Reasoning about Topics: towards a formal theory

Robert Demolombe & Andrew J.I.Jones

Abstract

It is suggested that one way of understanding **contexts** is to view them as **sets of topics**. Accordingly, a description of a given context would be a set of sentences about the set of topics concerned. The paper presents a model-theoretic and axiomatic characterisation of a logic for sentences of the type: " sentence 'p' is about topic t", and then discusses some features of that logic. The model-theory makes particular use of the three valued semantics of Bochvar. Finally, some remarks are made about application domains in which it would be useful to have a logic for reasoning about topics.

1 Introduction

What is a context? There may be several ways of approaching this question, but one plausible answer is surely this: that a context is a set of topics. For instance, when a decision is taken in a particular context -such as a decision of to re-introduce capital punishement in the context of rising crime-rates- the decision is made in the light of, and with reference to, the cluster of topics to which the proposition expressed by “rising crime-rate” refers. Those topics set the frame within which the decision-maker deliberates; they represent the relevant class of factors he/she is to take into account.

On this view, then, the sentences which describe a context will be sentences **about** some particular set of topics; and so to understand this view better, it will be important to investigate the logic of such sentences as “sentence ’p’ is about topic t”. This paper makes some preliminary steps toward the formulation of a logic for sentences of that kind.

Three fundamental assumptions guide our approach in the remainder of the paper:

(i) Sentences of the type “sentence ’p’ is about topic t” are shorthand for “the proposition expressed by the sentence ’p’ is about topic t”, where the notion of proposition is interpreted in a way which fits our use of a three-valued semantics of the Bochvar’s kind [1]. In this interpretation a proposition can be viewed as a set of partially defined worlds, that is, a set of worlds where only a given set of propositional variables are defined and where the proposition is true in this particular semantics. We shall use the notation ’p’ when the sentence p expresses a proposition in this three-valued semantics, and we shall use p when it expresses a proposition in classical two valued logic. We shall suppose that any given proposition concerns a set of topics. Our basic move is to claim that the sentence ’p’ is about some topic t provided that the following condition is met: t must be a member of that set of topics with which the proposition expressed by the sentence named by “p” is concerned.

(ii) Suppose that two sentences are equivalent in classical propositional calculus, and that they contain just the same atoms. Then we shall suppose that the one is about a given topic t if and only if the other is also about t. But note the proviso concerning “contains the same atoms”. For instance, although the following is a truth of classical propositional logic:

$$p \leftrightarrow (p \wedge (q \vee \neg q))$$

we are not prepared to allow that the sentences on the left and right of the equivalence sign are necessarily about the same topic.

(iii) We shall assume that the set of topics a **compound** sentence is about is not independent of its **mode of composition**. Consider, for instance, the following two sentences: ’The people of Greece want to call their country “Macedonia”’, and 'The people of Macedonia want to call their country "Macedonia"'. While neither of these two sentences, considered on its own is about the topic **conflict**, their conjunction may well be. The conjunctive mode of combining the two sentences may itself alter the class of topics concerned.

2 Syntax and Semantics

In this section we first present formal definitions of the language, of the models, and of the truth-conditions. We then discuss several additional properties which serve to give a richer structure to a set of sentences that are about some topic.

2.1 Formal definition of the logic

Let CPC be the language of classical propositional calculus. Let σ be a set of names for sentences of CPC. Let τ be a set of names of topics. Let A(-,-) be a sorted binary predicate whose first argument is of sort name of topic (τ), and whose second argument is of sort name of sentence (σ).

The language L is defined by the following rules:

Any atomic sentence of CPC, and any ground atom of the form A(x,y) are sentences of L. (The arguments x, y must of course satisfy the restrictions as to sort mentioned above.)

If p and q are any sentences of L, then $\neg$p and p $\vee$ q are sentences of L.

The logical connectives $\wedge$, $\rightarrow$ and $\leftrightarrow$, are introduced in the usual way, in terms of negation and disjunction.

All the sentences of L are defined by the above rules.

Notation: "p" names the sentence occuring between the quotation marks in "p". Intuitively, we read sentences of the form A(t,'p') as: "the sentence 'p' is about the topic t".

Some illustrative examples of sentences of L:

$\mathrm{A(t,'p') \wedge (p \rightarrow q)}$, "the sentence 'p' is about t and if p then q";

$\mathrm{A(t,'p \wedge q') \rightarrow A(t,'q \wedge p')}$, " if the sentence 'p $\wedge$ q' is about t, then the sentence 'q $\wedge$ p' is about t";

$\mathrm{A(t,'p') \rightarrow A(t','p')}$, "if the sentence 'p' is about t, then the sentence 'p' is about t' ".

