar logical and rhetorical relationships (e.g., ELABORATION, CONSEQUENCE). After producing these examples, the author went through the list and identified the pairs that seemed coherent, in the sense that it would be natural to present two such statements together; the selected patterns are marked in table 2 with an asterisk. I then gave the original list to two colleagues, both of whom chose exactly the same patterns. On analysis, it turned out that a simple rule explains our selections: we judged a pattern coherent either if the two statements had a set in common (i.e., if the pattern index contained the number 3), or if all relations across statements were disjoint (i.e., 144441).

The eleven patterns selected as coherent are presented again in table 3, with some organisation and renumbering. As will be seen, we follow Sanders, Spooren, and Noordman (1992) in taking account of the *order* of the two

assume a world in which all corgis are pets, and hence possessions, but some dogs are not; mice, like dogs, may or may not be pets. A pembroke is a kind of corgi.

Index	Profile	Example
Group I: Same subject, same predicate		
1	131231	Every corgi is a dog; every corgi is a dog
Group II: Same subject, different predicate		
2	131211	Every corgi is a dog; every corgi is an animal
3	131221	Every corgi is an animal; every corgi is a dog
4	131251	Every corgi is a dog; every corgi is a pet
Group III: Different subject, same predicate		
5	111231	Every corgi is an animal; every dog is an animal
6	121231	Every dog is an animal; every corgi is an animal
7	141231	Every corgi is a dog; every poodle is a dog
8	151231	Every corgi is a dog; every puppy is a dog
Group IV: Different subject, different predicate		
9	111311	Every corgi is a dog; every dog is an animal
10	123221	Every dog is an animal; every corgi is a dog
11	144441	Every corgi is a dog; every siamese is a cat

Table 3: Classification of the coherent patterns

statements in the text, with the result that some patterns (2-3, 5-6, 9-10) are inverses. We think this is appropriate since in all these cases, order has rhetorical significance. This will emerge in the next section, which considers more closely the 11 coherent patterns and their rhetorical implications.

3 Eleven coherence patterns

For each pattern we provide (a) a name, (b) further examples, and (c) some discussion of potential uses, with reference to earlier literature, especially Sanders, Spooren, and Noordman (1992) and the RST (Rhetorical Structure Theory) literature (Mann and Thompson, 1987; Carlson and Marcu, 2001).

3.1 Restatement (131231)

(1a) Every corgi is a dog; every corgi is a dog.
(1b) Pixie is a dog; Pixie is a domestic canine.
(1c) Every domestic canine – that is, every dog – lives in a kennel.
(1d) Mary is a pet-owner; that is, she owns a pet.

In verbalising an ontology there seems no reason to immediately repeat a statement in the same words, even though this rhetorical device is sometimes used in other kinds of discourse for emphasis (e.g., 'The government's priorities will be education, education and education'; 'He passed the exam. He really passed the exam!'). We might however have reason to repeat a statement using different terminology. This can arise when an ontology contains an *EquivalentClasses* statement stating that two classes denote the same set. Such statements are typically used for two purposes: (1) to record that two atomic terms are synonymous (this often happens when two ontologies are merged), and (2) to define an atomic term through a complex expression composed of simpler terms. Here are examples of each.

(1) *EquivalentClasses(Dog, DomesticCanine)*
A dog is defined as a domestic canine.

(2) *EquivalentClasses(PetOwner, ObjectSomeValuesFrom(Owns,Pet))*
A pet-owner is defined as someone who owns one or more pets.

In cases where one synonym is more colloquial or explicit than the other, one could imagine a restatement serving the purpose of explanation. However, without looking beyond the code of an ontology, it is not clear how one can tell which of two equivalent classes will be more readily understood by a reader.

This pattern corresponds to RESTATEMENT in RST, defined as follows in the Discourse Tagging Reference Manual (Carlson and Marcu, 2001):

> A restatement relation is always mononuclear. The satellite and nucleus are of (roughly) comparable size. The satellite reiterates the information presented in the nucleus, typically with slightly different wording. It does not add to or interpret the information.

