RECOGNIZING ADVICE, WARNINGS, PROMISES AND THREATS

Kevin Donaghy

School of Computer Science and Information Technology
Rochester Institute of Technology, Rochester, New York 14623
hkd@cs.rit.edu

It is argued here that utterances in the imperative mood typically are embedded either explicitly or implicitly in Aristotelean practical arguments, i.e., arguments whose conclusions specify an action to be performed by an agent and whose premises provide justification for that action. It is argued further that the illocutionary force of an imperative, e.g., advice, warning, request, etc., is determined by the structure of the practical argument in which it is embedded. Algorithms for reconstructing partial practical arguments are developed. Emerging from the discussion is a set of heuristics for identifying advice, warnings, conditional promises and conditional threats in natural language settings. Sample output from a test program employing these heuristics is presented. Finally, it is suggested that the techniques outlined in this paper point to the possibility of dialogue and story understanding systems which are more general and require significantly less domain specific knowledge than do current systems.

1. Practical Arguments

Consider the utterance "Don't go near the stove". Is this a warning, an order, a request or perhaps an instance of some other speech act category? Apart from context, it is impossible to tell. But once context is supplied the answer is typically quite evident.

1. If you touch the stove, you will burn yourself. So don't go near the stove. (warning)
2. The player who avoids touching the stove usually wins the game. So don't touch the stove. (advice)
3. I can't take another one of your casseroles. If you want to live don't touch the stove. (threat)

What these cases have in common is that they are all examples of what philosophers since Aristotle have called practical arguments, that is, arguments whose conclusions name an action to be performed by an agent and whose premises provide reasons for the agent to perform that action. Speech acts such as advice, warning, instruction, and moral exhortation are conceptually linked to practical arguments in the following way. In classifying an utterance as advice, warning, etc., we specify the type of practical argument in which that utterance is either implicitly or explicitly embedded.

Consider advice, for example. To advise X to do A is, among other things, to imply that X's interests will be served by doing A. When advice comes packaged

in the form of an imperative, that implication functions as the premise of a general practical argument whose conclusion is "X, do A." To warn X not to do A, on the other hand, is to imply that X's interests will suffer if X does A. That implication in turn functions as the premise of a general practical argument whose conclusion is "X, do not do A." To morally exhort X to do A is to imply that some moral end will be served should X do A. Once again, that implication is the premise of the general practical argument being advanced by the speaker.

A fundamental assumption of this paper is that arguments of the form "If X then Y. So (don't) do Z." comprise a small but important subset of practical arguments, for the reason that many if not all practical arguments with imperative or quasi-imperative conclusions can be recast in this form without loss of meaning or structure. This assumption is based on the Aristotelian means-end model of practical arguments as deliberations which "assume the end (viz. a desire need, interest or goal of the agent) and consider how and by what means it is to be attained." (Aristotle, 1915, 1112b15-31). Consider the following example.

The stove is hot. So don't touch it.

While readily understandable, this argument is incomplete. Fleshed out, it becomes

1. The stove is hot.
2. Hot things cause burns when touched.
3. If you touch the stove, you will burn yourself. (1,2)
4. (You wish to avoid burning yourself.)
5. So don't touch the stove. (3,4)

In the short version, the hearer's interests as well as the implications of the stove's being hot are so obvious that they are not mentioned. Note that in the long version, 1 is not even a premise of the main argument. Its role is to provide evidence for 3. If this example is typical, the form "If X then Y. So (don't) do Z" may well capture the deep structure of a large and significant class of practical arguments.

How does one go about reducing practical arguments to the form "If X then Y. So (don't) do Z"? To continue the example, suppose "Hot things cause burns when touched" has been stored in a knowledge base. The reduction of "The stove is hot. So don't touch it" can then be carried out as follows.

1. Assume that the real premise (RP) of the the argument is of the form "If X then Y" where X is the negation of the propositional content of the conclusion and Y is some as yet unspecified harm to H.
2. Also assume that the role of the stated premise (SP) is to provide evidence for RP.

This project is funded by Universal Energy Systems, Inc. and the United States Air Force, Contract No. F49620-88-C-0053/SB5881-0378.

