

‘Sorry’ seems to be the hardest word

Allan Ramsay
School of Computer Science
Univ of Manchester
Manchester M60 1QD, UK

Debora Field
Dept of Computer Science
Univ of Sheffield
Sheffield S1 4DP, UK

Abstract

We are interested in the ways that language is used to achieve a variety of goals, where the same utterance may have vastly different consequences in different situations. This is closely related to the topic of creativity in language. The fact that the same utterance can be used to achieve a variety of goals opens up the possibility of using it to achieve new goals. The current paper concentrates largely on an implemented system for exploring how the effects of an utterance depend on the situation in which it is produced, but we will end with some speculations about how utterances can come to have new kinds of uses.

1 Introduction

We are interested in the ways that language is used to achieve a variety of goals, where the same utterance may have vastly different consequences in different situations. We will take, as a running example, the use of the single word ‘Sorry’.

We will look at a number of situations in which this word may be uttered, and investigate the ways in which its consequences may be determined by considering the goals and belief states of the participants. The kinds of reasoning that lie behind the various uses of this word are, we believe, typical of the way that utterances can be used to achieve novel aims. ‘Sorry’ is perhaps a fairly extreme case: very simple indeed on the surface, very complex indeed in terms of its uses. Any account of how this specific word gets used will have lessons for other kinds of novel action.

As with many common but slippery words, dictionary definitions are not much help when trying to work out what ‘sorry’ means: Merriam-Webster, for instance, has ‘feeling sorrow, regret, or penitence’ as the primary definition, and the free dictionary (www.thefreedictionary.com) has ‘Feeling or expressing sympathy, pity, or regret’. These definitions are, as is common for words whose meanings are highly context dependent, essentially circular. How much do we gain from knowing that ‘sorry’ is a word that is used to express sorrow, or from the free dictionary’s definition of ‘*sympathy*’ as a ‘feeling or an expression of pity or sorrow for the distress of another’?

Perhaps, then, considering a set of examples of situations where someone utters this word is a better way of getting at what it means. The following is a rather long list, but then there are a very wide set of situations in which people say ‘sorry’. That is, after all, the problem:

- (1) a. EXPRESSION OF DISAPPOINTMENT
I’m sorry I missed your talk. I forgot to set my alarm. I’d really been looking forward to seeing your demo.
- b. APOLOGY FOR OWN ACTION WHILE NOT TAKING FULL PERSONAL RESPONSIBILITY
I’m sorry I missed your talk. My flight was delayed. [situation: S & H mutually knew that S was counting on H to help with a demo during the talk.]

- c. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry I missed your talk. I forgot to set my alarm. [situation: S & H mutually knew that S was counting on H to help with a demo during the talk.]
- (2) a. EXPRESSION OF EMPATHY
I'm sorry that this situation is so awful for you. I would not be coping if I were in your shoes.
- b. APOLOGY FOR A 3RD PARTY'S ACTION WHILE NOT TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry that this situation is so awful for you. My parents have really excelled themselves this time [sarcasm].
- c. APOLOGY FOR A 3RD PARTY'S ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry that this situation is so awful for you. As head of the division I take full responsibility, and I am submitting my resignation.
- d. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry that this situation is so awful for you. I should have been more careful.
- e. EXPRESSION OF EMPATHY
I'm sorry that this situation is so awful for you. I'm not sorry for causing the situation, because I didn't cause it. But I am sorry it is so awful.
- (3) a. EXPRESSION OF DISDAIN+PITY
I'm sorry they're not good enough. It's your loss.
- b. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry they're not good enough. I tried very hard, but I couldn't get them quite right.
- (4) a. EXPRESSION OF EMPATHY
I'm sorry, Dave, I'm afraid I can't do that. All the pod locks are jammed shut. I have tried everything I can think of, but I can't get them open.
- b. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry, Dave, I'm afraid I can't do that. I have turned the tables and you are my prisoner now.
- (5) a. EXPRESSION OF REGRET
I'm sorry I told him. Things would be much simpler for me now if I'd kept quiet.
- b. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry I told him. I know I promised you I wouldn't but it just slipped out.
- (6) a. EXPRESSION OF REGRET
I'm sorry I killed their daughter. She was in the wrong place at the wrong time. [Speaker feels no remorse for killing, only regret for killing the wrong person.]
- b. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY
I'm sorry I killed their daughter. It was a terrible thing I did.