3.2 Widening Elaboration (131211)

(2a) Every corgi is a dog, and more generally an animal.
(2b) Pixie is a corgi, and therefore a dog.
(2c) Pixie is a corgi, and therefore lives in a kennel.

I will use 'elaboration' in a name only when the subjects of the two statements denote the same individual or set. This departs slightly from the RST literature in which the second statement can also elaborate on the object of the first – see the relation ELABORATION-OBJECT-ATTRIBUTE in Carlson and Marcu (2001). So for widening elaboration, the subjects have the same reference, while the predicate of the second statement includes the predicate of the first. All pairs fitting this pattern are thus implications, since the second statement follows from the first (given the relationship between their predicates).

I can think of two rhetorical purposes for this pattern. One is that the author wishes to draw the second statement as a conclusion — for instance, to establish that Pixie lives in a kennel in example (2c). The other is that the author aims to locate the subject term in a classification scheme, thus using the second statement in order to generalise the first, as in example (2a). The latter move might be called GENERALISATION, signalled for instance by the connective 'more generally'; so far as I know this is not covered either in RST or other approaches.

In RST the widening elaboration pattern would presumably be covered by CONSEQUENCE (or EVIDENCE, I am not sure which); Sanders, Spooren, and Noordman (1992) would classify the relation as 'causal' because they misleadingly use this label to cover both logical implication and causality proper[6].

3.3 Narrowing Elaboration (131221)

(3a) Every corgi is an animal and, more specifically, a dog.
(3b) Pixie is owned by a royal — specifically, by the Queen.

As already pointed out, this is the inverse of pattern 2: the subjects are equivalent, but this time the predicate of the second statement is included in the predicate of the first (thus 'things owned by the Queen' is a subset of 'things owned by a royal'). As a result, the first statement follows from the second, and in parallel with pattern 2 we could posit a CONSEQUENCE relation, marked perhaps by 'because':

Pixie is owned by a royal, because she is owned by the Queen.

Looking for a more specific rhetorical purpose, I think this pattern would often be used for *refinement*: the author makes a generic statement, and then

[6]Perhaps Sanders et al. would treat logical implication as causal-pragmatic as opposed to causal-semantic. This would have to mean that the relation between (for example) 'Pixie is a corgi' and 'Pixie is a dog' is based on their illocutionary force, not their propositional content. Since the two propositions are obviously related irrespective of illocutionary force, this analysis seems implausible, so I assume that in such cases the relation is semantic, and the term 'causal' simply a mistake.

refines it by further narrowing the predicate. Perhaps this is the import of the relation ELABORATION-GENERAL-SPECIFIC in RST, defined by Carlson and Marcu (2001) as follows:

> The satellite provides specific information to help define a very general concept introduced in the nucleus.

3.4 Additive Elaboration (131251)

(4a) Every corgi is a dog, and also a pet.
(4b) Pixie is a corgi, and is also owned by a royal.

In the third elaboration pattern, the predicate terms overlap: thus some dogs are pets, some dogs are not pets, and some pets are not dogs. Again, some corgis are owned by royals, some are not, and royals own other things such as horses and palaces. I am tempted to call this pattern 'Overlapping Elaboration', but for consistency with Sanders, Spooren, and Noordman (1992) I have preferred 'Additive': the second statement provides further consistent information, without implying or being implied by the first.

The nearest counterpart in RST is ELABORATION-ADDITIONAL, except that in our scheme no distinction in importance can be drawn between the first statement and the second. In Sanders, Spooren, and Noordman (1992), the relation would be classified as additive-semantic-positive, and hence LIST.

3.5 Widening Comparison (111231)

(5a) Every corgi, and more generally every dog, is an animal.
(5b) Pixie is a dog, because every corgi is a dog.
(5c) Pixie lives in a kennel; so does every dog.

I will use 'comparison' in a name only when the statements have equivalent predicates. Thus in a widening comparison, the predicates corefer, while the subject of the second statement includes the subject of the first, with the result that the first statement follows from the second.

