PUTTING AFFECT INTO TEXT

Eduard H. Hovy
Yale University
Computer Science Department [1]

Abstract

How is affect communicated in language? Natural languages contain a large number of techniques for injecting affect into text, both explicitly and implicitly. This paper discusses some techniques that speakers use to slant their text, and describes what is required for a computer program to generate differently slanted versions of a single underlying representation.

1 Introduction

Affect — the speaker's sympathies and antipathies — can be communicated in language both explicitly and implicitly. Since an explicit statement of the speaker's opinion (such as "I don't like your shoes") may alienate the hearer, the speaker should skirt sensitive issues and achieve effects indirectly. Most languages have a large body of techniques for doing so. What are these techniques? This paper briefly describes how the program PAULINE (Planning And Uttering Language In Natural Environments) is able to produce differently slanted texts from a single underlying representation using affect.

PAULINE generates the examples used in this paper from the interpretations produced by the JUDGE program ([Bain 86]), a case-based expert system that models the sentencing behaviour of a judge. As input, JUDGE accepts the representation of a fight — a set of actions and resulting states — and as output to PAULINE it produces a set of interpretations. Each interpretation describes an action, its justifiability, and the culpability of the actor. Saying only the actions, a typical fight is:

```
FIRST, JIM BUMPED MIKE ONCE AND HURT HIM. THEN MIKE SMACKED JIM, HURTING
HIM. NEXT, JIM HIT MIKE ONCE. THE RESULT WAS THAT HE KNOCKED HIM DOWN.
AFTER THAT, MIKE SMACKED JIM SEVERAL TIMES AND KNOCKED HIM DOWN. JIM
SLAPPED MIKE SEVERAL TIMES. THE RESULT WAS THAT HE HURT HIM. AFTER THAT,
MIKE STABBED JIM. AS A RESULT, JIM DIED.
```

Clearly, in any real account of the fight, Jim's version will to differ appreciably from Mike's. Each speaker will make the decisions that slant the text in his favour. PAULINE makes affect-related decisions during each of its three stages: topic collection, topic organization into phrases and sentences, and topic realization into text.

[1]This work was supported in part by the Advanced Research Projects Agency monitored by the Office of Naval Research under contract N00014-82-K-0149.

2 Computing Affect

In order to communicate his opinions implicitly by slanting the text, the speaker must be able to determine which aspects of it he is sympathetic to, which aspects he dislikes, and which he does not care about. PAULINE uses three values of affect: GOOD, BAD, and NEUTRAL. Here, *affect* simply denotes something akin to "like". (Three values are sufficient to give the program interesting behaviour. In this regard it is similar to the narrative summarization work in [Lehnert 82].)

PAULINE gets its affects from two sources: The first source is the user, who lists one or more representation elements as *sympathies* or as *antipathies*. (For example, when PAULINE is to defend Mike, the concept "Mike" is GOOD and the concept "Jim" is BAD.) The second source is the intrinsic affect associated with generic representation types. (For example, in neutral context in the JUDGE domain, the concepts "hit" and "die" are BAD, the concept "unintentionally" is GOOD, and all other concepts, such as "Jim" and "Mike", are NEUTRAL. Similar information is used by the JUDGE program to determine its interpretation of each action.)

In order to determine its opinion about any arbitrary piece of input representation, PAULINE needs the ability to combine its affects with the intrinsic affects and to propagate the results along the relations between concepts. Though their exact form depends on the design of the representation, the **basic rules of affect propagation** are:

1. **affect is preserved when combined with NEUTRAL**
2. **like affects combine to GOOD**
3. **unlike affects combine to BAD**
4. **the combined affect inverts for certain relations between affect-bearing concepts, (e.g., the conceptual *patient* of a BAD act). This requires special rules**