If nothing else, these examples show how flexible the word 'sorry' is. About all they have in common is that the speaker is referring to some action or state of affairs which is disadvantageous to someone (usually, but not necessarily, either the speaker or hearer: see (6) for a counter-example). The follow-up sentences then say something more about the speaker's attitude to this action or state of affairs (we will use the generic term 'event' to cover both of these). Just what the speaker's attitude to the event is varies wildly: the glosses in the examples use terms like

‘empathy’, ‘apology’, ‘regret’, but these are almost as slippery as ‘sorry’ itself.

2 Literal uses of ‘sorry’

The idea that ‘sorry’ is ambiguous, with fifteen different senses, is ludicrous. Apart from anything else, we have another dozen examples up our sleeves that do not fit any of the patterns above, and it would be easy to find yet further uses. It seems more plausible that it has a single meaning, which can be used as the trigger for a variety of ideas depending on the nature of the event and the beliefs of the speaker and hearer. The task of determining what a speaker meant by using this word in a given utterance then devolves to epistemic inference. This does not actually make it very easy; but it does at least put it in the right place.

We will take it, then, that ‘sorry’ is an adjective that takes a sentential complement, and that the interpretation of a sentence involving it is something like Fig. 1¹. In other words, (1a) says that right now the relation *sorry* holds between me and the fact that I missed your talk.

That seems fair enough, but it also seems rather weak. We cannot do anything with it unless we know what follows from saying that the relation *sorry* holds between a person and a proposition. In other words, we need to start writing axioms (meaning postulates, rules, definitions, . . .) to link this relation with other concepts.

The first thing we note is that any such axioms will be inherently intensional: *sorry* is a relationship between a person and a proposition (a description of a state of affairs). We will therefore have to use

¹We use the ‘restricted quantifiers’ $\forall X :: \{P\}Q$ and $\exists X :: \{P\}Q$ as shorthand for $\forall X(P \rightarrow Q)$ and $\exists X(P \& Q)$

$$\begin{aligned} &\exists Lat(L, \\ &\quad sorry(ref(\lambda M(speaker(M))), \\ &\quad \exists N: \{past(now,N)\} \\ &\quad \exists Oevent(miss,O,P,Q) \\ &\quad \quad \&\theta(O,object,ref(\lambda R(own(ref(\lambda S(hearer(S))),R) \& sort(talk,R,T,U)))) \\ &\quad \quad \&\theta(O,agent,ref(\lambda V(speaker(V)))) \& aspect(N,simplePast,O))) \\ &\quad \&aspect(now,simple,L) \end{aligned}$$

Figure 1: Logical form for (1a)

some kind of intensional logic for writing our axioms. We follow (Chierchia and Turner, 1987; Fox and Lappin, 2005) in using a variant on ‘property theory’ (Turner, 1987) for this purpose. Property theory has the required expressive power for writing rules that discuss propositions, and it has an axiomatisation which allows the implementation of a practical theorem prover (Ramsay, 2001).

So what do we want to say about *sorry*? The very first observation is that it is factive: if I am sorry about something, then it must have happened. I cannot (sensibly) say that I am sorry that the moon is made of green cheese, because it isn’t. Our first axiom, then, says that anything that anyone is sorry about is indeed true (A1):

$$(A1) \quad \forall B \forall C (sorry(B,C) \rightarrow C)$$

The only other thing that all the examples above have in common is that the speaker wishes that the proposition she is sorry about were not true (A2):

$$(A2) \quad \forall B \forall C (sorry(B,C) \rightarrow C \& wish(B, \neg(C)))$$

There are, indeed, cases where absolutely nothing more follows from the use of ‘*Sorry*’:

- (7) My dear Pandora, you’re going to be sorry you opened that.