2.2 Models

A model M of the logic is a tuple
$\mathrm{M} =< \mathrm{W, I, J}, \mathcal{T}, \mathcal{S}, \mathrm{N, T, F} >$ where:

- W is a set of worlds.
- I is a function that assigns to each topic name in L an element of $\mathcal{T}$.
- J is a function that assigns to each sentence name in L an element of $\mathcal{S}$.

- $\mathcal{T}$ is a set of topics.
- $\mathcal{S}$ is the set of sentences in CPC.
- N is a function which assigns to pairs of sets of worlds sets of topics. The definition domain of N is:
  $2^{W} \times 2^{W} \rightarrow 2^{\mathcal{S}}$.
- T is a function which assigns to each atom in CPC a set of worlds.
- F is a function which assigns to each atom in CPC a set of worlds.

We impose on M the constraint that the name "p" is interpreted by the function J as the sentence p itself.

We also impose on the models M the constraint: $T(p) \cap F(p) = \emptyset$. The intuition is that, in a given world, a sentence p can be true or false or neither true nor false but it cannot be both true and false.

The functions T and F are extended to any sentence in L by the following rules:

$$T(\neg p) = F(p)$$
$$F(\neg p) = T(p)$$

$$T(p \vee q) = (T(p) \cap D(q)) \cup (T(q) \cap D(p))$$
$$F(p \vee q) = F(p) \cap F(q)$$

where D(p) is an abbreviation of $T(p) \cup F(p)$.

The functions T and F define a 3-valued logic for sentences in CPC. This logic was initially defined by Bochvar in [1]. Independently of our work the same logic has also been used by Buvač, Mason and Mc Carty in [2, 15] to characterize the set of meaningful sentences in a given context.

It has been shown by Demolombe in [8] that for any sentence p and q of CPC we have T(p)=T(q) and F(p)=F(q) for every model M iff we have $\vdash p \leftrightarrow q$, in classical propositional calculus, and p and q are formed with the same atoms. However, if we only have the property T(p)=T(q) for every model, then p and q are not necessarily formed with the same atoms. Consider, for instance, $p = a \wedge \neg a$ and $q = b \wedge \neg b$, where a and b are atoms.

If we give up the constraint $T(p) \cap F(p) = \emptyset$ we get a 4-valued logic of the kind suggested by Lewis in [14]. Unfortunately it has been shown in [8] that, even in this 4-valued logic the fact that $\vdash p \leftrightarrow q$ (where p and q are formed with the same atoms) is not equivalent to the fact that T(p)=T(q) for every model . See, for instance, $p = (a \wedge \neg a) \wedge b$ and $q = (a \wedge \neg a) \wedge \neg b$. In the 4-valued logic T(p) and T(q) are different whenever there exist a world w_1 where a is true and false and b is only true, and a world w_2 where a is true and false and b is only false.

These results indicate that in order to characterize the set of topics of a given sentence p it is not enough to consider T(p); we have to consider the pair $< T(p), F(p) >$.

The truth conditions for sentences in L are:

$M, w \models p$ iff $w \in T(p)$, if p is an atom of CPC.
$M, w \models \neg p$ iff $M, w \not\models p$.
$M, w \models p \vee q$ iff $M, w \models p$ or $M, w \models q$.
$M, w \models A(t,'p')$ iff $I(t) \in N(T(J(\text{“p”})), F(J(\text{“p”})))$.

In the last condition “p” is the name of the sentence p, and J(“p”) denotes the sentence which is the interpretation of the name “p”, that is p. In the following J(“p”) will be abbreviated to p.

The last condition intuitively says that A(t,'p') is true iff the topic t is one of the topics assigned by the function N to the proposition expressed by the sentence p.

2.3 Some interesting schemas

We understand validity of a sentence schema in the usual way, as truth in all the worlds in all the models. Among the schemas which the semantical conditions given in the previous section do **not** validate is:

$$\text{(i)}\quad A(t,'p \wedge q') \rightarrow (A(t,'p') \vee A(t,'q'))$$

This (negative) result is important in relation to our initial set of assumptions (Section 1), where we claimed that the mode of composition of a compound sentence may play a significant role in determining the class of topics the compound is about. Thus 'p $\wedge$ q' might be about the topic of conflict, whilst neither 'p' nor 'q' is about conflict.