3. The consequent of RP (viz. "you will burn yourself") can now be deduced from SP and the known fact that hot things cause burns when touched.

This and similar algorithms have been implemented in PASA (Practical Arguments & Speech Acts), a test program which accepts practical arguments as input and identifies their principal speech acts.

2. PASA

Several heuristics in PASA assist in the identification of speech acts and reduce substantially the need for domain specific or nonlinguistic knowledge that would otherwise be necessary. As a start, consider the following examples.

1. If you finish your homework, I will give you more castor oil to drink. So finish your homework.

2. If you don't finish your homework, I will give you more castor oil to drink. So finish your homework.

In neither case is there any difficulty in deducing S's views on castor oil. In the first example, it is promised as reward, and in the second is threatened as punishment. What makes these deductions possible is the relationship between the propositional contents of the imperative and the antecedent of the conditional. In the first instance they are identical; in the second, the one is the negation of the other. PASA utilizes both heuristics to identify speech acts and to deduce and record the evaluative stance of the speaker towards a given state of affairs. The former is instrumental in the identification of promises and advice, and the latter of threats and warnings. The next pair of examples illustrate another helpful heuristic which in fact is a generalization of the above.

3. If you finish your homework, I will give you more castor oil to drink. So get started.

4. If you don't finish your homework, I will give you more castor oil to drink. So get started.

Since getting started on a project increases (dramatically) the likelihood of finishing it, it is obvious that example 3 is a promise and example 4 a threat. Whenever possible, PASA makes similar inferences.

3. Sample I/O

PASA was designed as a testbed for the ideas presented in sections 1 - 2. As such, it is not a full-fledged natural language system, nor even a mature prototype for such a system. Inputs to the program are one premise practical arguments with imperative conclusions. In the initial set of examples, the principal speech act is determined by examining the structure of the argument and the syntactic form of premise's consequent. Program output is in boldface.

C-Prolog version 1.4

l ?- [pasa].

pasa consulted 33004 bytes 7.449997 sec.

l ?- talk.

l: If you finish your homework then I will let you watch television. So finish that math.

S promised to let you watch television if you finish

your homework.

l: I will ground you for a week if you fail the test tomorrow. So do not fail.

S threatened to ground you for a week if you fail the test tomorrow.

l: If you study for the exam then you will pass. So study hard.

S advised you to study hard since if you study for the exam you will pass.

In the next example, PASA first has to be taught the relationship between starting and finishing a project.

l: Start on your math right away. I will let you watch television tonight if you finish all your homework before six.

Let Y be the state of affairs described in the consequent of the premise. Which of the following most accurately describes the viewpoint of the speaker?

A. Y is in the interests of the hearer.

B. Y is not in the interests of the hearer.

C. Y does not affect interests of the hearer one way or the other.

l: A.

Let X be the state of affairs described in the antecedent of the premise and Z the action specified in the conclusion. Which of the following most accurately describes the viewpoint of the speaker?

A. The hearer should do Z in anticipation of Y.

B. By doing Z hearer would increase likelihood of X.

C. Neither of the above.

l: B.

S promised to let you watch television tonight if you finish all your homework before six. Starting on your math right away will make it more likely that you will finish all your homework before six.

PASA now knows what it needs to know in order to paraphrase similar cases. The next example is an enthymeme. PASA must generate the hidden premise before it can hazard a paraphrase.

l: The stove is hot. So do not touch the stove.

S warned you not to touch the stove since if you touch the stove you will be burned.

4. Towards Language Driven Understanding

PASA is a modest example of a language driven understanding system in which the need for domain specific knowledge is minimized. A methodological decision was made early on to appeal to nonlinguistic information only as a last resort. The motivation for this was twofold. In the first place, domain driven systems are inherently limited by the vast amounts of domain specific information required to process even the simplest texts. There appears little hope of generalizing these systems so that they are capable of exploiting structural commonalities between stories and dialogues from different domains. Secondly, reliance on domain knowledge for quick fixes to text processing problems tends to deaden sensitivity to linguistic information present in those texts.

I am convinced that linguistic cues play a far richer and more powerful role in natural language understanding than has been commonly supposed and that speech act analysis will prove a useful tool for systematically investigating those cues. I conclude with an illustration of how domain and language driven approaches to story understanding might differ. Consider the following story (Wilensky, 1978, pp. 2-3).