Why would one use widening comparison? One reason could be to establish the first statement without misleadingly implicating that it applied *only* to the subject referent. For instance, if we say 'Every corgi lives in a kennel', people might wonder what is so special about corgis as opposed to other kinds of dog — could it be that poodles *do not* live in kennels? Hence

‘Every corgi, indeed every dog, lives in a kennel’, which makes everything clear.

I cannot find any clear counterpart in the literature: the RST relation of COMPARISON in (Carlson and Marcu, 2001) comes closest. Of course, since the second statement implies the first, this could be classified as a CONSEQUENCE in the Sanders, Spooren, and Noordman (1992) scheme, or more precisely as causal-semantic-nonbasic-positive, with the usual qualifications regarding ‘causal’. ‘Nonbasic’ here refers to the order of two statements that have an implication relation; the ordering in which the first statement implies the second is regarded as basic, and the inverse thus nonbasic.

3.6 Narrowing Comparison (121231)

(6a) Every dog, and more specifically every corgi, is an animal.
(6b) Every dog, and thus every corgi, lives in a kennel.
(6c) Every corgi, including Pixie for example, likes the Queen.

This is the inverse of pattern 5: the predicates again corefer, but this time the subject of the second statement is included in the subject of the first. As a result, the implication follows the ‘basic’ order (just mentioned) in which the first statement implies the second. Accordingly, one purpose of the pattern would simply be to establish the second statement, by deducing it from the first (hence causal-semantic-basic-positive).

Another reason for using this pattern might be to exemplify the first statement. This would be appropriate in particular when the second statement is more familiar to the reader than the first, as in ‘Every amphibian — every frog for example — is cold-blooded’.

The nearest counterparts in RST are thus CONSEQUENCE (or EVIDENCE), covering the implication relationship, and also EXAMPLE; the latter so far as I know is not mentioned by Sanders, Spooren, and Noordman (1992).

3.7 Disjoint Comparison (141231)

(7a) Every corgi is a dog; so is every poodle.
(7b) Pixie, like Growler, is owned by the Queen.

In this pattern the predicates corefer and the subjects denote disjoint sets; the pattern therefore has no inverse (i.e., if we reverse the order of the statements, the relation remains disjoint comparison). Such a combination would plausibly be found in a passage comparing two distinct things – corgis and poodles let us say – to bring out their resemblances and differences; this kind of passage has been studied in particular by Milosavljevic (1997).

So far as I can tell, this pattern is not distinguished from other comparisons in RST, although perhaps it should be — for instance, 'Every dog, like every corgi, lives in a kennel' sounds odd since 'dog' and 'corgi' are not distinct.

3.8 Additive Comparison (151231)

(8a) Every corgi is a dog; every puppy is a dog.
(8b) Every musician listens to Pavarotti; so does every Italian.

In this pattern the predicates corefer and the subjects denote overlapping sets, hence 'additive'. I find this a relatively weak rhetorical relationship — probably the most dubious of the eleven that I selected. Perhaps it would normally be found when the focus of interest lies within the *predicate*; thus in (8b) we are interested in who listens to Pavarotti, not so much in musicians or Italians.

3.9 Forward Reasoning (111311)

(9a) Every corgi is a dog; every dog is an animal.
(9b) Pixie is a royal corgi; every royal corgi meets the Queen.

Here the predicate of the first statement is equivalent to the subject of the second, so that we obtain a chain from subclass to class to superclass (e.g., corgi-dog-animal). The two statements therefore seem to build up to the conclusion that the subclass is included in the superclass (e.g., 'Every corgi is an animal', or 'Pixie meets the Queen') — hence 'reasoning'.

There seems no clear counterpart in RST, the nearest being ELABORATION-OBJECT-ATTRIBUTE in cases where the first statement uses the copular. More generally, RST lacks any relation indicating that two propositions are co-premises — i.e., that they combine in order to support a conclusion.