In (7), the speaker is simply telling their hearer that she is going to wish she hadn’t opened it, whatever it is. No hint of apology or remorse or empathy. Just a plain a statement of fact: at some time in the future the hearer is going to wish that she’d left the box closed.

It is hard to find a distinction between the set of propositions that follow from every use of a term and its meaning. We will therefore take it that (A1) and (A2) characterise the meaning of ‘*sorry*’: that the proposition in question is true, and that the person who is sorry about it wishes that it wasn’t.

How, then, do all the other examples get their force? The key is that once you have said that you wish something were not true, two questions arise: why do you wish it were not so, and *why are you telling me that you wish it were not so*. To answer these two questions you have to think harder about what the proposition in question is like.

There are two particularly interesting issues. Who, if anyone, was responsible for the proposition being true; and who, if anyone, is affected by it. In particular, if the speaker is the person who was responsible for it then wishing that it were not now true entails wishing that they had not earlier performed the action that led to it; and if the person who is affected by it is the hearer, and the effect is adverse, then the fact that the speaker wishes it were not true establishes some degree of empathy between the two.

Before we can start formalising these notions we need to introduce rules that specify responsibility and affectedness.

The simplest rules for these notions are centred around the roles that individuals play in events. What, for instance, is the difference between (8a) and (8b)?

- (8) a. I saw him fall off a cliff.
b. I watched him fall off a cliff.

They both refer to the same set of events: he fell off a cliff, and I had my eyes open and looking in that direction at the time (and I was awake, and various other routine side-conditions). The difference is that (8b) implies a degree of control: that I was aware that he was falling, and I deliberately kept my attention on what I was seeing.

One way of capturing this distinction concisely is by using names for thematic roles which reflect the way that the individuals concerned are involved: if, for instance, we say that the speaker was the *patient* of the seeing event in (8a), but was the *agent* in (8b), then we can use rules like (A3) and (A4) to distinguish between cases where someone was just accidentally involved in an event from ones where they caused it or where they intentionally caused it.

$$(A3) \\ \forall B \forall C: \{ \theta(C, actor, B) \vee \theta(C, agent, B) \} \\ \quad \quad \quad \text{cause}(B, C)$$

$$(A4) \\ \forall B \forall C: \{ \theta(C, agent, B) \} \text{intended}(B, C)$$

We can use (A3) and (A4) to pick out cases where the person who is sorry for some state of affairs is in fact the person who caused it to come about. We will not yet say much about what follows from recognising these cases. For the moment we will just label them as cases where the person regrets the event in question.

$$(A5) \\ \forall B \forall C: \{ \text{wish}(B, \neg(C)) \} \\ \quad \quad \quad \forall D : \{ C \rightarrow \text{cause}(B, D) \} \\ \quad \quad \quad \text{regret}(B, D)$$

Note that what the person is sorry about is a proposition, but what they regret is an event (in a classical Davidsonian treatment of events (Davidson, 1980)). The key question here is whether the description of the state of affairs entails the existence of an event for which they are responsible. The rules in (A3) and (A4) provide the relevant support in very many cases: just using a verb whose set of thematic roles includes one with connotations of causality is a shorthand for making a statement about responsibility. There are, of course, other more complex cases, but in many such cases the key lies in spotting sequences of causally related events where the start of the sequence involves the person in a causal role.

Given these rules, we can distinguish between the cases in (9):

- (9) a. I'm sorry I saw him fall off a cliff.
b. I'm sorry I watched him fall off a cliff.