This works as follows: assume the current topic is the action ACT-6:

```
#{ACTION-UNIT  =  ACT-6
      [ACT :  HIT]
      [ACTOR :  MIKE]
      [TO :  JIM]
      [FORESEEABILITY :
            #{STATE  =  STATE-10
                [TYPE :  PHYSICAL-INTEGRITY-VIOLATION]
                [ACTOR :  JIM]
                [DEGREE :  SERIOUS-TEMPORARY]}]
      [NUMBER :  SINGLE]
      [DEGREE :  HARD]
      [INTENTIONALITY :  PRESENT]
      [RESULT :  #{STATE  =  STATE-11
                    [TYPE :  PHYSICAL-INTEGRITY-VIOLATION]
                    [ACTOR :  JIM]
                    [DEGREE :  KNOCK-DOWN]}]}
```

Figure 1

(this is a slightly modified and pruned version of the actual JUDGE representation). Stated neutrally, ACT-6 reads

(1) MIKE INTENTIONALLY HIT JIM HARD ONCE AND KNOCKED HIM DOWN

Assume PAULINE is sympathetic to Mike. In order to determine its opinion of ACT-6, the program combines the intrinsic affect for the type of ACT-6, HIT, (BAD) with its affect for Mike (GOOD) (noting that, since Mike is the ACTOR, the affect doesn't invert), to get the affect BAD. That is to say, in ACT-6, Mike looks bad and, when speaking, the action should be omitted if possible or mitigated if not. One of the ways PAULINE could say this is:

(2) MIKE JUST TAPPED JIM ONCE

If, on the other hand, PAULINE's sympathies are for Jim, it combines the affect for HIT (BAD) with its affect for Jim (GOOD), giving BAD. Since Jim fills the role TO (the conceptual role *patient*), this result must be inverted, and so the final affect is GOOD. That is to say, ACT-6 is GOOD for the case against Mike. In this case PAULINE should enhance the topic. Furthermore, the RESULTant state, STATE-11, carries the same affect, because Jim (GOOD) suffers a PHYSICAL-INTEGRITY-VIOLATION (intrinsically BAD). This result was intentionally (INTENTIONALITY PRESENT) caused by Mike (BAD). These three affects GOOD, BAD, and BAD combine to produce GOOD for Jim in STATE-11, causing it to be enhanced too. Thus, when defending Jim, PAULINE produces sentences such as

(3) MIKE PURPOSELY SMASHED JIM AND KNOCKED HIM DOWN

3 Generating with Affect

3.1 Topic Organization into Phrases

Before it says anything, the generator performs a number of planning tasks. One task is to cast sentence topics together into phrases. In addition to the planning criteria used by other generators (such as focus in [McKeown 82] and hearer knowledge in [Cohen 78] and [Appelt 81]), PAULINE uses affect to control the juxtaposition of sentence topics, since certain phrases are very useful for conveying affect implicitly. The following phrases can be called *enhancers*:

(a) "*Not only* did Pete play the game, *but* he hit five home runs"
(b) "Pete played the game; *what's more*, he hit five home runs"

Clearly, these phrase forms imply that Pete's playing and his home runs carry the same affective value (either both GOOD or both BAD), and in fact that the value is to be strengthened due to their juxtaposition. In contrast, the following sentences carry no such cumulative affective import:

(c) "Pete played the game and he hit five home runs"
(d) "When Pete played the game he hit five home runs"

When an enhancing phrase juxtaposes two affect-laden sentences, the affect is strengthened; when it juxtaposes an affect-laden sentence with a neutral one, the affect is imputed to the latter. Thus, in addition to stressing affective concepts, a speaker can strengthen his case by imputing affect to neutral concepts too! This is, for example, what PAULINE does to produce

(4) NOT ONLY DID JIM EXPECT NO THREAT FROM MIKE ANY LONGER, BUT HE COULD FORESEE THAT HE WOULD HURT HIM IF HE REALLY SLAPPED HIM

when defending Mike. Here Jim's not perceiving a threat from Mike is simply NEUTRAL, but his ability to foresee the BAD result of his action, coupled with the fact that he did it anyway, is BAD for him. However, when juxtaposed in this way, *both* sentences seem BAD for Jim — exactly what PAULINE wants.

