ATTITUDE EMERGENCE - AN EFFECTIVE INTERPRETATION SCHEME FOR PERSUASIVE DISCOURSE

HORNG-JYH P. WU and STEVEN L. LYTINEN
Artificial Intelligence Laboratory
The University of Michigan
Ann Arbor, MI 48109. U.S.A.

ABSTRACT

Previous approaches have used a reasoning mechanism called *belief percolation* to determine the actual speech intent of the speaker (e.g., (Wilks and Bien 1979)). In this paper, a similar mechanism, called *attitude emergence*, is proposed as a mechanism for inferring a speaker's attitude toward the propositions in a persuasive discourse. It is shown that in order to adequately interpret the statements in advertisements, associations of relevant semantic information, through *bridging inferences*, are to be *percolated up* through attitude model contexts to enhance and calibrate the interpretation of statements. A system called BUYER is being implemented to recognize speech intents through attitude emergence in the domain of food advertisements taken from *Reader's Digest*. An example of BUYER's processing is also presented in the paper.

Introduction

One of the most significant characteristics of persuasive discourse is that it involves the expression of people's beliefs, desires, preferences, etc. These beliefs, desires, and preferences constitute a model of *mental attitudes* (or, an attitude model) which characterizes the mind of the speaker engaging in a persuasive discourse. An attitude model is important for figuring out what the speaker means; i.e., his *speech intent*. More specifically, often when the speaker expresses his beliefs, desires or preferences in persuasive discourse, he means to induce a reaction, in the forms of comparable mental attitudes, on the part of a hearer. For example, an expression of the speaker's belief can be intended to induce such an belief in the hearer. However, in general, inferring the speech intent through attitude model reasoning is complex. For instance, in our domain of persuasive discourse – advertisements – a major statement may be followed by minor statements, as demonstrated in the following passage.

(1.1) Nabisco is great.
(1.2) It is nutritious whole wheat,
(1.3) low in cholesterol and saturated fat,
(1.4) has plenty of fiber and vitamin.

Arranged in this fashion, the expression of the speaker's preference in (1.1) – Nabisco is great – comes to be supported by statements (1.2) through (1.4). This makes the acceptance of (1.1) as the hearer's own much more compelling, intending to induce in the hearer the same preference expressed in (1.1). Although this explanation sounds intuitively simple and correct, the question remains: how do statements such as (1.2) to (1.4), which are on the surface, "disjoint" expressions of the speaker's belief, come to have a real psychological impact on the hearer? Some sort of reasoning must have been employed to *bridge* them with (1.1) and to produce the persuasive effects.

Previously, model-based reasoning has been investigated for many tasks, such as belief ascription and metaphor understanding (Ballim *et al.* 1991), logical reasoning (Dinsmore 1987), and natural language understanding in general (Fauconnier 1985). One previous approach to attitude model reasoning concentrates exactly on issues of inferring speech intents. As discussed concisely in (Wilks and Bien 1979), the statement "Frank is coming tomorrow" can be interpreted in many ways, depending on the context. For instance, if the hearer believes that the speaker believes that Frank is hostile to the hearer, and the hearer has no personal knowledge about Frank, then this statement might be interpreted as a *threat* to the hearer.

To account for different possible interpretations of statements like these, Wilks *et al.* propose an *attitude percolation* mechanism, in which a statement is *pushed down* to the frame of the system's belief to create the attitude context – the system's belief of the speaker's belief. Since in this more specific context, the following statements are simultaneously present[1]: Frank comes tomorrow and Frank is hostile to the system. Thus, the system can infer that Frank may harm the system, which in turn, allows the system to interpret the original statement as a threat to itself.

In our domain of advertisement persuasive discourse, much fewer facts are privately known compared to what are mutually known[2]. Therefore, it is argued

[1] If the system has personal knowledge otherwise about Frank, then this may affect the reasoning process.

[2] It is also an essential criteria for "good" advertisements to adhere to fashionable viewpoints.

that the *push down* operation, which investigates into the personal knowledge about one another, will be limited. Instead, in this paper, a more relevant mechanism, called *attitude emergence*, is investigated. It is a more comprehensive treatment of attitude models, utilizing mutually known semantic and world knowledge to convey speech intent, as observed in the interpretation of our earlier example. More specifically, attitude models for statements are constructed, and assimilated through bridging inferences based on mutual knowledge. The semantic association resulting from assimilation then gets *percolated up* along attitude model contexts so that the proper interpretation of speech intent is recognized. This attitude emergence mechanism is being implemented in a system called BUYER, which understands food advertisements taken from *Reader's Digest*. A corpus of 120 advertisements has been collected. Part of the ads are used for constructing the system, while the rest of them are used to verify the generality of BUYER's knowledge base. In this paper, we present an example of BUYER processing one of these ads.