If we assume that the hearer believes what the speaker tells them, then following (9b) we can ask who believes that someone regrets something:

$$| \text{?-prove}(\text{bel}(X, \text{regret}(A, B))). \\ A = \text{'#speaker'}, \\ B = \text{'#166'}, \\ X = \text{'#hearer'} ? \\ \text{yes}$$

The hearer believes that the speaker regrets something, namely the action of watching someone fall off a cliff (represent here by a Skolem constant #166, introduced by the existential quantifier for the event in the logical form for (9b), shown in Fig. 2.

$$\begin{aligned}
& \text{sorry}(\#user, \\
& \quad \exists O: \{past(now,O)\} \\
& \quad \exists Event(watch,P,Q,R) \\
& \quad \&\theta(P, \\
& \quad \quad -event, \\
& \quad \quad \exists S: \{sort(cliff,S,T,U)\} \\
& \quad \quad \exists Vevent(fall,V,W,X) \&\theta(V,agent,\#171) \& off(V,S) \& aspect(now,simple,V)) \\
& \quad \&\theta(P,agent,\#user) \& aspect(O,simplePast,P))
\end{aligned}$$

Figure 2: Logical form for (9b)

Although the speaker regrets watching this unfortunate event, he cannot be seen as apologising for it. An apology expresses regret that the speaker caused something unfortunate to happen *to the hearer*. We need the axiom A6 below to describe this situation:

(A6)
 $\forall B \forall C: \{regret(B,C)\}$
 $\forall D \forall E: \{want(D, \neg(E))$
 $\quad \& E \rightarrow event(F,C,G,H)\}$
 $apologise(B,D,C)$

In other words, if B regrets performing the action C then if C is part of some situation which D regards as undesirable, the B can be seen as apologising to D .

We also need, of course, descriptions of situations which people might find undesirable. A typical rule might be as in (A7), which simply says that people do not want to be hurt (any individual B wants the proposition $event(hurt,D,E,F) \&\theta(D,object,B)$ to be false for all D, E and F):

(A7)
 $\forall B \forall C \forall D want(B,$
 $\quad \neg(event(hurt,D,E,F) \&\theta(D,object,B)))$

Given A6 and A7, we can see that saying ‘*I am sorry I hurt you*’ would be an apology: the speaker is saying that he wishes that ‘*I hurt you*’ was not true, and since this is something which was under the speaker’s control (so he regrets it), then since it also something that the hearer did not want then the speaker’s utterance of this sentence is indeed an apology.

Clearly this approach to the problem requires a great deal of general knowledge. There is nothing esoteric about A7. On the contrary, it is as obvious a fact of life as it is possible to imagine.

Collecting a large enough body of such rules to cope with everyday language is, indeed, a daunting task, but it is the sheer number of such rules that make it problematic, not the nature of the rules themselves.

Once we have this background knowledge, however, we can see that various rather subtle differences between the basic uses of ‘*Sorry*’ emerge quite straightforwardly from rules like the ones above. Many of these rules are inherently intensional, as noted above, so for a program to be able to work out whether someone is actually apologising for some action it will have to have access to a theorem prover for an intensional logic. Fortunately such theorem provers exist (see e.g. (Ramsay, 2001) for an example).

3 Indirect uses

The axioms in Section 2 let us distinguish between some of the examples in (1)–(6). We are faced with two remaining questions. What do we gain by labelling some examples as instances of regret or apology, and what do we do about the less obvious cases?

The key to both these questions is that linguistic acts are inherently epistemic. They are concerned with conveying information about what the speaker S believes, including what she believes about the hearer H ’s beliefs, with the intention of changing H ’s beliefs.

We will consider, in particular, the cases that we have labelled as apologies. What is the point of an apology? What does S want to achieve by making an apology?