3.10 Backward Reasoning (123221)

(10a) Every dog is an animal; every corgi is a dog.
(10b) Every royal corgi meets the Queen. Pixie is a royal corgi.
(10c) Every amphibian is cold-blooded. A frog, for example, is an amphibian.

Backward reasoning is the inverse of forward reasoning (i.e., the order of statements is reversed), so that the subject of the first statement is equivalent to the predicate of the second. As before, the two statements build up to a conclusion, but with a less clear chain of inference owing to the backward order. Why would anyone want to say (10b) rather than (9b)? One reason might be topic flow: the author might have been talking about royal corgis, and wish to switch the focus to Pixie (e.g., to elaborate on a particular occasion on which Pixie met the Queen).

3.11 Contrast (144441)

(11a) Every corgi is a dog; every siamese is a cat.
(11b) Mozart lived in the 18th century, Wagner in the 19th.

In the final pattern, the subjects denote disjoint sets, and so do the predicates. This is the only selected pattern in which there is no equivalence relation across the two statements. Like pattern 7 (disjoint comparison), it is typically found in passages that compare two distinct but related individuals or classes (Milosavljevic, 1997), with disjoint comparison specifying resemblances and contrast specifying differences.

The closest relation in RST (unsurprisingly) is CONTRAST, although the definition given here is more specific. The following example of CONTRAST from Carlson and Marcu (2001) clearly fits the 144441 pattern:

> [But from early on, Tigers workers unionized,] [while Federals never have.]

We have four sets here — S_1=Tiger's workers, S_2=Federal's workers, P_1=people that unionized, P_2=people that did not unionize — with S_1 disjoint from S_2 and P_1 from P_2.

Perhaps our definition of contrast is actually too broad, since it does not require any link between the subjects (or predicates) except that they are disjoint. An example like 'Every corgi is a dog; every prime number is an integer' would therefore fit the pattern, despite the absurdity of comparing corgis with prime numbers. This suggests that our classification of relationships between sets needs some refinement — we might, for instance, decide to limit the definition of contrast to cases in which the disjoint subjects are siblings in the taxonomy.

4 Coverage of RST relations

The last section looked at coherence patterns based on set relations derivable from DL-based ontologies, and tried to link them to rhetorical relations as classified for example in RST. I now approach the problem from the other direction by listing the main RST relations, and considering which are relevant in a discourse based on an ontology, and which are not.

RST relations are classified as either presentational, informational, or multinuclear. For presentational and informational relations, the relation applies to two statements of different importance, a nucleus (N) and a satellite (S). 'Presentational' and 'informational' correspond to 'pragmatic' and 'semantic' in Sanders, Spooren, and Noordman (1992). Relation names and definitions are taken from Mann and Taboada's RST website[7].

4.1 Presentational relations

Antithesis *N and S contrast, writer favours N not S.* Probably inapplicable.

Background *S improves comprehension of N.* This might be assigned when N employs a defined class explained by S. Thus: 'Mary is a pet-owner. A pet-owner is someone who owns one or more animals.' This would belong to the pattern 111333, which we did not consider in the last section since in the second statement subject and predicate are equivalent (S_2P_2=3, not 1).

Concession *N is affirmed although unexpected given S.* We could assign this relation to the pattern 131241 (disjoint elaboration), which was omitted from table 2 (and hence from the last section) because it is logically inconsistent — for instance, 'Pixie is a corgi even though she is a cat'. A consistent use of this relation would be possible for the pattern 131251 (additive elaboration) if we knew that the overlap between the two predicates was *small*, relative to either predicate. For instance, suppose there exist a few vegetarian dogs, even though nearly all dogs are non-vegetarians and nearly all vegetarians are non-dogs. Armed with this knowledge, we could assign CONCESSION to the 131251 pair 'Pixie is a vegetarian, even though she is a dog'.

[7] `http://www.sfu.ca/rst/01intro/intro.html`

Enablement *Understanding S helps reader perform action in N.* Inapplicable.