The basic framework of attitude emergence

As observed above, the recognition of speech intent in persuasive discourse is rather complex. Various sorts of mutual knowledge may be employed in bridging the statements and bringing out the speech intents. In (Wu and Lytinen 1991), we proposed a three-step procedure of attitude emergence:

Step 1: Construct the initial *attitude model* (or A-model).
Step 2: Assimilate the successive statements coherently into one A-model (if possible) and recognize the semantic association between the models.
Step 3: Percolate upwards along A-model context to effect attitude change due to the semantic association recognized in Step 2.

A-models are recursive structures of information with attitude contexts, each layer of which consists of an agent and an attitude he holds toward the deeper level information. A simple passage is analyzed in (Wu and Lytinen 1991),

(2.1) Peter loves antique cars.
(2.2) His favorite model is the 1887 Duryea.

The evolution of A-models through this three-step procedure can be summarized as follows: At the end of Step 1 (see Fig. 1), an A-model is constructed consisting of an attitude context – (Report Speaker ...), which embeds another attitude context – (Love Peter ...), which contains an object – antique_cars. Note that, these linear formula are just short hands for models. We take that in implementation, models are formula plus indexing and encapsulation, as the boxes in Fig. 1 is intended to capture: the indices on agents are called *S-boxes* and on attitudes, *A-boxes*.

Each attitude context creates its own environment for simulative reasoning (Wilks and Bartley 1990),

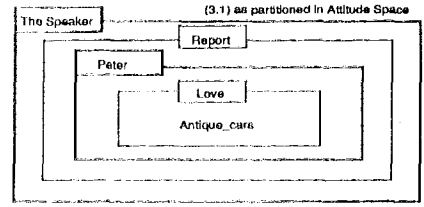

Figure 1: The initial attitude model after (2.1) has been processed.

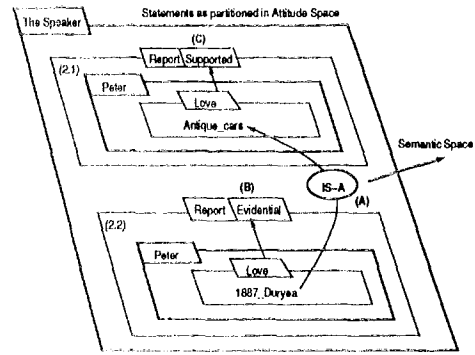

Figure 2: The attitude model and reasoning when (2.2) is assimilated.

which involves only entities with comparable attitude status. In turn, simulative reasoning can produce results applicable to the attitude context where it takes place, and may sometimes affect related contexts, e.g., causing re-evaluation of the attitude in the embedding attitude contexts. Thus, in Step 2, while statement (2.2) is being assimilated with statement (2.1), some reasoning takes place marked as (A), (B) and (C) in Fig. 2.

At point (A), an IS-A semantic relation is recognized between "antique_cars" and "1887_Duryea" in the semantic space which, as depicted as orthogonal to attitude space in Fig. 2, stores attitude-independent semantic information. In order to trigger emergence of more attitude related information, the recognition of the IS-A relation is percolated up along layers of attitude contexts to reach the (Report Speaker ...) context, where (B) and (C) occur. At points (B) and (C), statements (2.2) and (2.1) are found to be related; in particular, the (Report Speaker ...) context of (2.2) is calibrated to be "evidential" (to (2.1)), while, that of (2.1) to be "supported" (by (2.2)). That (B) and (C) occur is due to the following world knowledge (WK1):

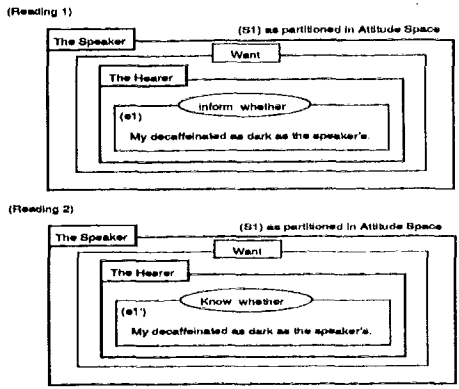

Figure 3: The two attitude models for the readings of *S1*.