By contrast a good case can be made for maintaining that where two or more sentences are each about the same topic, then their conjunction is also about this topic, i.e.,

$$\text{(ii)}\quad (A(t,'p') \wedge A(t,'q')) \rightarrow A(t,'p \wedge q')$$

The validity of (ii) may be secured by imposing the following constraint on function N:

(cii) If $I(t) \in N(T(p), F(p)))$ and
$I(t) \in N(T(q), F(q)))$ then
$I(t) \in N(T(p \wedge q), F(p \wedge q)))$.

Furthermore, we have found no good reason to deny that if a sentence is about some topic t, then its negation is also about that topic:

$$\text{(iii)}\quad A(t,'p') \rightarrow A(t,'\neg p')$$

The corresponding constraint on N is:

(ciii) If $I(t) \in N(T(p), F(p)))$ then
$I(t) \in N(T(\neg p), F(\neg p)))$.

As was explained above in 2.2, when two sentences containing just the same atoms are classically equivalent, then the one is about a given topic if and only if the other is about that topic. That is:

(iv) If $\models p \leftrightarrow q$ and p and q contain just the same atoms then $\models A(t,'p') \leftrightarrow A(t,'q')$

From (iv) and (iii) it follows that:

$$\text{(v)}\quad A(t,'p') \leftrightarrow A(t,'\neg p')$$

is a valid schema.

Given now that our basic intuition, as outlined in Section 1, dictates that schema (i) not valid, it follows that the following schema is also not valid:

$$\text{(vi)}\quad A(t,'p \vee q') \rightarrow (A(t,'p') \vee A(t,'q'))$$

For suppose $A(t,'p \vee q')$; then by (iv) $A(t,'\neg(\neg p \vee \neg q)')$ and by (v) and (iv) $A(t,'\neg p \vee \neg q')$. Acceptance of (vi) would then yield $(A(t,'\neg p') \vee A(t,'\neg q'))$, which by (v) gives $(A(t,'p') \vee A(t,'q'))$, which would mean that (i) is valid.

The existing literature (see concluding overview) hardly reflects a consensus regarding the logical properties of "aboutness"; so, no doubt, some will disagree with our claims regarding which schemas are valid, and which are not. Our **firmest** intuitions lead us to reject (i) and to accept (ii) and (iv). This means that (iii) and (iv) cannot **both** be accepted.

These are just preliminary results. Much remains to be done by way of further investigation of possible constraints on the function N.

2.4 Axiomatics

The axiomatics of the logic is the result of a rather straightforward translation of the schemas which are retained for the semantics.

In the axiomatics, in addition to the axiom schemas of classical propositional calculus and to the inference rule Modus Ponens, we have the following weakened rule of equivalence:

$$\text{(REA)}\quad \frac{\vdash p \leftrightarrow q \;\text{ and }\; p \text{ and } q \text{ contain the same atoms}}{A(t,'p') \leftrightarrow A(t,'q')}$$

If we accept schemas (ii) and (iii) we have the corresponding axiom schemas:

(aii) $(A(t,'p') \wedge A(t,'q')) \rightarrow A(t,'p \wedge q')$
(aiii) $A(t,'p') \rightarrow A(t,'\neg p')$

A proof of completeness of the logic similar to the proof of completeness given in [4, 3], and based on results in [8], should work without any particular difficulties. The proof of validity is quite easy.

3 Application domains for topics

We present below three application domains where we have to characterize sets of sentences in terms of their meaning, and where topics seem to be appropriate for this purpose.

When an agent puts a query to a database system, the system can help the agent by providing him with additional information relevant to the query. This kind of system behavior is usually called "cooperative answering". In [5, 6] Cuppens and Demolombe have developed a method for cooperative answering based on the use of topics. Roughly speaking, if the query is the sentence p, and if 'p' is about the topic t, then t is identified as a topic of interest for the agent, and the system returns, in addition to the answer to p, other sentences that are about the agent's topics of interest, and are consequences of the database. Here, topics permit the characterization of the set of sentences an agent is interested in. It is worth noting that, if the structure of the set of sentences about the same topic is such that A(t,'p') implies $A(t,'p \vee q')$, the result is that if the agent is interested in p he is also supposed to be interested in $p \vee q$.