Evidence *Knowing S increases belief in N.* Could be assigned to any of the patterns in which one statement implies the other, such as widening elaboration (e.g., 'Mary owns a cat, so she is a pet-owner'). However, it is not clear how a nucleus-satellite distinction can be made here.

4.2 Informational Relations

Circumstance *S sets a framework for interpreting N, where N and S are about the same situation.* Probably inapplicable.

Condition *Realisation of N depends on S.* Inapplicable.

Elaboration *S provides additional detail to N.* Our model covers several kinds of elaboration, although without providing a criterion for distinguishing nucleus from satellite; however, this is not especially clear in RST either.

Evaluation *S provides assessment of N.* Inapplicable.

Interpretation *S provides interpretation of N, relating it to some other framework.* Inapplicable.

Cause *S is a cause (volitional or nonvolitional) of N.* Inapplicable if we have to rely entirely on generic features of OWL, which lacks any notion of causality.

Result *N is cause of S.* See above.

Purpose *S is realised through N.* See above.

Restatement *N restates S in a more relevant way (but of approximately equal length).* This corresponds quite closely to our 131231 pattern, with the usual qualification that we have not yet proposed any criterion for distinguishing nucleus from satellite. It is also hard to give any meaning (for ontologies) to 'in a more relevant way'.

Solutionhood *N presents a solution to problem N.* Inapplicable.

Summary *S restates N more briefly.* This could be assigned to a restatement (pattern 131231) in which the same predicate set is expressed more concisely in the second statement. For instance, if the ontology contains the definition 'A pet-owner is a person who owns one or more animals', then SUMMARY might be assigned to 'Mary is a person who owns one or more animals; in short, she is a pet-owner'.

4.3 Multinuclear Relations

Conjunction *Two related Ns of equal importance.* This is so weak that it could be assigned to any of our patterns (except perhaps contrast).

Contrast *Two statements with similarities and differences.* See pattern 11 (144441).

Disjunction *Statements are alternatives.* This cannot be assigned as a relation between two OWL statements. Single statements might contain disjunctions in the form of classes constructed using *ObjectUnionOf*, as in 'Pixie is a corgi or a poodle'; this could be construed as a disjunction of two statements ('Pixie is a corgi or Pixie is a poodle'), in which case DISJUNCTION could be assigned.

Sequence *Conjunction with a sequencing criterion such as order in time.* The problem would be to find an ordering criterion. This might perhaps be done for statements using numerical data values (e.g., 'Hydrogen has a valency of 1; Oxygen has a valency of 2; ...').

4.4 Discussion

Our model covers nearly half these relations, along with some others that are often used although not listed on the RST website (e.g., COMPARISON, EXAMPLE). The main gaps in coverage are as follows:

- At present OWL has no treatment of time, or causality, or probability, or propositional attitudes (belief, volition, etc.). This prevents us from assigning relations like CAUSE, UNTIL, CONCESSION, ENABLEMENT, EVALUATION which depend on these factors.

- Even when we can approximate to presentational or informational relations in RST, we cannot draw the nucleus-satellite distinction. Whether this is an important limitation is unclear — there is no N-S distinction in Sanders, Spooren, and Noordman (1992), for instance.

Note that although OWL itself has no treatment of time, causality, probability etc., ontology developers often have to model these phenomena, and so conventional treatments (sometimes called Ontology Design Patterns) have begun to emerge. With sufficient consensus, these could provide a basis for widening the model given here.

5 Conclusion

I have proposed a method for assigning a rhetorical relation to a pair of OWL statements. Briefly, statements are interpreted as subject and predicate sets, and the six set relationships within and across the two statements are classified to derive a 'coherence pattern'. Of 46 consistent patterns with subjects included in predicates, we selected eleven which were considered coherent by three different judges. We found that these patterns were related to traditional rhetorical concepts like consequence, elaboration, comparison, contrast, and example, although we were not able to assign relations depending on causality, time, or probability. The merit of the proposed method is that it makes use of the kind of information that an OWL ontology *does* provide — namely, statements *elsewhere* in the ontology that are relevant to the pair of statements under consideration.