> **If** the speaker provides *more detailed information* about a claim (Y) in a statement (X)
> **Then** Y is "supported" by "evidence" X and becomes *more believable*

As demonstrated by (A) – (C), semantic association emerges from embedded attitude contexts to calibrate the attitude in higher contexts - i.e., how attitude emergence happens.

In this simple example, the attitude emergence mechanism has involved, nonetheless, a large amount of knowledge. This knowledge is briefly reviewed below. Yet, more sophisticated reasoning is required to process real world ads (see below). First, there is knowledge concerning A-model construction based on the following mapping rules: (1) Sentence types to A-boxes, e.g., a declarative sentence type maps into a belief; (2) Attitude verbs to A-boxes, e.g., "loves" into a preference; (3) Evaluative predicates to A-boxes, e.g., "favorite" into a preference; (4) Adverbs to A-models, e.g., "certainly" into a belief; (5) Cue phrases to A-models, "it is time that" into a desire (recommendation).

Secondly, there is knowledge concerning A-model assimilation and bridging inferences. For example, in passage (2), the expression "his favorite model" is resolved through the following three separate bridging inference steps:

1. That "his" refers to things pertinent to Peter.
2. That "favorite" is an evaluative predicate translated into a "prefer" attitude box.
3. That cars have models.

For step (1), a focusing mechanism is required to locate the S-box – Peter, since a pronominal expression usually (but not always) refers to some object in the focus. For step (2), the A-model helps to guide the resolution further into the most relevant A-box; in other words, an A-model itself can serve as a marker to find the most relevant A-model earlier in the discourse. For step (3), a basic semantic association occurs to recognize possible semantic relations. In summary, the resolution of the expression "his favorite model," and similarly for all other expressions, is achieved by bridging inferences which synthesize many knowledge sources, including focusing, A-models themselves, and semantic association.

A real world example from BUYER.

In this section, a real world advertisement is presented to demonstrate the application of attitude emergence in processing persuasive discourse. As discussed above, a simple version of A-model construction and assimilation has to be extended to include more general world knowledge. The following ad, which BUYER has processed, demonstrates this.

The Folgers ad.

S1. Is your decaffeinated as dark as ours?
S2. Start with one teaspoon of both.
S3. But just because the amounts are equal doesn't mean the results will be.
S4. Mountain Grown Folgers dark, sparkling Crystals are the difference.
S5. So dark and rich, shouldn't you switch?

In the process of understanding this ad, the system has to figure out many things, for example: Is *S1* a question or a prompt for suggesting an action (for the hearer to perform)? Is *S2*, given the proper interpretation of *S1*, an order or a recommendation; Does *S3* affect the status of the attitude expressed in *S1*? and so on. In order to answer these questions concerning speech intents, the attitude model of each statement has to go through more involved calibration and enhancement than the simple version presented in the previous section. First, the reading of statement *S1* is ambiguous. It could mean that the speaker wants the hearer to inform him as to whether the hearer's decaffeinated is as dark as ours. Or, it can have another reading: the speaker wants the hearer to know whether the hearer's decaffeinated is as dark as ours. The latter reading is inferred by the following world knowledge (WK2)[3]:

> **If** the advertiser already knows everything about his products (which is reasonable to assume),
> **Then** a *question* concerning the product is actually an intention to *inform*.

The two possible attitude models for the two readings of *S1* are depicted in Fig. 3. Note that the two variables *e1* and *e1'* stand for the different events specified by (Inform-whether Hearer ...) and (Know-whether Hearer ...), respectively. The two events are indexed by their event types, as demonstrated by the ovals in the boxes in Fig. 3. Then, when *S2* is processed, from its sentence type (imperative), it is inferred that its attitude context is compatible with both readings of *S1*,

[3]For simplicity and readability, we represent BUYER's knowledge in this and following examples as English-like rules.

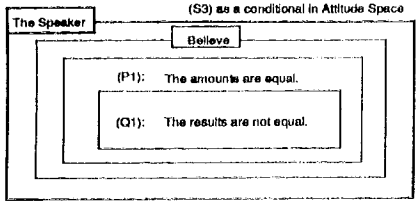

Figure 4: The attitude models of the conditional *S3*.