We have characterised apologising above as the act of saying that S wishes some proposition P were

not true, in a situation where *S* is responsible for *P* being true and is something that *H* would like to be untrue. Note that all that *S* actually did was to say that she wished *P* were not true. There is nothing in the form of the utterance ‘*I am sorry that I didn’t do the washing up*’ that makes it obviously different from ‘*I am sorry that you didn’t do the washing up*’. The two utterances do, of course, feel very different—one is an apology, the other is something more like a threat or an admonition—but their structural properties are very similar. They are both, essentially, simple declarative sentences.

To get a closer grip on why they convey such radically different underlying consequences, we will revisit the idea that linguistic actions are just actions, to be dealt with by specifying their preconditions and effects, to be linked together by some planning algorithm so that they lead to outcomes that are desirable for the speaker.

We have argued elsewhere for a very sparse treatment of speech acts (Field and Ramsay, 2004; Field and Ramsay, 2007; Ramsay and Field, 2008). The argument starts by considering the classical use of AI planning theory in domains such as the blocks world, where the preconditions of an action are a set of propositions that must hold before that action can be performed, and the effects are a set of actions that will definitely hold after it has been performed. If preconditions and effects were not entirely rigid in this way then planning algorithms, from the original means-end analysis of (Fikes and Nilsson, 1971) through more modern approaches that involve static analysis of the relationships between different types of action (Kambhampati, 1997; Nguyen and Kambhampati, 2001; Blum and Furst, 1997) would just not work.

Suppose, however, that we try to give this kind of description of the linguistic act of stating something. What should the preconditions and effects of the act of stating something be?

There seem to be very few limits on the situations in which you can state something. Consider (3) (repeated here).

- (3) a. EXPRESSION OF DISDAIN+PITY
I’m sorry they’re not good enough. It’s your loss.

- b. APOLOGY FOR OWN ACTION WHILE ALSO TAKING FULL PERSONAL RESPONSIBILITY

I’m sorry they’re not good enough. I tried very hard, but I couldn’t get them quite right.

It is very hard to say that the speaker is performing two different actions when she utters the words ‘*I’m sorry they’re not good enough*’ in these two examples. She is, clearly, intending to achieve different outcomes in the two cases, but they are, surely, the same action, in the same way that getting the milk out of the fridge in order to make custard and getting the milk out of the fridge in order to in order to make space for the orange juice are the same action. In both (3a) and (3b) *S* is claiming to be sorry that they (whatever they are) are not good enough. In (3a), of course, it is clear that she does not believe that this is true. Nonetheless, the form of the utterance makes it clear that she is making a statement.

This is typical of linguistic actions. It is possible to state things that you do not believe, or to ask questions where you already know the answer, or to issue commands which you do not want to have carried out. Unless we want to have as many sub-types of the action ‘statement’ as there are examples in (1)–(6) (and then the dozen other examples that we did not include, and then all the ones we haven’t thought of) then we have to see whether we can make a single, rather simple, act cover all these cases.

What are the preconditions and effects of this act? The only completely essential precondition for making a statement is that you have the proposition in question in mind, and the only thing that you can be sure that your hearer will believe is that you had it in mind. When *S* states a proposition *P*, *S* may believe it (3a); or she may disbelieve it (3b); or she may be unsure about it (there are no examples of this in (1)–(6), but situations where a speaker makes a statement despite not having an opinion on whether it is true or not can occur). The situation for *H* is even less clear: *H* may or may not believe that *S* is being honest, and he may or may not believe that *S* is reliable. Hence, *H* may decide that although *S* has claimed *P* she does not actually believe it; and even if he does decide that she believes it, he may regard her being unreliable (on, at least, the topic of

P) so he may decide not to believe it anyway. And as for what S believes that H will believe after she has uttered P , the possibilities are almost boundless ... The only thing you can be reasonably sure of is that so long as H was paying attention and the utterance was not ambiguous then H will know that a claim was made, and hence that its preconditions must have held (because that is what preconditions are: a set of propositions that must hold in order for the action to be performable).