I think there are two ways in which this model could be extended, while continuing to rely only on the logical functors of OWL (and not for instance on common domain patterns). First, there might be information within an ontology on the *importance* of a statement. One criterion could be whether the statement is about a core concept, using some generic notion of centrality like the number of statements employing the concept. Another could be whether the statement is strong or weak, as measured by the distance in the taxonomy between the subject and predicate classes: thus 'Pixie is a corgi' is stronger than 'Pixie is an animate object', because the latter embraces a chain that might include corgi, dog, canine, mammal, and animal. Second, the patterns presented here assume that the predicate of a statement includes the subject (i.e., that $SP = 1$, whereas theoretically any of the five relationships is possible:

$SP = 1$: Every corgi is a dog
$SP = 2$: Only dogs are corgis
$SP = 3$: A dog is a domestic canine
$SP = 4$: No corgi is a poodle
$SP = 5$: Some dogs are pets (and some are not)

We cannot express the last of these in OWL, but it would be worth extending the coverage to statements with $SP = 3$ or $SP = 4$, which correspond to the OWL functors *EquivalentClasses* and *DisjointClasses*, both commonly used.

References

Carlson, Lynn and Daniel Marcu. 2001. Discourse tagging manual. Technical report, ISI Tech Report ISI-TR-545.

Mann, W. and S. Thompson. 1987. Rhetorical structure theory: a theory of text organization. In L. Polyani, editor, *The structure of discourse*, Norwood, NJ. Ablex.

Milosavljevic, Maria. 1997. Augmenting the user's knowledge via comparison. In *In Proceedings of the 6th International Conference on User Modelling*, pages 119–130.

Power, Richard and Allan Third. 2010. Expressing owl axioms by english sentences: dubious in theory, feasible in practice. In *Proceedings of the 23rd International Conference on Computational Linguistics.*

Sanders, T., W. Spooren, and L. Noordman. 1992. Toward a taxonomy of coherence relations. *Discourse Processes*, (15):1–35.

Schwitter, Rolf, Kaarel Kaljurand, Anne Cregan, Catherine Dolbear, and Glen Hart. 2008. A Comparison of three Controlled Natural Languages for OWL 1.1. In *4th OWL Experiences and Directions Workshop (OWLED 2008 DC)*, Washington, 1–2 April.