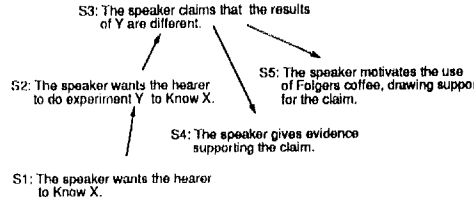

Figure 5: The statements of the Folgers ad in view of their attitude models.

i.e., (Want Speaker (P Hearer ...)). Hence, the semantic association begins to emerge when the events *e2*, specified as "start with one teaspoon of both," and *e1* or *e1'* (see Fig. 3) are being assimilated. The coherence reasoning component of BUYER (Wu and Lytinen 1990) is able to recognize the following Enable coherence relation:

> That the hearer starts with one teaspoon of both kinds of coffee can *enable* that he knows which coffee is better.

That is, the action *e2* suggested in *S2* is an *experiment* to find out something. The choice between *e1* and *e1'* is now clear, due to the following world knowledge (WK3):

> If an agent has a goal to find out something about X,
> Then he can perform an experiment with X.

Since *e1'* – (Know-whether Hearer ...) – is acquired as a goal for the hearer according to Reading 2, *e1'* and hence, Reading 2 is determined as the speech intent. Then, when *S3* is processed, it logically means:

It is not the case that P_1 implies Q_1.

where P_1 stands for – the amounts are equal; Q_1 – the results are equal. The common sense logical reasoning employed in constructing the attitude model of *S3* proceeds as follows: assume P_1 is a fact/observation, should we assume Q_1 or $\neg Q_1$ according PF1? If Q_1 is assumed, then it produces: P_1 implies Q_1, which would be contradictory. So, the only alternative is to choose $\neg Q_1$. The result is then the attitude model shown in Fig. 4. Following attitude percolation, the attitude model of Fig. 4 is pushed down into the one in Fig. 3 for Reading 2, creating the attitude context of (Want Speaker (Believe Hearer ...)). At this point, BUYER is able to reason that P_1 – the amounts are equal – is a fact, by resolving "the amounts" to "the amounts of coffee used in *e2*." Given that the conditional PF1 has a satisfying antecedent, the consequent $\neg Q_1$ (or "the results are not equal") is derived as a fact. Note that, similar to how "the amounts" is resolved, "the results" would be resolved as "the results in the experiments with the two coffees."

Next, when *S4* is processed, by pushing down and resolving "the difference" to "the difference between the results in the experiment," it is recognized that *S4* is supporting the implicit speech intent made in *S3* – the speaker wants the hearer to believe that the results of the experiments are not equal. This is due to the following world knowledge (WK4):

> If the speaker gives the physical cause of a consequence statement,
> Then the consequence statement becomes more *believable*.

Thus, the intended belief that the results are not equal is further enhanced. Finally, *S5*, like *S3*, is not a literal conditional statement. It logically means:

P_2 implies should?($\neg Q_2$)

where P_2 is that Folgers is so dark and rich and Q_2, switch to Folgers. However, the intended meaning is, obviously – switch to Folgers. The derivation hinges on the following "dogma" about "abnormal vs. normal states": (1) Unfamiliar external states may be abnormal; (2) If unknown external states are indeed abnormal, people query about them using "should X?"; (3) Abnormal states are to be corrected. Due to (1) – (3), the intended meaning can be derived, since $\neg Q_2$ is abnormal and to be corrected; in addition, P_2 is also a fact. Fig. 5 summarizes the speech intents reasoned by attitude emergence for the Folgers ad:

Related work

Work on belief percolation ((Wilks and Bien 1979), (Wilks and Bien 1983), (Wilks and Ballim 1987), (Ballim *et al.* 1991)) has strongly inspired our work. However, most of this work concentrates on one single operation of attitude/belief percolation – the "push down" operation. Although this operation is important for investigative reasoning and assimilating attitude models, effects on attitude models themselves due to attitude percolation are more important for our domain of persuasive discourse. Thus, attitude emergence stands as a more relevant mechanism to reason about persuasive speech intent. Moreover, the proposed three-step procedure for computing attitude emergence proves to be a general framework for recognizing speech intents.

The mapping rules proposed in (Hinkelman and Allen 1989), as well as those in (Gerlach and Sprenger 1988), are similar to the attitude model construction rules, while deeper reasoning may underlie some of

their rules, e.g., the interpretation of should?($\neg Q_2$) as Q_2, as we have done to ours. In (Hinkelman and Allen 1989), they also proposed plan-recognition to further identify the speech intent. However, we believe that in persuasive discourse, identifying speech intent through A-models can be done more locally using coherence and bridging inferences. In this sense, our approach is closer to those proposed in (Cohen 1987) and (Mann and Thompson 1988), while (Cohen 1987) considered only support relations and (Mann and Thompson 1988) remained as a descriptive theory, and both did not consider A-models as essential for inferring speech intent.