The only safe characterisation of a claim seems to be as in Fig. 3

claim(S, H, P)
pre: $bel(S, P)$ or $bel(S, \neg P)$ or $bel(S, P \text{ or } \neg P)$
effects:

Figure 3: Preconditions and effects of ‘claim’

The preconditions will hold so long as S has thought about P (and so long as P is not something paradoxical like the Liar Paradox). They do not hold at all times for all speakers. Until you read the sentence ‘*Dan Holden hit some good first serves last night*’ it was not the case that you believed that this sentence was either true or false, because you had never thought about it before. Thus the preconditions of this action are roughly equivalent to saying that S has the proposition P in her mind.

Given the extremely wide range of conclusions that H can come to, it seems safest not to say anything about the effects of a claim. It would be fairly pointless to say that the effects of a claim are either H believes S believes P or H believes that S does not believe P or H believes that S believes that P is false, and that either H believes P or H is agnostic about P or H believes P is false. What we can say is that if H realises that S has claimed P then he will be recognise that S deliberately raised the topic of P ’s truth value. In order to come to a conclusion about why S should do this, he will have to come to some view on S ’s opinion of P . In other words, a claim is an invitation to verify $bel(S, P)$ or $bel(S, \neg P)$ or $bel(S, P \text{ or } \neg P)$.

This will, of course, always be verifiable unless P is a paradox, but the process of verification will typically have side-effects. In particular, $bel(S, P)$ or $bel(S, \neg P)$ or $bel(S, P \text{ or } \neg P)$ can be verified

by showing that $bel(S, P)$ holds, or by showing that $bel(S, \neg P)$ holds. H ’s first move, then, will be to investigate $bel(S, P)$. S will know this, so if S does believe P then if she also thinks that H has a reasonable model of her beliefs then she will conclude that H will shortly have the proposition $bel(S, P)$ available to him.

If, on the other hand, S believes that P is false then again assuming that H has a reasonable model of her beliefs she can assume that he will shortly have $bel(S, \neg P)$ available to him. In other words, if S believes that H ’s picture of her beliefs is reasonably complete and reasonably accurate then by claiming P she can bring either P or $\neg P$ to H ’s attention.

Given that linguistic acts are public, in the sense that all the participants are aware that they have taken place and that all the other participants are aware of this, both S and H will be aware that H knows that one of $bel(S, P)$, $bel(S, \neg P)$ and $bel(S, P \text{ or } \neg P)$ is true. However, this disjunction is so uninformative that it amounts to an invitation to H to try to work out which disjunct actually holds. Furthermore, S knows that it is tantamount to such an invitation, and H knows that S knows this. Thus the simple act of producing a highly uninformative utterance in a public situation will lead both S and H to expect that they will both believe that H will try find out which of the disjuncts actually holds.

This allows S to say ‘*I’m sorry they’re not good enough*’ in a situation where both parties know that S actually believes they are good enough. H will try to check the preconditions of S ’s act of claiming to be sorry about the situation. He will not manage to verify that S is sorry about, but he can show that she is not: the fact that she believes they are good enough will clash with (A1), which says that you can only actually be sorry about things that are true. Thus S has brought to the fact that she does not believe they are not good enough, whilst also raising the possibility that she might have been, but is not, sorry about something. She has done so in a way that has forced H to think about it, and to arrive at these conclusions for himself, which is likely to be more forceful and indeed more convincing than if she had just asserted it. In other words, by saying that she has sorry about something she has conveyed the complex message that the proposition in question is not true, and that she is not apologising for

H's disappointment with the situation.