This consequence is perfectly correct for this type of application because, if it happens that the system cannot answer the query p, but it can answer the query $p \vee q$, it will be useful for the agent to know at least whether $p \vee q$ is true.

Another application domain for topics is the characterization of the reliability of agents who insert data in a database. Informally, Demolombe and Jones in [9] define an agent to be reliable for the sentence p iff we are guaranteed that, if he inserts "p" in the database, then p is true of the world. It might be much more convenient to define reliability of agents in terms of topics rather than in terms of sentences. For instance, in a company, an agent may be known to be reliable for all the sentences about the topic "accounting", another one for all the sentences about the topic "health". Here, if we accept that A(t,'p') implies $A(t,'p \vee q')$, it follows that if an agent is reliable for t, he is reliable for p and also for $p \vee q$. This consequence in general is not acceptable because the reason why the agent believes p, and the reason why he believes $p \vee q$, may be completely independent, and the former can be perfectly justified, while the latter is not. Then, in this application domain, the property: $A(t,'p') \rightarrow A(t,'p \vee q')$ should be rejected.

Finally, a third possible application domain of topics is the characterization of the set of sentences an agent is permitted or prohibited to access in a database. For example, an agent may be permitted or prohibited to access all the data which are about the topic "nuclear energy". In [7] Cuppens and Demolombe have shown that if an agent is permitted to know p, he is also permitted to know all the consequences of p, and, in particular $p \vee q$, but he is not necessarily permitted to know the antecedents of p. On the other hand, if he is prohibited to know p, he is also prohibited to know all the antecedents of p, but not necessarily its consequences. Here, we see that the property $A(t,'p') \rightarrow A(t,'p \vee q')$ is acceptable for the definition of permission , but not for prohibition, and the property $A(t,'p \vee q') \rightarrow A(t,'p')$ is acceptable for prohibition, but not for permission.

These three examples of application domains indicate that we must not confuse the rules for reasoning about "aboutness", and the rules for reasoning about concepts such as interest, for cooperative answering, or reliability, or permission or prohibition to know. We have to select, for reasoning about "aboutness", only the rules that are clearly justified for this concept, and the rules for reasoning about the other concepts have to be selected independently.

It turns out that for aboutness these rules are very weak. Nevertheless they are extremely useful. For instance, from the rule (REA), and from the fact that the sentence $'p \wedge \neg(q \vee r)'$ is about t, we can infer that all the sentences: $'p \wedge \neg(r \vee q)'$, $'\neg(q \vee r) \wedge p'$, $'\neg(r \vee q) \wedge p'$, $'p \wedge \neg q \wedge \neg r'$, $'p \wedge \neg r \wedge \neg q'$, $'\neg q \wedge \neg r \wedge p'$, $'\neg r \wedge \neg q \wedge p'$, and $'\neg r \wedge p \wedge \neg q'$, are all about t. Without the rule (REA) we should have to say explicitly for all these sentences that they are about t. That would be highly inefficient.

4 Concluding overview

In [10] Epstein defines several logics (called Relatedness logics or Dependence logics) in an attempt to formalize the notion of "topic" or "subject matter" or "referential content". For each logic two possible formalizations are proposed.

The first one is based on a "relatedness" relation on sentences. Two sentences A and B are related, and this is denoted by R(A,B), if their subject matters have something in common. For instance the subject matter of "Ralph is a dog" is related to the subject matter of "Dogs are faithful", or of "Georges is a duck", but it is not related to the subject matter of the sentence "2+2=4". It is assumed that one can say, for each pair of sentences, whether R(A,B) holds or not. A very important feature of the relatedness relation is that it is independent of the logical connectives that appear in the sentences and depends only on the relatedness of their atoms.

The second one is based on a function s(p) which assigns to each sentence p a set of topics. It is assumed that a set of topics is assigned to each atom, and that the set of topics of a compound sentence is just the union of the sets of topics assigned to each component. Here again we see that logical connectives play no role in the determination of the topics a sentence is about.

It is possible, from a given relatedness relation, to define the set assignment function s, and vice versa.

Buvač and Mason in [2] define a vocabulary (a set of atoms) associated to a context. A sentence is meaningful in a given context if its vocabulary is included in the vocabulary of this context. We notice that this notion of meaningfulness is independent of the logical connectives that appear in a sentence. In the semantics the set of meaningful sentences in a given context is characterised with the Bochvar three-valued logic. In their approach the vocabulary is only part of the definition of a context. A context is also defined by a set of sentences that hold in this context.