Similarly, *mitigators* are phrases with weakening effect. When a mitigating phrase juxtaposes two sentences carrying opposite affect, the resulting affect is that of the first sentence, weakened; when it juxtaposes an affect-laden sentence with a neutral one, the opposite affect is imputed to the latter. In the following sentences, if "John whipped the dog" carries BAD affect, then, if we know nothing more, "he remembered the cat" becomes GOOD:

(e) "*Although* John remembered the cat, he whipped the dog"
(f) "John remembered the cat. *However*, he whipped the dog"

(The first part could just as well have been made BAD by using an enhancer: "*Not only* did John remember the cat, *but* he whipped the dog".)

A number of constraints must be met before two topics can be juxtaposed in an enhancer or mitigator phrase:

1. **Two-predicate enhancer and mitigator phrases can only be used when the parts carry consistent affects; that is,**
   - **in enhancer phrases the two predicates must carry like affect**
   - **in mitigator phrases the two predicates must carry unlike affect**
   - **or else one predicate must be NEUTRAL**
2. **Two-predicate enhancer and mitigator phrases can only be used when the topics in both parts focus on the same concept**
3. **The predicates in two-predicate enhancer and mitigator phrases should match in as many aspects as possible**

3.2 Realization

Sentence subject selection: Since the sentence subject is a prominent position, it must be chosen with care. Grosz [Grosz 77], Sidner [Sidner 79], and McKeown [McKeown 82], among

others, describe rules for choosing subjects in order to produce flowing text. These rules are often underspecific; McKeown's algorithm simply picks the default (a predefined entry for each predicate) when a number of focus candidates exist with the same number of implicit links to the potential focus list. Affect can be used as an additional criterion for subject choice — either at a low level, simply to help winnow out candidates, or at a high level, to help slant the text very strongly. When PAULINE has the goal to convince the hearer or to make known its affects, it use, in addition, the following rule:

> 4. **The new focus should be selected from GOOD candidates for sentences with GOOD affect and from BAD candidates for sentences with BAD affect**

For example, when PAULINE is defending Jim, the following sentence topic is BAD for him; by this rule, Mike, an antipathy, is chosen as subject, thereby de-emphasizing Jim's role:

(5) MIKE ONLY WAS BUMPED BY JIM ONCE

Sentence Inclusion: Adverbs: Some adverbial stress words specifically function as enhancers or mitigators; in the JUDGE domain, PAULINE uses the following:

intentionality: "intentionally", "purposely" — "accidentally"
degree: "hard" — "lightly" (hit)
number: "repeatedly" — "once" (stabbed)
stress: "really" — "only", "just"

When these words are used to modify concepts that do not already carry affect, they seem strange, forcing the hearer to postulate affect; an additional constraint is:

> 5. **Adverbial stress words can only be used to enhance or mitigate expressions that carry some affect already**

During the realization of a sentence, the program collects the aspects of the topic it can legitimately say as adverbs (it cannot, for example, misrepresent the topic to say "lightly" when the DEGREE is HARD). From these, it selects at most two (the two least recently said types), since when affective adverbs are overused the effect is unnatural, to give, for example:

(6) MIKE *JUST* HIT JIM *ONCE*
(7) JIM COULD FORESEE THAT HE WOULD HURT MIKE IF HE *REALLY* SLAPPED HIM. HE SLAPPED HIM *REPEATEDLY* AND HURT HIM

Word Choice: Verbs: Often an action can be described by a number of verbs of which most carry some affect. For example, PAULINE accesses "hit" from the representation element HIT and then uses its affective values to discriminate to "tap" as a mitigator and "smash" as an enhancer (the lexicon contains discrimination nets, as in [Goldman 75]):