A Program

```
% subject.pl Generate all subject-predicate combinations for two statements

% Written by Richard Power
% Copyright (c) 2010 by Open University

%================================================

:- module(subject, [run/0]).

:- use_module(library(lists)).
:- use_module(library(ordsets)).
:- use_module(library(clpfd)).
:- use_module(save, [saveList/2]).

% Module code

% We have two sentences 1 and 2, each composed of a subject and
% predicate, interpreted as set-denoting expressions. This means
% we have four sets S1, P1, S2, P2, with the constraints that
% S1 is narrower or equal to P1, and S2 is narrower or equal to
% P2. Can we enumerate every possible set relationship between
% S1-S2, P1-P2, S2-P1 and S1-P2? The relationships are narrower,
% wider, equal, distinct, or partially overlapping. This looks like
% a CSP.
%
% 1 A is narrower than B
% 2 A is wider than B
% 3 A is equal to B
% 4 A is distinct from B
% 5 A partially overlaps B

run :-
  findall(Solution, generateSolution(Solution), Solutions),
  length(Solutions, N), write(N), nl,
  filter(Solutions, Filtered),
  reorder(Filtered, Ordered),
  length(Filtered, F), write(F), nl,
  saveList('subject.txt', Ordered).

generateSolution(Solution) :-
  initialise(Solution, Vars),
  constrain(Vars),
  labeling([], Vars).

initialise([rel(s1-p1, S1P1),
            rel(s1-s2, S1S2),
            rel(s1-p2, S1P2),
            rel(p1-s2, P1S2),
            rel(p1-p2, P1P2),
            rel(s2-p2, S2P2)], [S1P1,S1S2,S1P2,P1S2,P1P2,S2P2]) :-
  S1P1 in 1..3, S1P1 #\= 2,
  S1S2 in 1..5,
  S1P2 in 1..5,
```

```
  P1S2 in 1..5,
  P1P2 in 1..5,
  S2P2 in 1..3, S2P2 #\= 2.

constrain([S1P1,S1S2,S1P2,P1S2,P1P2,S2P2]) :-
  triangle(S1P1, P1S2, S1S2),
  triangle(S1P1, P1P2, S1P2),
  triangle(S1S2, S2P2, S1P2),
  triangle(P1S2, S2P2, P1P2).

triangle(AB, BC, AC) :-
% AB BC => AC
  (AB #= 1) #/\ (BC #= 1) #=> (AC #= 1),
  (AB #= 1) #/\ (BC #= 3) #=> (AC #= 1),
  (AB #= 1) #/\ (BC #= 4) #=> (AC #= 4),
  (AB #= 2) #/\ (BC #= 1) #=> (AC #\= 4),
  (AB #= 2) #/\ (BC #= 2) #=> (AC #= 2),
  (AB #= 2) #/\ (BC #= 3) #=> (AC #= 2),
  (AB #= 3) #/\ (BC #= 1) #=> (AC #= 1),
  (AB #= 3) #/\ (BC #= 2) #=> (AC #= 2),
  (AB #= 3) #/\ (BC #= 3) #=> (AC #= 3),
  (AB #= 3) #/\ (BC #= 4) #=> (AC #= 4),
  (AB #= 3) #/\ (BC #= 5) #=> (AC #= 5),
  (AB #= 4) #/\ (BC #= 1) #=> (AC #\= 2) #/\ (AC #\= 3),
  (AB #= 4) #/\ (BC #= 2) #=> (AC #= 4),
  (AB #= 4) #/\ (BC #= 3) #=> (AC #= 4),
  (AB #= 4) #/\ (BC #= 5) #=> (AC #\= 2) #/\ (AC #\= 3),
  (AB #= 5) #/\ (BC #= 1) #=> (AC #= 1) #\/ (AC #= 5),
  (AB #= 5) #/\ (BC #= 2) #=> (AC #= 4) #\/ (AC #= 5),
  (AB #= 5) #/\ (BC #= 3) #=> (AC #= 5),
  (AB #= 5) #/\ (BC #= 4) #=> (AC #\= 1) #/\ (AC #\= 3),

% AB AC => BC
  (AB #= 1) #/\ (AC #= 1) #=> (BC #\= 4),
  (AB #= 1) #/\ (AC #= 2) #=> (BC #= 2),