In [12, 13] Lakemeyer gives a formal definition to sentences of the form "sentence p is all the agent i knows about the subject matter s". A subject matter is defined as a set of atoms, and a sentence is about a given subject matter if its set of atoms is included in the set of atoms of the subject matter. That means that the fact that a sentence is, or is not, about a subject matter is independent of the logical connectives that appear in the sentence.

There are two important differences between our work and the works referred to above. The first one is that we take into account the logical connections between the components of a sentence, because we think they play an important role in the definition of their meaning. The second one is that the concept of aboutness is explicitly represented in the language via sentences of the form A(t,'p').

Maybe Goodman was one of the first authors to point out (in [11]) that "logically equivalent statements are not always about the same things", but from his statement: "is there no reasonable sense in which a statement may be about a given a class and yet not about its complementary class?", it is clear that he did not accept the schema (iii) .

Finally, to close this comparison with other related works, we have to mention the work of Lewis. In [14] Lewis defines a "subject matter" as "part of the world in intension". An equivalence relation on worlds is defined from a given subject matter as follows: two worlds are in the same equivalence class if and only if they are exactly alike for that part of the world defined by the subject matter. Moreover, a proposition corresponds to each equivalence class. This equivalence relation suggests an extensional view of subject matters.

Then Lewis informally analyses the structure of subject matters. This analysis is based on the inclusion relation on subject matters. From the intensional point of view a subject matter M1 is includeed in a subject matter M2, if the part of the world described by M1 is included in the part of the world described by M2. For instance the subject matter "1680" is included in the subject matter "17th century", since the description of the 17th century contains the description of every year in the 17th century, and in particular of the year 1680.

Notice that, as in our approach, Lewis accepts schema (iii) (" a proposition and its negation should be exactly alike with respect to what they are about"), and he does not accept schema (i). In order to avoid assigning the same subject matter to all the noncontingent propositions he suggests considering a four valued logic, but no details are given about the formal definition of this logic.

Acknowledgements: We are very grateful to Sylvie Cazalens whose previous joint work with us played a role in the development of this paper.

References

[1] D.A. Bochvar. Two papers on partial predicate calculus. Technical report STAN-CS-280-72, Stanford University, 1972. Translation of two papers from 1938 and 1943.

[2] S. Buvač, V. Buvač, and I.A. Mason. Metamathematics of contexts. *Fundamenta Informaticae*, 23(3), 1995.

[3] S. Cazalens. *Formalisation en Logique non standard de certaines méthodes de raisonnement pour fournir des réponses coopératives, dans des systèmes de Bases*

de Données et de Connaissances. PhD thesis, Université Paul Sabatier, Toulouse France, 1992.

[4] S. Cazalens, R. Demolombe, and A. Jones. A Logic for Reasoning about Is About. Technical report, ONERA-CERT, 1992.

[5] F. Cuppens and R. Demolombe. Cooperative Answering: a methodology to provide intelligent access to Databases. In *Proc. of 2d Int. Conf. on Expert Database Systems*, Tysons Corner, Virginia, 1988.

[6] F. Cuppens and R. Demolombe. How to recognize interesting topics to provide cooperative answering. *Information Systems*, 14(2), 1989.

[7] F. Cuppens and R. Demolombe. Normative Conflicts in a Confidentiality Policy. In *Proc. of ECAI'94 Workshop on Artificial Normative Reasoning*, 1994.

[8] R. Demolombe. Multivalued logics and topics. Technical report, CERT-ONERA, 1994.

[9] R. Demolombe and A. Jones. Deriving answers to safety queries. In R. Demolombe and T. Imielinski, editors, *Nonstandard queries and nonstandard answers*. Oxford University Press, 1994.

[10] R.L. Epstein. *The Semantic Foundations of Logic, Volume1: Propositional Logic*. Kluwer Academic, 1990.

[11] N. Goodman. About. *Mind*, LXX(277), 1961.

[12] G. Lakemeyer. All you ever wanted to know about Tweety. In *Proc. of the Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992.

[13] G. Lakemeyer. All they know about. In *Proc. of the 11th National Conference on Artificial Intelligence (NCAI-93)*, 1993.

[14] D. K. Lewis. Relevant implication. *Theoria*, LIV(3), 1988.

[15] J. McCarthy and S. Buvač. Formalizing contexts. Technical Report STAN-CS-TN-94-13, Stanford University, 1994.