(8) JIM JUST *TAPPED* THAT JERK MIKE ONCE
(9) JIM PURPOSELY *SMASHED* MIKE AND KNOCKED HIM DOWN

Table 1: Defending Mike

FIRST, *THAT JERK* JIM BUMPED MIKE *HARD* AND HURT HIM. THEN MIKE *JUST TAPPED* JIM *ONCE.*	enhancing noun, stress mitigating adverbs, verb
AFTER THAT, JIM DID NOT EXPECT THAT MIKE WAS GOING TO HURT HIM ANY LONGER; *ALSO;* JIM COULD FORESEE THAT HE WOULD INJURE MIKE IF HE PURPOSELY STRUCK HIM. HE *REALLY SMASHED* HIM. THE RESULT WAS THAT HE INJURED HIM.	NEUTRAL and BAD in enhancer phrase enhancing stress enhancing verb
NEXT, MIKE HIT JIM, KNOCKING HIM DOWN.	suppress all but action
NOT ONLY DID JIM EXPECT NO THREAT FROM MIKE ANY LONGER, *BUT* HE COULD FORESEE THAT HE WOULD HURT HIM IF HE REALLY SLAPPED HIM. HE SLAPPED HIM *REPEATEDLY* AND HURT HIM.	enhancer phrase enhancing adverb
MIKE REALIZED THAT JIM *REALLY* HURT HIM; *ALSO*, HE EXPECTED THAT JIM'S GOAL WAS TO HURT HIM. JIM *WAS STABBED.*	enhancing stress, phrase mitigating passive
AS A RESULT, JIM *ONLY* DIED.	mitigating (!) stress

Table 2: Defending Jim

JIM COULD NOT EXPECT THAT MIKE WOULD BE HURT IF JIM *ACCIDENTALLY* BUMPED HIM; ALSO, A REASONABLE PERSON COULD NOT FORESEE THAT IF HE BUMPED HIM HE WOULD HURT HIM. HE HAD NO INTENTION TO BOTHER MIKE. MIKE *ONLY WAS BUMPED* BY JIM ONCE. THE ACTION WAS AN ACCIDENT.	mitigating adverb mitigating stress mitigating passive
THEN MIKE REALIZED THAT JIM HURT HIM. *IN ADDITION*, MIKE DID NOT EXPECT THAT JIM WAS GOING TO HURT HIM ANY LONGER. MIKE'S GOAL WAS TO INJURE JIM. MIKE COULD FORESEE THAT HE WOULD INJURE HIM IF HE PURPOSELY HIT HIM ONCE. HE HIT HIM. THE RESULT WAS THAT HE INJURED HIM. HE REQUIRED JUSTIFICATION FOR CAUSING HIM TO BE INJURED. THE ACTION WAS AN ESCALATED RETALIATION.	enhancer phrase include all topics that are BAD for Mike...
NEXT, JIM REALIZED THAT MIKE INJURED HIM. JIM *JUST TAPPED* MIKE ONCE. THE ACTION WAS A SIMPLE RETALIATION.	mitigating stress mitigating verb
...	etc...
AS A RESULT, JIM DIED.	

4 Conclusion

In summary, consider the generation of a complete fight, from the set of interpretations of each action PAULINE gets from JUDGE. In table 1, PAULINE's sympathies are for Mike, so that it stresses Jim's part of the fight; in table 2, from the same input, PAULINE defends Jim.

5 Acknowledgements

Thanks to Bill Bain for JUDGE, and to Jeff Grossman and Ashwin Ram for lots of hard work on PAULINE.

6 Bibliography

1. Appelt, D.E.,
Planning Natural Language Utterances to Satisfy Multiple Goals, Ph.D. dissertation, Stanford, 1981.

2. Bain, W.M.,
Case-based Reasoning: A Computer Model of Subjective Assessment,
Ph.D. dissertation, Yale University, forthcoming 1986.

3. Cohen, P.R.,
On Knowing What to Say: Planning Speech Acts, Ph.D. dissertation, University of Toronto, 1978.

4. Grosz, B.J.,
The Representation and Use of Focus in Dialogue Understanding, Stanford Research Institute Technical Report no. 151, 1977.

5. Goldman, N.M.,
Conceptual Generation, in **Conceptual Information Processing**, Schank, R.C. (ed), North-Holland Publishing Company, 1975.

6. Lehnert, W.G.,
Plot Units: A Narrative Summarization Strategy, in **Strategies for Natural Language Processing**, Lehnert, W.G. & Ringle, M.H. (eds), Lawrence Erlbaum Associates, 1982.

7. McKeown, K.R.,
Generating Natural Language Text in Response to Questions about Database Queries, Ph.D. dissertation, University Of Pennsylvania, 1982.

8. Sidner, C.L.,
Toward a Computational Theory of Definite Anaphora Comprehension in English, Massachusetts Institute of Technology Technical Report AI-TR-537, 1979.