"One day John went through a red light and was pulled over by a cop. John had just gotten a summons for speeding the previous week, and was told that if he got another violation, his license would be taken away. Then John remembered that he had two tickets for the Giant's game on him. He told the cop that he would give them to him if he forgot the whole incident. The cop happened to be a terrific football fan. He took John's tickets and drove away.

Q1 Why did John offer the cop a couple of tickets?

A1 Because he was afraid he was going to lose his license if he got another summons."

Wilensky has this to say about the story.

"Consider what is involved in making the inference that John offered the cop his football tickets to prevent the loss of his license. First, the reader would have to infer that the cop was going to give John a traffic ticket. This inference requires the knowledge that a policeman is supposed to ticket people who break traffic rules... Now the reader must interpret John's statement to the cop as an attempt to prevent him from giving him a ticket. To interpret this sentence as an offer, a reader must know that one way to prevent someone from doing something is to persuade him not to do it...by offering him something desirable in exchange for his cooperation. The understander can (then) infer that football tickets are desirable to a football fan, since football tickets are necessary for getting into a football game."

Wilensky is setting the stage for a domain driven theory of understanding in which large stores of non-linguistic knowledge are required for story comprehension. Ironically, this very type of knowledge often impedes rather than assists comprehension, a fact known to every writer who has employed the O'Henry formula to surprise his audience. One can imagine such a writer adding a short paragraph to Wilensky's story in which it is revealed that John, desperately hoping to lose his license, had threatened the cop with football tickets, a threat which proved pathetically ineffective. A major drawback in the domain driven approach is that it limits the understander to one interpretation of a story when several may be possible. Consider the same story from a language driven perspective. The story is presented in schematic form to insure that the understander has little or no domain knowledge available to him.

1. One day A did B and was approached by C.
2. A had just been given a D for doing E the previous week, and was told that if he got another D, then F would happen.
3. Then A remembered that he had a G with him.

4. A told C that he would give him a G if C did not give him a D.

5. C accepted the G and did not give A a D.

A = John

B = ran a traffic light

C = cop

D = ticket

E = speeding

F = John loses license

G = football tickets

An understander would make little headway with the schematic version of the story until reaching statement 4, where it becomes evident that A has either promised or threatened to give C a G if C does not give him a D. At this point the schema lends itself to two quite different interpretations.

Suppose that 4 is a conditional promise. Given the purpose of such promises, it follows that A does not want C to give him a D. Thus, from A's point of view, getting another D is undesirable. Now some sense can be made of statement 2. There is a strong probability that the reason why getting another D is undesirable is because it would lead to F. So F too is most likely undesirable from A's perspective. Given this interpretation of 4, the understander now knows all he needs to know about the schematic story to answer the sample question.

Q Why did A say that he would give G to C if C did not give him a D?

A1 Because A was afraid that F would happen if C gave D to A.

Now suppose that 4 is a conditional threat. In this case, it follows that A wants C to give him a D. Given statement 2, the likelihood is that the reason why getting another D is desirable from A's perspective is that it would lead to F. So F too is likely desirable. The appropriate answer to Q in this case is

A2 Because A hoped that F would happen if C gave D to A.

It is of some interest that on this method of analysis, the inherent ambiguity of the story is preserved. This example and others like it provide suggestive evidence that understanding the general structure of such stories requires far less domain specific knowledge than Wilensky would lead us to believe. Clearly, many of the necessary inferences can be drawn from linguistic cues gleaned from the text. Domain knowledge facilitates and enriches comprehension. However, it may not be as fundamental to the task of understanding as some researchers have suggested.

REFERENCES

- Aristotle (1915). *Ethica Nichomachea*. In W.D. Ross (Trans.), *The Works of Aristotle Translated into English* (Vol. ix). London: Oxford University Press.
- Cohen, P.R., & Levesque, H.J. (1985). Speech acts and rationality. *Proc. ACL*, 49-59.
- Searle, J. (1969). *Speech Acts*. Cambridge: Cambridge University Press.
- Wilensky, R. (1978). Understanding goal-based stories. Ph.D Thesis. Y140.