  (AB #= 1) #/\ (AC #= 3) #=> (BC #= 2),
  (AB #= 1) #/\ (AC #= 4) #=> (BC #\= 1) #/\ (BC #\= 3),
  (AB #= 1) #/\ (AC #= 5) #=> (BC #= 2) #\/ (BC #= 5),
  (AB #= 2) #/\ (AC #= 1) #=> (BC #= 1),
  (AB #= 2) #/\ (AC #= 3) #=> (BC #= 1),
  (AB #= 2) #/\ (AC #= 4) #=> (BC #= 4),
  (AB #= 2) #/\ (AC #= 5) #=> (BC #\= 2) #/\ (BC #\= 3),
  (AB #= 3) #/\ (AC #= 1) #=> (BC #= 1),
  (AB #= 3) #/\ (AC #= 2) #=> (BC #= 2),
  (AB #= 3) #/\ (AC #= 3) #=> (BC #= 3),
  (AB #= 3) #/\ (AC #= 4) #=> (BC #= 4),
  (AB #= 3) #/\ (AC #= 5) #=> (BC #= 5),
  (AB #= 4) #/\ (AC #= 1) #=> (BC #\= 2) #/\ (BC #\= 3),
  (AB #= 4) #/\ (AC #= 2) #=> (BC #= 4),
  (AB #= 4) #/\ (AC #= 3) #=> (BC #= 4),
  (AB #= 4) #/\ (AC #= 5) #=> (BC #\= 2) #/\ (BC #\= 3),
  (AB #= 5) #/\ (AC #= 1) #=> (BC #= 1) #\/ (BC #= 5),
  (AB #= 5) #/\ (AC #= 2) #=> (BC #\= 1) #/\ (BC #\= 3),
  (AB #= 5) #/\ (AC #= 3) #=> (BC #= 5),
  (AB #= 5) #/\ (AC #= 4) #=> (BC #\= 1) #/\ (BC #\= 3),
```

```
% BC AC => AB
  (BC #= 1) #/\ (AC #= 2) #=> (AB #= 2),
  (BC #= 1) #/\ (AC #= 3) #=> (AB #= 2),
  (BC #= 1) #/\ (AC #= 4) #=> (AB #= 4),
  (BC #= 1) #/\ (AC #= 5) #=> (AB #\= 1) #/\ (AB #\= 3),
  (BC #= 2) #/\ (AC #= 1) #=> (AB #= 1),
  (BC #= 2) #/\ (AC #= 2) #=> (AB #\= 4),
  (BC #= 2) #/\ (AC #= 3) #=> (AB #= 1),
  (BC #= 2) #/\ (AC #= 4) #=> (AB #\= 2) #/\ (AB #\= 3),
  (BC #= 2) #/\ (AC #= 5) #=> (AB #= 1) #\/ (AB #= 5),
  (BC #= 3) #/\ (AC #= 1) #=> (AB #= 1),
  (BC #= 3) #/\ (AC #= 2) #=> (AB #= 2),
  (BC #= 3) #/\ (AC #= 3) #=> (AB #= 3),
  (BC #= 3) #/\ (AC #= 4) #=> (AB #= 4),
  (BC #= 3) #/\ (AC #= 5) #=> (AB #= 5),
  (BC #= 4) #/\ (AC #= 1) #=> (AB #= 4),
  (BC #= 4) #/\ (AC #= 2) #=> (AB #\= 1) #/\ (AB #\= 3),
  (BC #= 4) #/\ (AC #= 3) #=> (AB #= 4),
  (BC #= 4) #/\ (AC #= 5) #=> (AB #\= 1) #/\ (AB #\= 3),
  (BC #= 5) #/\ (AC #= 1) #=> (AB #\= 2) #/\ (AB #\= 3),
  (BC #= 5) #/\ (AC #= 2) #=> (AB #= 2) #\/ (AB #= 5),
  (BC #= 5) #/\ (AC #= 3) #=> (AB #= 5),
  (BC #= 5) #/\ (AC #= 4) #=> (AB #\= 2) #/\ (AB #\= 3).

% Filter out solutions for which the statements have no term in common

filter([], []).
filter([S|Ss], [NewS|NewSs]) :- relevant(S, NewS), filter(Ss, NewSs), !.
filter([_|Ss], NewSs) :- filter(Ss, NewSs).

relevant([rel(s1-p1,A),rel(s1-s2,B),rel(s1-p2,C),rel(p1-s2,D),
          rel(p1-p2,E),rel(s2-p2,F)], A-B-C-D-E-F) :-
  (memberchk(3, [B,C,D,E]); [B,C,D,E]=[4,4,4,4]).

% Order by S1P1 and S2P2

reorder(Solutions, Numbered) :-
  recode(Solutions, Recode),
  sort(Recode, Sorted),
  replace(Sorted, Ordered),
  renumber(1, Ordered, Numbered).

recode([], []).
recode([S|Ss], [NewS|NewSs]) :- S=A-B-C-D-E-F, NewS=A-F-B-C-D-E, recode(Ss, NewSs).

replace([], []).
replace([S|Ss], [NewS|NewSs]) :- S=A-F-B-C-D-E, NewS=A-B-C-D-E-F, replace(Ss, NewSs).

renumber(_, [], []).
renumber(N, [X|Xs], [(N=X)|NewXs]) :- M is N+1, renumber(M, Xs, NewXs).
```