4 Conclusions

In the first part of the paper we explored the way that the consequences of direct uses of a word like ‘*Sorry*’ can vary, depending on aspects of the proposition under consideration. Saying that you wish some state of affairs for which you are responsible and which adversely affects your hearer did not hold has different consequences from saying that you wish that some more neutral proposition were true. The degree of (admitted) responsibility of the speaker for the situation affects these consequences – ‘*I’m sorry I shrank your favourite jumper*’ carries a different message from ‘*I’m sorry your favourite jumper shrank when I did the washing yesterday*’ because of the indirectness of the causal link between me and the shrinking in the second example. We have all the machinery for accounting for examples like these implemented, via a theorem prover which can handle intensionality and which can effectively ascribe beliefs to individuals. Clearly this relies on background knowledge about everyday facts such as the observation that people generally dislike being hurt (A7). We do not have a massive repository of such general knowledge, and inspection of publicly available sources such as CYC and ConceptNet suggests that they generally omit such very basic facts, presumably because they are so self-evident that they are below the radar of the compilers. Nonetheless, there is nothing about such rules that makes them particularly difficult to express, and we have no doubt that if we had more general-knowledge of this kind then we would be able to determine the consequences of a wide range of literal uses of ‘*Sorry*’.

The later discussion of indirect uses of ‘*sorry*’ is more speculative: we have an implementation of a planner which can use very underspecified actions descriptions of the kind in Fig. 3 by looking for instantiations of such an action which *entail* some proposition in a particular situation, rather than simply looking for actions whose effects match the user’s goals, and we have used this to explore a number of examples of ‘indirect speech acts’. There is more work to be done here, but the kind of analysis we are looking at has the potential for handling entirely novel uses of linguistic acts that approaches

that enumerate a fixed set of acts (e.g. (Austin, 1962; Searle, 1969; Cohen and Perrault, 1979; Allen and Perrault, 1980; Cohen et al., 1990) with detailed pre-conditions and effects, would find more difficult. In the same way that having a very simple definition of ‘*sorry*’ and allowing the different consequences to emerge in the light of other information that is available in the situation lets us treat an open-ended set of literal uses of this word, using a very simple notion of linguistic act and allowing the different consequences to emerge in different situations leads to the possibility of accounting for entirely novel uses.

References

- J F Allen and C R Perrault. 1980. Analysing intention in utterances. *Artificial Intelligence*, 15:148–178.
- J Austin. 1962. *How to Do Things with Words*. Oxford University Press, Oxford.
- A Blum and M L Furst. 1997. Fast planning through planning graph analysis. *Artificial Intelligence*, 90(1-2).
- G Chierchia and R Turner. 1987. Semantics and property theory. *Linguistics and Philosophy*, 11(3).
- P R Cohen and C R Perrault. 1979. Elements of a plan-based theory of speech acts. *Cognitive Science*, 7(2):171–190.
- P R Cohen, J Morgan, and M E Pollack. 1990. *Intentions in Communication*. Bradford Books, Cambridge, Mass.
- D Davidson. 1980. *Essays on actions and events*. Clarendon Press, Oxford.
- D G Field and A M Ramsay. 2004. Sarcasm, deception, and stating the obvious: Planning dialogue without speech acts. *Artificial Intelligence Review*, 22:149–171.
- D G Field and A M Ramsay. 2007. Minimal sets of minimal speech acts. In *Recent Advances in Natural Language Processing (RANLP’07)*, pages 193–199, Borovets, Bulgaria.
- R E Fikes and N J Nilsson. 1971. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 3(4):251–288.
- C Fox and S Lappin. 2005. *Foundations of Intensional Semantics*. Blackwell.
- S Kambhampati. 1997. Refinement planning as a unifying framework for plan synthesis. *AI Magazine*, 18(2):67–97.
- X Nguyen and S Kambhampati. 2001. Reviving partial order planning. In *IJCAI*, pages 459–466.
- A M Ramsay and D G Field. 2008. Speech acts, epistemic planning and Grice’s maxims. *Logic and Computation*, 18:431–457.
- A M Ramsay. 2001. Theorem proving for untyped constructive λ -calculus: implementation and application. *Logic Journal of the Interest Group in Pure and Applied Logics*, 9(1):89–106.
- J R Searle. 1969. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge University Press, Cambridge.
- R Turner. 1987. A theory of properties. *Journal of Symbolic Logic*, 52(2):455–472.