A-models also are somewhat similar to model-theoretical approaches to semantics. While theories of the more formal kind – e.g, *Discourse Representational Theory* (DRT) (Kamps 1981) – and of the more cognitive – e.g., *Mental Spaces* (Fauconnier 1985) – emphasize the fundamental issues of reference and presupposition, attitude emergence sees application of model reasoning to recognize speech intents. These also demonstrate that mental attitudes serve as only one (though important) way to organize information. There are other ways information should be organized. For example, our formulation of conditionals (for *S3*) is organized not according to attitudes, but principles studied in DRT and mental spaces.

Conclusion and future work

In this paper, a mechanism called attitude emergence is discussed. The basic framework of attitude emergence, which consists of attitude model construction, assimilation, and effects propagation, was first proposed in (Wu and Lytinen 1991) with limited operations. The mechanism is further improved and extended to recognize more indirect speech in persuasive discourse, by adding other common sense and logical reasoning. The generality of attitude emergence is demonstrated by a real world ad which BUYER has processed. BUYER is the computer implementation of attitude emergence and is implemented as a rule-based system which currently has 348 rules organized in 10 problem-solving modules. These problem-solving modules are organized in a way that rules are both forward- and backward-chained, depending the deductive and abductive nature of the rules, respectively, and allows efficient backtracking.

One future work of attitude emergence lies on the further systematization on the *dynamics* of attitude models. Only the force aspect of a statement is present in the current formulation. A fuller formulation of attitude dynamics should include both force and counter force, reflecting the enforcement and the resistance toward an expressed attitude. For example, in Folgers ad, the "counter force" induced by *S2* – the inertia of the hearer not to be told what to do – can be overcomed by the (attracting) force expressed in *S1* – the speaker wants the hearer to know something important – and the force due to the common knowledge that the experiment urged in *S2* can enable the attainment of such knowledge. Relating statements to psychological forces are steps toward explaining the psychological reality of persuasive force and pressure. Along this line, it is found that the work on "force dynamics" proposed in (Talmy 1988) is highly relevant. We are currently looking into the relation between the two.

References

Ballim, A; Wilks, Y.; and Barnden, J. 1991. Belief ascription, metaphor, and intensional identification. *Cognitive Science* 15:133–171.

Cohen, R. 1987. Analyzing the structure of argumentative discourse. *Computational Liguistics* 13:11–23.

Dinsmore, J. 1987. Mental spaces from a functional perspective. *Cognitive Science* 11:1–12.

Fauconnier, G. 1985. *Mental spaces.* The MIT Press, Cambridge, MA.

Gerlach, M. and Sprenger, M. 1988. Semantic interpretation of pragmatic clues: Connective, modal verbs, and indirect speech acts. In *Proceddings of the Eleventh International Conference on Computational Linguistics.* 191–195.

Hinkelman, E. and Allen, J. 1989. Two constraints on speech act ambiguity. In *Proceedings of the Twenty Seventh Annual Meeting of the Association for Computational Linguistics.* 212–219.

Kamps, H. 1981. A theory of truth and semantic representation. In Groenendijk, J.; Janssen, T.; and Stokhof, M., editors 1981, *Formal methods in the study of language.* Mathematish Centrum, Amsterdam. 277–322.

Mann, W. and Thompson, S. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text* 8:243–281.

Talmy, L. 1988. Force dynamics in language and cognition. *Cognitive Science* 12:49–100.

Wilks, Y. and Ballim, A. 1987. Multiple agents and heuristic ascription of belief. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence.* 118–124.

Wilks, Y. and Bien, J. 1979. Speech acts and multiple environments. In *Proceedings of the Sith International Joint Conference on Artificial Intelligence.*

Wilks, Y. and Bien, J. 1983. Beliefs, points of view and multiple environments. *Cognitive Science* 8:120–146.

Wilks, Y. and Hartley, A. 1990. Belief ascription and model generative reasoning: joining two paradigms to a robust parser of messages. In *The 1990 DARPA Workshop.* 219–239.

Wu, H-J. and Lytinen, S. 1990. Coherence reasoning in persuasive discourse. In *Proceedings of the Twelfth Conference of the Cognitive Science Society,* Cambridge, Massachusetts. 503–510.

Wu, H-J. and Lytinen, S. 1991. Attitude and coherence reasoning in persuasive discourse. In *Proceedings of the 1991 AAAI Spring Symposium on Argumentation and Belief,* Stanford, California.